City Research Online

City, University of London Institutional Repository

A study of lawyers' information behaviour leading to the development of two methods for evaluating electronic resources

Public archive image

Stephann Makri

A thesis submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy

UCL

2008

Declaration of originality

I, Stephann Makri, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

In this thesis we examine the information behaviour displayed by a broad cross-section of academic and practicing lawyers and feed our findings into the development of the Information Behaviour (IB) methods - two novel methods for evaluating the *functionality* and *usability* of electronic resources.

We captured lawyers' information behaviour by conducting naturalistic observations, where we asked participants to think aloud whilst using existing resources to 'find information required for their work.' Lawyers' information behaviours closely matched those observed in other disciplines by Ellis and others, serving to validate Ellis's existing model in the legal domain. Our findings also extend Ellis's model to include behaviours pertinent to legal information-seeking, broaden the scope of the model to cover information *use* (in addition to information-seeking) behaviours and enhance the potential analytical detail of the model through the identification of a range of behavioural 'subtypes' and levels at which behaviours can operate.

The identified behaviours were used as the basis for developing two methods for evaluating electronic resources – the IB *functionality* method (which mainly involves examining whether and how information behaviours are currently, or might in future be, supported by an electronic resource) and the IB *usability* method (which involves setting users behaviour-focused tasks, asking them to think aloud whilst performing the tasks, and identifying usability issues from the think-aloud data).

Finally the IB methods were themselves evaluated by stakeholders working for LexisNexis Butterworths – a large electronic legal resource development firm. Stakeholders were recorded using the methods and focus group and questionnaire data was collected, with the aim of ascertaining how usable, useful and learnable they considered the methods to be and how likely they would be to use them in future. Overall, findings were positive regarding both methods and useful suggestions for improving the methods were made.

Contents

List of figures

List of tables

List of appendices

Publications based on work in this thesis

Makri, S. (2006). Studying Law Students' Information-Seeking Behaviour to Inform the Design of Digital Law Libraries. Presented as a Poster Presentation at the 10th ECDL Conference, September 17-22, Alicante, Spain.

Makri, S. (2007). Studying Academic Lawyers' Information-Seeking to Inform the Design of Digital Law Libraries. In IEEE Computer Society Bulletin of the Technical Committee on Digital Libraries 3(3). Available online: http://www.ieee-tcdl.org/Bulletin/current/makri/makri.html

Makri, S., Blandford, A. & Cox, A.L. (2006). A Study of Legal Information-Seeking Behaviour to Inform the Design of Electronic Legal Research Tools. In Blandford A. and Gow, J. (Eds.). Proceedings of the 1st International Workshop on Digital Libraries in the Context of Users' Broader Activities (DL-CUBA), 15th June 2006. pp. 33-36. JCDL 2006, Chapel Hill, NC, USA.

Makri, S., Blandford, A. & Cox, A.L. (2007). 'I'll just Google it!': Should Lawyers' Perceptions of Google Inform the Design of Electronic Legal Resources? In proceedings of the Web Information-Seeking and Interaction (WISI) Workshop. SIGIR 2007. 27 July 2007, pp. 5-8. Amsterdam, The Netherlands.

Makri, S., Blandford, A. & Cox, A.L. (2008). Investigating the Information-Seeking Behaviour of Academic Lawyers: From Ellis's Model to Design. Information Processing and Management 44(2), pp. 613-634.

Makri, S., Blandford, A. & Cox, A.L. (In Press). Using Information Behaviours to Evaluate the Functionality and Usability of Electronic Resources: From Ellis's Model to Evaluation. To appear in the Journal of the American Society for Information Science and Technology.

Makri, S., Blandford, A. & Cox, A.L. (In Preparation). Trial by Fire: A Case Study on the Development of Lawyers' Information Expertise. In preparation for submission to Information Processing and Management.

Acknowledgements

Firstly, I would like to thank my supervisors, Ann Blandford and Anna L. Cox for their valuable guidance. Next, I would like to thank the academic lawyers, practicing lawyers and law librarians that participated in my empirical study and the pilot participants and staff at LexisNexis Butterworths for their assistance with the development and evaluation of the Information Behaviour methods. I would also like to thank Simon Attfield for running the IB method focus group sessions.

Finally, I would like to thank my family and Gayle for their support and Sanjay for choosing a career path which seems to have paralleled my academic path.

This thesis is dedicated to the memory of Maria Hambi, who always wanted me to become a Doctor (albeit of a different kind).

Chapter 1: Introduction

This chapter at a glance…

In this chapter we:

- Present our motivation for the research on lawyers' information behaviour described in this thesis.
- Summarise our thesis and contributions to research.
- Explain the structure of the remaining chapters in the thesis.

1.1   Research motivation

This work is motivated by the human-centred ethos that in order to design interactive systems that truly support information work, it is first necessary to gain a detailed understanding of the work that these systems might be designed to support. In this vein, we examine the information behaviour displayed by lawyers when using a range of existing electronic legal resources to find information required for their work. Whilst our findings could have been used to *directly* inform the design of new or existing electronic legal resources (for example through making and implementing design suggestions), we recognised a 'gap in the market' for *methods* to support the evaluation of this type of resource. Therefore, based on a similar user-centred ethos, we chose to feed our findings into the development of two novel methods – one for evaluating the *functionality*, the other the *usability* of electronic legal resources. By doing so, our work *indirectly* informs the design of electronic legal resources, by facilitating their evaluation and subsequent improvement. In the remainder of this section, we discuss why information is important for lawyers and why they are an important user group to study.

Law is a highly knowledge-intensive domain and obtaining accurate and up-to-date legal information can mean the difference between winning or losing cases. The information work carried out by lawyers can be complex, often involving finding and working with a wealth of different types of information. This 'wealth' of legal information spans different types of documents (e.g. law reports/legal cases, legislation, commentary articles, forms and precedents etc.), a wide range of legal topic areas and a range of jurisdictions (i.e. geographical areas to which the law applies).

Nowadays, much of this information is obtained from electronic legal resources, ranging from general and law-specific internet search engines (largely unorganised indexes of web pages), to digital libraries (large, *organised repositories* of information and knowledge - Lynch and Garcia-Molina, 1996), to legal citator tools (another type of index, this time of the current or historical status of legal documents). The need to obtain a wide range of (often complex) information requires lawyers, who usually do not have a strong information background, to make effective use of these resources. This, however, is not always easy. Most digital law libraries (including LexisNexis Butterworths and Westlaw, which are two that are frequently used by lawyers) contain materials covering hundreds of years of law. As these resources contain so much information, of different types, they must provide sophisticated tools to help lawyers find the information within.

The screenshots in figure 1 illustrate two such tools available in the Justis digital law library, which specialises in providing access to mostly UK, Irish and European case law. The left-hand image illustrates the facility to search for legal case reports through the use of various segmented field searches. The right-hand image illustrates the facility to browse for legal journal articles by source and year of publication.
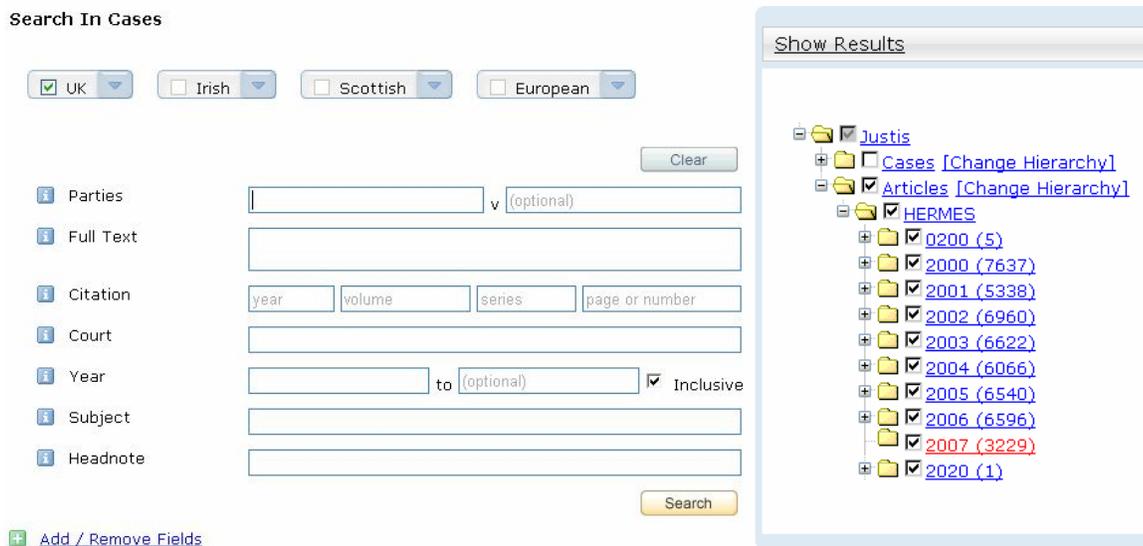
Figure 1: Partial screenshots of the Justis electronic legal resource illustrating the facility to search for legal case reports through the use of various segmented field searches (left) and the facility to browse for legal journal articles by source and year of publication (right). Screenshots included with permission from Justis UK management.

Whilst the provision of sophisticated tools such as these aims to support lawyers' information-seeking, it also has the potential of loading these resources with functionality, making them complicated and therefore difficult for lawyers to use. This is an interesting paradox as it illustrates the trade-off between providing simplicity and functionality in electronic legal resources. This paradox is also the main motivation for our work. The complexity of many electronic legal resources (coupled with the importance of information work for lawyers) makes law a good example of a domain with the potential to benefit substantially from improvements in electronic resource design.

The traditional view in legal research education is that lawyers themselves (or shortcomings in their electronic resource training) are to blame for any information-seeking difficulties that they face. Howland and Lewis (1990) interviewed a number of law firm librarians about recent law graduates' competency in the use of LexisNexis and Westlaw. The general perception was that summer clerks and first-year associates were not efficient or cost-effective users of electronic resources, in spite of having some training using the resources. One librarian commented: *"It is not a lack of intelligence that is the problem but rather an unwillingness to learn efficient searching – it's amazing that while most associates are so careful about most aspects of their work, they are so sloppy about computer searching"* (p.387). However, as researchers in the field of HCI, it is not our aim to 'design users' to operate in a way that more closely matches the way that the system intends them to interact with it and we do not regard this as solely an information literacy or training issue. Instead, we are motivated by the prospect of *improving* electronic legal resources by ensuring that they are designed to fit users' needs. This approach is particularly important for electronic *legal* resources as, according to Sutton (1994), *"both Lexis and Westlaw were designed with no apparent attention being paid to the information-seeking behaviour of attorneys"* (p.198). As we stated at the outset, we believe that it is only truly possible to ensure that electronic resources are designed to fit users' needs if we gain a thorough understanding of our users and how they look for information in the context of their work.

There are, however, few existing comprehensive studies of how lawyers look for and use information and even fewer studies aimed at informing the design of electronic legal resources (whether directly, by using the findings to make suggestions for the design or re-design of these resources, or indirectly, by feeding the findings into tools or techniques that can be used to support design or re-design). This is a traditional Human-Computer Interaction (HCI) approach for user-centred design – an approach aimed at understanding how people use interactive systems and feeding that understanding into the design of these systems. This is an approach which

we follow in this thesis. More specifically, we conduct a series of naturalistic observations of academic and practicing lawyers using a variety of electronic legal resources to find information required for their work. We then use the findings from this study as a basis for developing the Information Behaviour (IB) methods - two novel methods for evaluating the functionality and usability of electronic legal resources.

In the remainder of this chapter, we summarise our thesis and the contributions to research that it makes, before providing an overview of the remaining chapters and an indication of how they relate to one another.

1.2   Thesis summary and research contributions

Our thesis makes two important contributions to research. The first of our contributions is the theoretical and practical findings from an empirical study into academic and practicing lawyers' information behaviour. This study involved conducting naturalistic observations, where the academic and practicing lawyers were asked to think aloud whilst using existing electronic legal resources. The observation process, combined with probing questions (before, during and after the observation) provided an insight into their information behaviour (i.e. *what* they were doing and *why* when using the existing resources). The interviews/observations were transcribed and analysed using an approach based on the open and axial coding elements of Grounded Theory, not with the aim of generating theory per se, but with the aim of identifying behaviours that might inform the design and evaluation of electronic legal resources. It was found that both academic and practicing lawyers' information behaviour closely matched behaviours observed in other disciplines by Ellis and other researchers (see Ellis, 1989; Ellis, Cox and Hall, 1993; Ellis and Haugan, 1997; Meho and Tibbo, 2003).

Our empirical study led to the validation of Ellis's (1989) behavioural model of information-seeking in the new academic domain of law and outside of an academic setting in the professional domain of law. Our study also led to the extension of Ellis's model to include behaviours pertinent to legal information-seeking, a broadening of the scope of Ellis's model through the coverage of information use (as opposed to information-seeking) behaviours and the enhancement of the potential analytical detail of Ellis's model through the identification of mutually-exclusive pairs of behavioural subtypes and different levels at which many behaviours can operate.

The behaviours that were identified were then used as the basis of two novel methods for evaluating electronic resources - the Information Behaviour (IB) functionality and IB usability methods. The development and formative evaluation of the IB methods is the second contribution made by this thesis. The IB *functionality* method involves evaluators examining whether and how information behaviours are supported (or might be supported in future) by a particular electronic resource, at a number of different levels, and considering whether it is still necessary to support all of the behaviours/levels that the resource currently supports. The IB *usability* method involves setting behaviour-focused tasks to intended or actual users of the resource, asking them to think aloud whilst performing the tasks, and identifying usability issues from the resultant think-aloud data.

Both the IB functionality and usability methods were evaluated by a small team of stakeholders (including some usability experts) working for LexisNexis Butterworths, a large electronic legal resource development firm. The stakeholders were given a one-day tutorial on how to use the methods and given the opportunity to practice using the methods by evaluating the functionality and usability of two electronic legal resources that they had developed. After using the methods, the stakeholders answered short questionnaires and participated in two short focus group sessions, where they were asked questions surrounding how usable, useful and easy to learn they deemed the methods to be and how likely they were to use the methods in future. The output from the evaluations was also collected and analysed. Overall, findings were positive regarding all of the above success measures and for both the IB functionality and the IB usability methods. In addition, tutorial attendees made useful suggestions for how the methods might be improved.

To summarise, the theoretical and practical contributions of this thesis are:

1. An empirical study into academic and practicing lawyers' information behaviour leading to:

a. The validation of Ellis's behavioural model of information-seeking in the new academic domain of law (through observations of academic lawyers looking for information) and outside an academic setting in the professional domain of law (through observations of practicing lawyers looking for information).
b. The extension of Ellis's model to include behaviours pertinent to legal information-seeking.
c. Broadening of the scope of Ellis's model through the coverage of information use (in addition to information-seeking) behaviours.
d. Enhancement of the potential analytical detail of Ellis's model through the identification of mutually-exclusive pairs of behavioural subtypes and different levels at which many behaviours can operate (these concepts are explained in chapter 5).

2. The development and formative evaluation of the Information Behaviour (IB) methods - two novel methods for evaluating the functionality and usability of electronic legal resources (theoretically underpinned by the information behaviours identified in our empirical study).

1.3   Overview of remaining chapters

As a broad overview, our thesis begins by discussing the literature that motivates our empirical work. We then present our methodology and findings and discuss them in relation to previous work. Next, we discuss the development and evaluation of the IB methods before concluding by summarising the thesis and examining its impact and the potential for future work in related areas.

We now provide a more detailed overview of our thesis. Chapters 2 and 3 review relevant literature. In chapter 2, we review the literature on lawyers' information behaviour and their attitudes towards and use of electronic legal resources. In this chapter, we review three sets of studies: user-centred studies on lawyers' attitudes towards and use of electronic legal resources, studies of lawyers' information behaviour and studies on lawyers' information use and re-use behaviour. In doing so, we highlight the research gap that our thesis aims to fill. This chapter also includes a brief introduction to the system of law in England and Wales (known as English Common Law) and to the range of legal information sources and electronic legal resources available to lawyers. In chapter 3, we examine the potential for using information-seeking and behaviour models to inform electronic resource design and evaluation. This involves a review of selected models from the Information Science domain and a discussion on how useful each model is for providing a theoretical lens on lawyers' information behaviour *that will yield data capable of informing the design and evaluation of electronic resources.* We suggest that Ellis's behavioural model of information-seeking is the most suitable choice for this purpose, mainly due to the fact that the model is based on concrete and observable behaviours as opposed to cognitive or physical processes that demand a further level of filtration to the underlying behaviour level so that they may truly inform design and evaluation.

In chapter 4, we present the methodology of our naturalistic study of academic and practicing lawyers' information behaviour. This includes discussion of our data collection and analysis approach and a discussion of the theoretical influences on our approach.

In chapter 5, we present our findings on lawyer's information behaviour and discuss them in relation to previous information work, particularly work by David Ellis and his colleagues.

Next, in chapter 6, we use the findings from our naturalistic study to inform the development of the Information Behaviour (IB) methods – two novel methods for evaluating the functionality and usability of electronic legal resources. In this chapter, we present both methods, along with worked examples and discuss the benefits and limitations of using the methods. We also discuss the development and early testing of the methods, which includes an account of pilot sessions with three lawyers and an electronic resource developer and the insights gained as a result. In chapter 7, we present the methodology and findings of a study aimed at formatively evaluating the IB methods with a group of stakeholders working for LexisNexis Butterworths – a large electronic

legal resource development firm. The aim of this study was to find out how useful, usable and easy to learn the stakeholders deemed the methods to be and how likely it was that they would use the methods in future.

Finally, in chapter 8, we summarise the thesis and re-visit the research contributions listed in this chapter. We also discuss the implications of our thesis work for shaping future studies of information behaviour and for informing the design and evaluation of electronic resources. This is followed by a discussion of the potential for future work in the areas of legal information behaviour, behavioural research evaluation and information behaviour in general. Figure 2 summarises the structure of the remainder of this thesis.
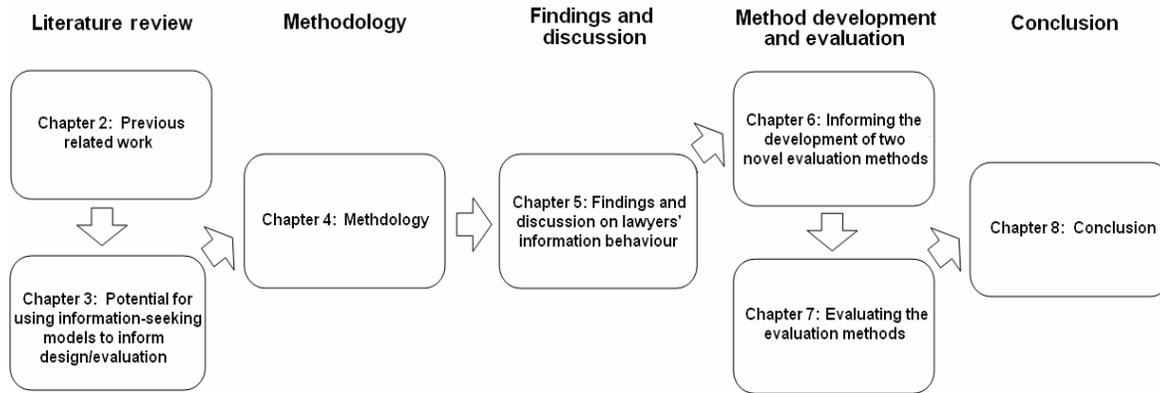
Figure 2: Flow diagram to summarise the structure of this thesis.

Chapter 2: Previous related work

This chapter at a glance…

In this chapter we:

- Provide an introduction to English law and to electronic legal resources.
- Highlight the gap in research on lawyers' information behaviour that our thesis aims to fill.
- Examine previous work on lawyers' attitudes towards and use of electronic legal resources to further motivate our work.

2.1   Overview

The work that currently exists surrounding the information behaviour of lawyers and their use of electronic legal resources is scattered amongst a variety of domains. These include the domains of Artificial Intelligence and Law (where studies have been primarily concerned with designing electronic systems to support legal reasoning and argumentation), Information Retrieval (where, as Sutton, 1994 highlights, studies have been primarily concerned with measuring the 'performance' of digital law libraries such as LexisNexis and Westlaw using relevance assessment based on test sets of static authority), Information Science (where studies have been primarily concerned with understanding lawyers' information behaviour) and Human-Computer Interaction (where studies have been primarily concerned with designing interactive systems to support legal work). In this chapter we focus on select literature from the domains of Information Science and Human-Computer Interaction, which are the domains most closely matched to our overall aim of understanding lawyers' electronic information behaviour and feeding this understanding into the design and evaluation of electronic legal resources.

Before we begin the review of the literature, we provide a brief introduction to the system of law in England and Wales (known as English Common Law) and to the range of legal information sources and electronic legal resources available to lawyers. This grounding is not only useful for understanding aspects of lawyers' information behaviour as presented in this chapter as part of our literature review, but also for understanding elements of the information-seeking episodes of lawyers that took part in our empirical study as presented in chapter 5 of this thesis. After this introduction, we examine the previous work on lawyers' information behaviour and their attitudes towards and use of electronic legal resources. We begin, in section 2.3.1 by examining a series of user-centred studies on lawyers' attitudes towards and use of electronic legal resources. Next, we review of a number of studies of lawyers' information behaviour, followed by studies on lawyers' information use and re-use behaviour.

2.2   Introduction to English Common Law and to legal sources and resources

As summarised by Cheatle (1992), England and Wales follow the 'Common Law' legal system (also referred to as case law and judge-made law), where much law remains unwritten but remains based on legal rules laid down in court by the judiciary. Courts administer the law by interpreting legislation and are bound by the doctrine of judicial precedent, which compels them to follow decisions of higher courts. This helps to explain why legal information-seeking is often a highly involved task. Finding the answer to a legal question may not be as simple as consulting an authoritative version of the law (e.g. legislation such an Act of Parliament), but instead might involve examining a range of case law in order to form an opinion on where the law currently stands.

Legal sources can be divided into primary and secondary sources, as summarised by Andrews (1993). Primary sources are authoritative records of the law made by law making authorities. Secondary sources pertain to the law, but are not authoritative records of the law (i.e. they are not official texts). There are three primary sources of UK law: parliament, the courts and the European Community. Acts of Parliament, known as Statutes, are a powerful information source and take precedence over other legal sources. Parliament often delegates legislative powers to other bodies (such as local councils or government departments). These powers are set down in the form of Legal Rules – a form of secondary legislation, usually published as Statutory Instruments. In addition, parliament confers ministerial powers which are used in the regulation of public bodies (published as Directions, Guidance, Circulars and Codes of Practice). Acts of Parliament are published by the Office of Public Sector Information (OPSI), which also publishes Statutes in Force, which aims to provide a comprehensive subject-based list of all statutory material which is currently in force. Important secondary sources for lawyers include textbooks, legal journals (which include a variety of both practical and academic articles) and commentary materials (which summarise the law related to particular legal areas). A popular commentary source, mentioned several times in this thesis is 'Halbury's Laws of England and Wales.' A fundamental grounding in the types of legal sources available to lawyers is likely to be useful in order to understand the information behaviour displayed by lawyers in chapter 5.

It is also useful to have a basic understanding of the electronic legal resources available to lawyers. LexisNexis Butterworths and Westlaw are two of the largest and best known digital law libraries. Although both of these resources contain a wide type of legal sources as described above, as highlighted by Norman (2004), LexisNexis and Westlaw are both owned by multinational firms and this has resulted in the withdrawal of rights to each others' sources. Access rights to sources has also led to several subtle differences between the content carried by both electronic resources, for example LexisNexis Butterworths has access to a wider range of unreported UK cases than Westlaw. In general, however, both of these resources have similar coverage (in terms of both age of material, the jurisdiction or geographical area that the material applies to, the topical area of the material and the type of material as described previously). New versions of both of these resources were released to the public during the period that our work was carried out. In addition, the LexisNexis group have recently started to phase out LexisNexis Professional (the predecessor to the LexisNexis Butterworths electronic resource). However, at the time of writing, both were available to academia and practice as separate commercial products. There are also a number of smaller digital law libraries, such as Justis and Lawtel, that provide access to a range of legal materials (particularly UK and EU cases).

A wide variety of other electronic legal resources are also available to lawyers, often for specialised purposes. For example, the HeinOnline digital library provides access to older legal journal articles that are unlikely to be available on LexisNexis Butterworths or Westlaw, whilst other digital libraries cater for particular areas of law (e.g. Kluwer Arbitration). Indexing and citator services such as Current Legal Information (CLI), cater for other niches. For example, apart from containing case digests from all reported cases from 1947 onwards, CLI includes a comprehensive case and legislation citator to make it easier for lawyers to gain a current or historical understanding of the law. Although such a wide range of electronic legal resources exist to support legal work, few user-centred studies have been carried out to examine lawyers' use of these resources. In the next section, we review the few existing studies in the area.

2.3   Previous work on lawyers' information behaviour and their attitudes towards and use of electronic legal resources

In this section, we review previous work on lawyers' attitudes towards and use of electronic legal resources and highlight a gap in research in this area. We begin by examining a series of user-centred studies on lawyers' use and perception of electronic legal resources. These studies include a survey by Elliott and Kling (1996) into electronic resource usage, a log analysis by Yuan (1997) of law students' searches using LexisNexis's Quicklaw electronic resource and studies by Andrews (1993) and Oulanov and Pajarillo (2003) on lawyers' and law librarians' perceptions of the LexisNexis electronic resource. We then take a slightly broader focus and review

the literature on lawyers' electronic information-seeking behaviour (i.e. literature on how lawyers use electronic legal resources to find information) and their electronic information use and re-use behaviour (i.e. literature on how lawyers use and re-use the electronic information that they find).

2.3.1   User-centred studies on lawyers' attitudes towards and use of electronic legal resources

Many existing studies that have involved lawyers using electronic legal resources have been system rather than user-centred – focusing on quantitative search performance factors such as completion time, accuracy and efficiency (see Dempsey et al., 2000 as an example). However, there are a small number of user-centred studies on lawyers' attitudes towards and use of electronic legal resources that motivate our work (which is also user-centred). We discuss these studies below.

Vollaro and Hawkins (1986) conducted a series of interviews with patent attorneys working at the AT&T Bell Laboratories. The interviews focused at discovering what types of attorneys might choose to conduct their own searches for legal information and which might turn to a librarian intermediary and whether those who conducted their own searches were satisfied with the results. They found that attorneys who chose to conduct their own searches had longer average search times than intermediaries, despite becoming frequent users of online databases. It was found that patent attorneys only sought to attain a basic level of search competence and, once this had been achieved, were satisfied with their results. This was despite the fact that these attorneys had only familiarised themselves with a few information sources within the database and therefore were potentially (and perhaps unknowingly) missing out on results that might have been relevant by only grasping the basics of searching the system. Vollaro and Hawkins (1986) explain this attitude by the fact that *"from the point of view of the attorney or other end-user, information is a means to an end..."* whilst *"information professionals may value the information as an end in itself and therefore may be more zealous in using many sources to fill a request"* (p.69).

It was also found that nearly all attorneys mentioned difficulty in finding appropriate search terms and in remembering the special features of each database, especially when use was infrequent. Other problems encountered were not knowing when all possible avenues had been pursued and forgetting commands. In addition, little mention was made by participants of workarounds to overcome information-seeking problems, nor was there any mention of special features of the database, such as British spellings in INSPEC (one of the databases used). Vollaro and Hawkins therefore concluded that although this group of patent attorneys could search successfully, even with only limited training, they had not progressed to a point of expertise to consider important details such as workarounds and advanced database features.

In another user-centred study, Yuan (1997) monitored the LexisNexis Quicklaw searches of a group of law students over the period of a year. Yuan examined several aspects of their searching behaviour, including the increase of their command and feature repertoires, their change in language usage, increase of search speed and change of learning approaches. Yuan found that experience did not result in searchers making fewer errors or being able to recover from more errors. Yuan also found that although participants with higher levels of Quicklaw experience used a greater variety of commands and features than those with lower levels of experience, some commands remained rarely or never used. Despite this, however, participants were able to accomplish many tasks by knowing a core set of commands and features. As a result of this finding, Yuan suggests that system designers should consider which system features are made explicit to users, which are hidden and how defaults are set. Yuan also suggests the need for improved interface design which provides explicit information of what functionality the system provides (along with better user documentation and online help).

In order to identify the 'fit' of digital law libraries to various organisations, Elliott and Kling (1996, 1997) conducted a qualitative study in the early 1990s on digital library usage. This was at a time when electronic legal resource use was becoming increasingly important. The authors interviewed forty-six legal professionals (including judges, District Attorneys, Public Defenders and Criminal Defence Attorneys) based in three

courtrooms in the same county within the California Superior Courts System. (Elliott and Kling, 1997). Most of the participants had access either to LexisNexis, to Westlaw or to both.

The semi-structured interviews included questions about the lawyers' general work (e.g. the type of work they perform, amount of work they delegate, time spent on research), their general computer usage (e.g. their first experience with a computer and the training they received/would have liked to receive/would still like to receive, the tasks they perform on a computer and their attitude if computer work was to be removed), their digital library usage (how they learned to use them, how often they used them, the impact of digital libraries on their work and their use of paralegals – legal assistants - for delegation purposes). They also asked interviewees about their general attitude towards technology and future predictions for technology in the courtroom.

Elliott and Kling (1996) found that the hierarchical role of legal professionals often influenced their access to digital library systems. For example, the elevated role of judges enabled them to have online access to LexisNexis at any time, even from home through a modem link. Similarly, the role of District Attorneys as 'fighters of criminals' versus Public Defenders as 'defenders of criminals' influenced funding for computer equipment in general and, in particular, digital library access. Elliott and Kling also found that, whilst some legal professionals still preferred paper sources (such as one of the judges who would much prefer to browse a paper-based case on his bench rather than excuse himself to use LexisNexis in his chamber), the discontinuation of some paper sources in favour of electronic ones effectively 'forced' legal professionals to increasingly turn to electronic information sources.

Andrews (1993) examined user perceptions of LexisNexis by administering questionnaires and structured interviews to eighteen legal professionals and law librarians. He asked interviewees about the usability of LexisNexis as it stood in 1993. The usability of interface was regarded as a 'significant barrier to usage.' When asked how LexisNexis could be improved, many interviewees suggested the need for search strategy formulation support and support in search syntax usage. This indicates recognition for the need for search-related support rather than support for using specific parts of the interface. This seems to be supported by the fact that the law librarians that were interviewed, who presumably have formed useful search strategies over the years, thought LexisNexis was reasonably easy to use. However, even interviewees who were satisfied with searching on LexisNexis noted problems related to free-text searching (which facilitates searching anywhere in a particular document, including the title and is still achievable in the current versions of LexisNexis Butterworths and LexisNexis Professional at the time of writing). One barrister explained that because law is concept based, it does not sit easily with free-text searching, potentially leading to 'output overload' or difficulties in narrowing down the search so that a balance can be sought between retrieving a manageable number of results but not missing something relevant. Andrews also identified difficulty in retrieving useful search results when free-text search was used, as this type of search does not provide context for the search query. This is in contrast to 'segmented field searches,' which provide context by searching within a particular part of a document or meta-data field.

Oulanov and Pajarillo (2003) also conducted a study on perceptions of LexisNexis, this time using structured questionnaires that were issued to eight academic librarians at Queensborough Community College in New York City. Although the authors do not include a copy of the questionnaire, the included tables reveal that the questionnaire was predominantly based on a five point Likert grading scale aimed at uncovering the librarians' perceptions on three aspects of the resource: its 'retrieval features,' its 'effectiveness' and other usability-related aspects (which the authors rather ambiguously categorise under the heading 'user effort perception criterion'). Questions related to the 'retrieval features' of LexisNexis asked the librarians to rate the resource in its 'use of retrieval techniques,' 'use of Boolean operators,' 'use of fuzzy queries.' Questions related to the 'effectiveness' of LexisNexis asked the librarians to rate it on 'relevance,' 'recall' and 'precision' and questions related to perceived 'effort' asked the librarians to rate the 'ease of use' of the system, 'ease of learning' of the system, and 'necessity of training' amongst others. The authors concluded that LexisNexis 'fared well in almost all questions asked,' with the exception of low ratings on the 'necessity of training' question (which might refer to how far the librarians believe that LexisNexis training is necessary for end-users or how far end-users believe that they require further training in order to use the resource). As the authors do not explain the questions asked

in any detail and only use a small sample size of eight librarians, we cannot draw any useful conclusions from this study. In addition, it can be argued that asking only *librarians* about their perceptions about LexisNexis is unlikely to yield reliable results as librarians generally have high information literacy skills and undergo significant training using electronic resources (the sample of librarians chosen had an average six years of LexisNexis experience).

The user-centred studies reviewed above highlight that it is possible to gain an understanding of lawyers' attitudes towards and use of electronic legal resources by observing them using these resources and by asking them questions about their resource use. We advocate this type of user-centred approach, particularly because we believe in its value in helping us to truly understand users, their information needs and their work. However, none of the above studies have sought to gain a deep understanding of lawyers' use of electronic legal resource with the aim of *improving* these resources. The study by Vollaro and Hawkins (1986) was aimed at discovering when lawyers might search electronic legal resource themselves and when they might turn to intermediaries. Yuan (1997) was interested in lawyers' information search behaviour, but did not have a wider aim of better supporting this behaviour. The other studies reviewed in this section were aimed at gaining an understanding of user attitudes towards and perceptions of electronic legal resources. Therefore none of these studies share our motivation of gaining an understanding of lawyers' use of electronic legal resources in order to feed this understanding into ways of improving these resources (in the case of this thesis, methods for evaluating their functionality and usability).

2.3.2 Studies on lawyers' information-seeking behaviour

As we have illustrated, none of the user-centred studies of lawyers' use and perceptions of electronic legal resources above share our HCI-focused motivation of understanding information behaviour in order to inform the design of tools or methods to support this behaviour. The same is almost the case regarding work in the Information Science domain, which focuses on understanding lawyers' electronic information-seeking behaviour. However, as an exception to the rule, Kuhlthau and Tama (2001) make some recommendations for the design of electronic legal resources based on their insights gained into lawyers' information-seeking behaviour. In this section, we review existing studies of lawyers' information-seeking behaviour from the Information Science domain. None of these studies are focused primarily on electronic information-seeking, however all provide an insight into lawyers' information work. Otike (1999) provides a comprehensive review of several studies, many of which are unpublished theses. We summarise these studies (and Otike's own study) below and then review the highly-cited work on lawyers' information-seeking behaviour by Cole and Kuhlthau (2000) and Kuhlthau and Tama (2001). We also discuss Leckie, Pettigrew and Sylvain's (1996) model of professionals' information-seeking (which was informed by the literature on lawyers' information behaviour) and review subsequent empirical studies aimed at ascertaining whether the model *does* actually apply to lawyers.

Kidd (1978) studied the information needs of solicitors in a private practice in Scotland. Kidd's ultimate objective was to examine whether the solicitors' needs could be supported through the introduction of electronic legal resources. Kidd found that solicitors work in an information intensive environment and were prone to constant interruption. He found that solicitors sought information in order to assist in solving legal cases and in order to keep abreast of the law, often remaining unaware that much processing of information was going on. Kid concluded that although a computerised information retrieval system might support the work of solicitors and was regarded by them as a welcome development, none of the lawyers that were interviewed expressed a great need for a computerised service. Cheatle (1992) also studied the information needs of solicitors in private practice, this time in a London firm. She noted that lawyers with little experience in a particular branch of law sought information frequently because they were still on a learning curve but older, more experienced lawyers did not. Cheatle attributes this to the balance between the lawyer's experience and the complexity of the work undertaken.

Feliciano (1984) administered a closed questionnaire to thirty lawyers in the Philippines in order to find, amongst other things, why they need information and which types of information they need. The vast majority

of respondents claimed they needed information to provide specific information for work in progress, to provide introductionary information needed for work in progress, to 'improve abilities' and to keep informed of work developments. Each of these needs was highlighted by at least 70% of respondents, indicating that these lawyers use information in order to satisfy a broad variety of needs (but indicating little else apart from this).

Published after Otike's review of the literature, Haruna and Mabawonku (2001) administered a similar questionnaire to 361 lawyers in Lagos, Nigeria. The questionnaire had a similar focus on information needs and included questions on the type of information sought by Nigerian lawyers, the types of information sources used (e.g. law reports, journals, Internet) and factors hindering the effective utilisation of these sources. The three highest ranking types of information sought were knowing 'the latest decisions of superior courts' (which 98% of respondents selected), knowing 'most recent legislation' (which 96% of respondents selected) and obtaining 'information on local and international seminars' (which 92% of respondents selected). However Haruna and Mabawonku's use of pre-defined categories for selection might have restricted or biased participants' responses. Similarly, although the authors claim that the questionnaire was 'pre-tested' and 'validated,' they provide no specific detail on how the categories of information type (and the other pre-selected categories used throughout the questionnaire) were selected.

Hainsworth (1992) examined the information-seeking behaviour of fifty-seven appellate judges in Florida. Her aim was to identify, isolate and describe the factors contributing to their behaviour. She found that the quality and depth of information-seeking undertaken by the judges was guided by their internal feeling of satisfaction towards their resulting opinions. Hainsworth found that judges seek information independently and individually, remaining sceptical of information that is provided to them. Although the judges were found to use computers to support their information behaviour, this was mainly to support the reading and writing function of their work. Interestingly, regarding electronic information-seeking (and not mentioned in Otike's article), Hainsworth (1992) also found that the age of a judge did not influence whether he or she might use electronic legal resources. She found that some younger judges, even at the age of forty, did not have a tendency to use electronic resources, whilst some older judges used them frequently.

Otike herself (see Otike, 1999) conducted semi-structured interviews with nine academic lawyers (law teachers and researchers) and twenty-four practising lawyers (mostly solicitors), all based in London and the East Midlands. These lawyers varied in age and legal experience. Otike asked questions surrounding which types of information the lawyers require to meet their needs, the reasons prompting them to seek information, where they seek information from and the factors that influence their information needs and seeking habits.

Consistent with previous findings, over half the lawyers interviewed by Otike delegated their work to trainee lawyers, library staff and law clerks. The common reason for delegation was due to 'insufficient time' available for research. Delegation of research was common among practicing lawyers, but less popular amongst academic lawyers (who argued that since their research often needed to be extensive and incorporate information from many different sources, they were the only people who could obtain the exact information that they need).

Otike (1999) found that lawyers varied greatly in their use of information. She found that frequency of information use varied depending on the type of work the lawyer undertook and the experience that they had in their particular work role and legal area. She also found that practicing lawyers did not consult information as often as academic lawyers because practicing lawyers' needs were found to be confined to a limited area and hence less information was required to satisfy those needs. Rather unsurprisingly, Otike found that lawyers' information needs were greatly influenced by the type of work that they do. She found that practicing lawyers undertook the work roles identified by Leckie et al. (1996) of advocacy, drafting, counselling and managing, whilst academic lawyers undertook work roles involving teaching, research and consultancy. The model by Leckie et al. is discussed later in this chapter.

Regarding types of information required, Otike found a split between the need for detailed and researched information (usually obtained from law journals, law reports or textbooks) and brief, factual information (obtained from a variety of sources including statutes, case summaries or digests). Some lawyers were found to

use both types of information. One lawyer, for example, pointed out that he started with brief factual information to 'get the basic facts right' before moving to case law and textbooks. Otike found that much of this information was obtained through internal sources, most of which were paper-based. Obtaining information from information sources such as colleagues, personal contacts, or at seminars or conferences was also popular. Although obtaining information using electronic resources was mentioned, Otike suggests that their use was confined to large law firms and law school libraries which were able to afford them. As Otike does not state the size of the law firms in which she interviewed practicing lawyers, we cannot make any firm conclusions about whether the cost of electronic resources played a factor in the information-seeking behaviour described by these particular lawyers.

Cole and Kuhlthau's Study on the information-seeking behaviour of 'novice and expert' lawyers

Cole and Kuhlthau (2000) conducted a study with fifteen practicing lawyers in Montreal and New Jersey. They interviewed one group of lawyers at the beginning of their career and compared their concept of task, information and information-seeking with another group of lawyers who had practiced a specific branch of law for over seven years. Cole and Kuhlthau found that lawyers at the beginning of their careers tended to treat problem recognition and solution separately, regarding the information required to form a legal case as objective or 'fact like.' Lawyers further on in their careers tended to conceptualise possible solutions to problems whilst conceptualising the problem.

Cole and Kuhlthau found that an early conceptualisation of possible 'solutions' to a case or client problem enabled the 'expert' lawyers to add value to the information they collect. The authors identified four ways in which experienced lawyers added value. Firstly, lawyers with many years of experience were able to find (and possibly exploit) facts that might be presented in a certain way in order to *"seem real to the judge and jury"* (p. 109). Cole and Kuhlthau identified this type of value added as *"packaging the new knowledge and understanding so that it is effectively communicated to the client or jury and judge"* (p. 109). Secondly, Cole and Kuhlthau noted that extra value could also come from constructing new knowledge and understanding from the information in order to benefit the client, judge or jury. For example, one experienced lawyer spoke of never failing to find *"the key case that will win the case"* (p.109). This lawyer asserted that gaining an understanding of the key issues in a case allowed him to look for holes in a potential opposing argument. Thirdly, the authors highlighted that *"finding cost-effective data that can be processed into value added"* (p. 109) was another way that experienced lawyers added value to a case. For example, one experienced lawyer sought information from other lawyers who had tried a similar case, but had been appointed on the other side of the case. This lawyer said that this allowed her to see their take on the argument, what cases they cited, and what their adversaries said in response to the argument. Finally, Cole and Kuhlthau found that value could be added through the lawyers packaging their knowledge so that they can be as confident as possible that the client, judge or jury will act on the communication (for example by shaping the case to fit the charges that the judge permits the jurors to consider for a particular case). Cole and Kuhlthau argued that even further value could be added by presenting the information that has been found in a way that maximises the chances of the client, judge or jury acting on the communication in a way that exploits the lawyer's expertise. This can be considered as a powerful dynamic projection of how the case might unfold, and how the facts and people in the case can be used to the lawyer's advantage, before the case has even been prepared.

The concept of 'value added' led Cole and Kuhlthau to define legal information-seeking as *"a process of constructing new knowledge and understandings to add value to an enterprise (i.e. a client, jury or judge)"* (p.111). The authors suggest that if systems designers view legal information-seeking in this light, this might lead to the implementation of mechanisms and systems to support legal information-seeking at each stage of the value adding process.

Kuhlthau and Tama's Study on the information-seeking behaviour of lawyers

Kuhlthau and Tama (2001) studied lawyers' information-seeking behaviour, with a particular focus on the variety of information tasks that lawyers undertake, how they use information to accomplish their work and the role that mediators play in the process of legal information-seeking and use. Kuhlthau and Tama conducted structured interviews with eight practicing lawyers in New Jersey, identified as early career experts with six to ten years' experience in their areas of practice. The participants worked in small to medium sized law firms, specialising in a variety of types of law including complex tort, personal injury, contract disputes, criminal matters, environmental cases, real estate matters and landlord-tenant disputes.

The authors found that lawyers' work included both routine tasks (such as dealing with matters that were settled out of court and did not require extensive pre-trial or trial preparation) and complex tasks (which involved preparing a case for trial). Complex tasks were described as being accomplished in stages, moving from fact gathering to defining the theory of a case, to resolving the matter through trial. Participants described 'figuring out a strategy' for a complex case, viewing the task as 'a puzzle to unravel' where they were likely to be aware of a missing slot to fill, but not necessarily of what will fill it (due to facts and evidence that are not readily apparent on the surface). When the participants first started their work, they felt anxious and uncertain about such situations, whilst now that they had experience, they expressed enthusiasm for more complex tasks as this allowed for creativity in presenting cases.

Within the process of undertaking a complex task, Kuhlthau and Tama found that lawyers used sources of information in different ways throughout the process. Initially sources provided overview and background knowledge. They then served the purpose of helping them construct a theory or strategy in the case. The lawyers completed their work when they determined they had used sufficient information to create a persuasive presentation in court.

Many of the participants indicated a strong preference for paper over electronic sources of information. One participant attributed their preference to the ease with which browsing summaries and Shepardising (citation tracking aimed at finding cases or statutes that have been overruled) could be achieved using paper sources compared with electronic sources. Another attributed the preference to the ability to view more documents at once in paper form. Several participants also attributed it to the need to have specific keywords in mind when querying electronic resources. Kuhlthau and Tama identified the need to support the simultaneous presentation of an array of cases, to present information outside a traditional keyword relevancy approach (in order to allow for individual creativity in developing a case) and the need to provide users with a sense of control in doing research and avoid them becoming 'lost' in the computerised information.

Kuhlthau and Tama's findings highlight an interesting question - why did these eight lawyers still prefer to use paper-based sources over electronic resources in the twenty-first century? The evidence above suggests the answer may be related to usability. A book is immediately usable. It is easy to browse tables of contents and indexes to find relevant sections and equally easy to follow references to other books. These were points echoed by participants of our own study of lawyers' information behaviour (discussed later in this thesis) even though, in general, use of electronic legal resources was widespread amongst both the academic and practicing lawyers that took part.

Leckie et al.'s model of professionals' information-seeking and related studies of how the model applies to lawyers

Leckie et al. (1996) propose a model of information-seeking which they claim is generalisable across all professional domains. The model was derived from an analysis of the information-seeking literature rather than from empirical studies and is based on literature surrounding the information-seeking behaviour of engineers, healthcare professionals and, most pertinent to our study, lawyers.

Leckie at al. highlight that professionals play many distinct roles throughout any given day, not only those concerning the provision of expertise and knowledge related to their domains, but other more general roles such as managing, counselling, supervising, planning etc. According to Leckie et al., these roles result in distinct types of activities which in turn shape the type of information needed, the way in which it is retrieved and the ultimate use of that information. Leckie et al. identify four roles of legal professionals from the literature:

1. **Advocacy** – persuading someone (usually a tribunal of some kind) what the law should be, what law should be applied or how the law should be applied. This involves tasks such as determining relevant cases and precedents (which itself requires the professional to search primary and/or secondary legal sources).
2. **Drafting** – preparing documents and correspondence. This gives rise to other lower-level tasks such as determining whether the firm has prepared documents on the same issue previously, or what prior research has been done on the topic (which requires the professional to search the firm's internal sources).
3. **Counselling** – helping and advising clients. Tasks include interviewing clients, responding to telephone queries and representing clients in court.
4. **Managerial** – Managing the firm's resources. Tasks include monitoring the firm's financial situation, training students or delegating work to secretarial staff.

Leckie et al. also highlight, as part of their model, a number of 'intervening factors' that influence professionals' information needs. These include demographics, context, frequency, predictability, importance and complexity. According to Leckie et al, *"any of the components of the model can occur simultaneously, thus representing the true complexity of a professional's work life" (p.180)*.

Wilkinson (2001) applied Leckie et al.'s model by investigating the information-seeking behaviour of 154 practicing lawyers in Ontario, Canada. She used proportionate samples of both corporate and private practicing lawyers from four different sizes of law firm and employed the Critical Incident Technique (see Flanagan, 1954), asking the lawyers to discuss, in detail, a problem which they had recently encountered connected with the practice of law. Wilkinson found that the problems split into two categories: problems related to the administration of law practice (client instructions, errors and omissions, conflicts of interest, communications, relations with other lawyers, representing the clients and the administration of the law practice directly) and problems involving substantive areas of law (administrative law, immigration, corporate and commercial practice, civil and criminal litigation, family law, wills and trusts or real estate). Wilkinson argues that this suggests only two professional roles for lawyers (rather than the four identified by Leckie et al., 1996) - the service provider (when lawyers are engaged with the substantive areas of law in meeting their clients' needs) and the administrator/manager. We argue, however, that Wilkinson's role of 'service provider' may subsume Leckie et al.'s roles of 'advocacy' 'drafting' and 'counselling.' This is because the problems listed above related to the administration of law practice could well be classified under Leckie's original roles. Indeed, one might argue that 'advocacy' 'drafting' and 'counselling' are all sub-roles of 'service provision' and that Wilkinson simply regrouped Leckie's roles at a different level of abstraction rather than changing them outright. Wilkinson also found that none of the problems described by the lawyers seemed to relate to other professional roles mentioned in Leckie et al. (1996), such as 'researcher,' 'educator' or 'student.' This finding seems sensible, since Leckie et al. do not mention these roles in direct relation to lawyers, but instead in relation to roles that are frequently mentioned in other studies of professionals' information-seeking behaviour (studies that, in fact, referred to the information-seeking behaviour of Local Authority Social Service workers, Cardiovascular nurses and other health professionals).

Kerins, Madden and Fulton (2004) applied Leckie et al.'s model to an academic (as opposed to professional) context. They conducted an empirical study with twelve postgraduate Irish law students and found 'similar patterns' in the information-seeking behaviour of students studying to become professionals that Leckie et al. (1996) found in their study of professional lawyers (although the precise details of what makes these patterns 'similar' to the work by Leckie et al. remains unclear). They argue that law students perform similar roles to practicing lawyers as, like professionals, they need to stay abreast with published literature relating to their area of study. Kerins et al. also argue that the essential information skills of legal professionals (locating primary and

secondary materials, evaluating the relevance, applicability and value of these materials to the task at hand, managing the information found and using it for a specific purpose) are also likely to be information skills that are required of law graduates on completion of their education.

Whilst these studies surrounding Leckie et al's model provide an insight into lawyers' work, the model itself is highly abstract in nature as it is aimed at *"[capturing] the complexity of the information-seeking activities of professionals"* (p. 187). This means that the model is not suitable for predicting or describing lawyers' information-seeking behaviour in detail and we regard it more as a broad framework than as an information-seeking model. This is, in effect, how Leckie et al's work has been subsequently used by Wilkinson (2001) and Kerins et al. (2004) – as a framework for organising the insights they have gained from observing academic and practicing lawyers looking for information. This is why we discuss Leckie et al.'s model in this section (as part of a series of related studies examining lawyers' information-seeking behaviour) rather than in the next chapter, were we examine the potential for using information-seeking models to inform the design and evaluation of electronic legal resources.

2.3.3 Studies on lawyers' information use and re-use

In this final section examining previous work on lawyers' attitudes towards and use of electronic legal resources, we examine a series of studies from the HCI domain that have focused on lawyers' information use and re-use. In contrast to the literature already reviewed in this chapter, these studies *do* share a similar motivation to our own – all of the studies involved gaining an understanding of legal work in order to feed this understanding into either recommendations for or prototype interactive systems to support this work.

The first of these studies details Blomberg, Suchman and Trigg's (1996) collaboration with a business division within a Silicon Valley law firm that was involved in developing products that bridged paper and electronic documents. This involved designing a filing cabinet prototype to support practicing lawyers in re-locating and re-using documents that they had previously produced. In designing the prototype, Blomberg et al. found that a central tenet of the firm was to avoid drafting anything from scratch if at all possible. This often involved retaining documents from previous transactions that might prove useful in the future and 'walking the halls,' asking colleague if they have ever drafted a particular type of document or one with specific provisions. This led to the authors designing a system based on scanned versions of documents from the frequently accessed folders in a particular lawyer's actual filing cabinet. When designing the search mechanism for the filing cabinet, Blomberg et al. noted that words like 'corporation' and 'agreement' are, to the legal information seeker, as non-distinguishing as 'and' and 'the.' This led the authors to allow users to define their own stop words on a corpus-by-corpus basis. The authors also used thumbnail images of the scanned documents, which allowed users to quickly identify the 'genre' or style of the legal document. This led to a search paradigm of combined pictorial and text searching.

Komlodi and Soergel (2002) also focused on information use and re-use, specifically legal information seekers' use of their memory and externally recorded search histories to inform their later searches. Komlodi and Soergel found, like Kuhlthau and Tama (2001) and Blomberg et al. (1996), that during the legal research process, law students not only needed to consult electronic legal resources, but also to return to their personal research files. Komlodi and Soergel developed a set of search-history-based user interface tools to support the recording, categorisation and annotation of search results. This was achieved through the system keeping track of user actions and results in an electronic resource and using this expanded history to encourage easier information reuse and future search tasks. This work led to an (albeit limited) form of search histories being incorporated into the Westlaw electronic legal resource.

Marshall et al. (2001) also conducted an information-use related study that involved observing a group of law students prepare for Moot Court (a pretend legal trial). During these observations, Marshall et al. identified the continued importance and authority of books in students' legal research process. Often paper-based sources were used alongside electronic sources (for example a paper book, might list a particular citation, which can then

be retrieved directly from an electronic resource). They also found that many of the users' information-seeking strategies involved following links rather than conducting explicit searches. Finally, they highlighted the utility of using electronic resources for case evaluation. Marshall et al. noted that students began their Moot Court research by identifying key cases, described as a 'launching pad' or 'looking for a thread to pull.' The students then continued to use citations as a point of departure, either as obvious links to a precedent if they came across the citation several times or as a way of determining whether the cases are still 'good law' (i.e. whether they have been overturned, cited more recently or are sufficiently authoritative). Marshall et al. fed their insights into the design of an e-book prototype, which acted as a wireless access device to electronic information resources and supported a wide range of lawyers' reading-related activities such as annotation.

Finally, Jones (2006) conducted Contextual Inquiry observations of eight students and an instructor working in an academic U.S. Legal Aid Clinic. She analysed transcripts and videotapes of the lawyers working with clients and examined the lawyers' LexisNexis and Westlaw search logs and documents produced, with the long-term aim of feeding these findings into the design of a system to support legal information-seeking and use. Her preliminary findings were that these lawyers relied heavily on collaboration (in this case contacting practicing attorneys for assistance and advice) and on knowledge management activities involving locating and looking over documents in a client file. She found that although these documents were rarely annotated, detailed memos were produced when conversations were had with outside experts. Jones concluded that *"extensive collaboration and a heavy reliance on informal sources of information such as listservs and the advice of local experts allowed the students to cope with complex cases which evolved over time"* (p. 358). She suggested that future systems designed to support lawyers in a legal aid clinic such as this should focus on the social nature of legal research by acting as online repositories that facilitate the sharing, annotation and tagging of documents so that they can be located more easily.

These studies all serve to illustrate the approach of gaining an understanding of lawyers and their work in order to inform the design of interactive systems to support this work. This is an approach which closely mirrors our own (although in our case, our empirical study of lawyers information behaviour feeds into the design of two *methods* to support the evaluation and subsequent re-design of electronic legal resources rather than into the design of an electronic tool). These studies also serve to illustrate that lawyers' information behaviour constitutes more than just their information-seeking activities (but also their information use and re-use activities). This suggests the possibility of taking a holistic approach towards examining the full range of lawyers' information behaviour, not just their information-seeking behaviour. This is indeed the approach we take in our empirical study of lawyers' information behaviour (described in chapters 4 and 5).

2.4   Summary and conclusion

In this chapter, we have reviewed the existing literature surrounding lawyers' information behaviour and their attitude towards and use of electronic legal resources. Whilst a few of these studies share a similar motivation to our own – of gaining an understanding of lawyers and their work in order to feed this understanding into the design or evaluation of tools to support their work, these studies have mainly resulted in the design of specialised systems to support information use (as opposed to the design or re-design of electronic legal resources). To the best of our knowledge, no previous studies have aimed to understand lawyers' use of *existing* electronic legal resources in *order to feed this understanding into improving these resources* (whether directly, resulting in design interventions, or indirectly, resulting in guidelines or tools to support design and evaluation). This highlights a gap in research that our thesis aims to fill (through the development of two methods for evaluating the functionality and usability of electronic legal resources). These studies also serve to illustrate the importance of information in lawyers' work and highlight the potential of a user-centred approach towards understanding lawyers' information behaviour. This approach motivates and drives the work described in the remainder of this thesis.

Chapter 3: Potential for using information-seeking models to inform design and evaluation

---

This chapter at a glance…

In this chapter we:

- Motivate our work on lawyers' information behaviour by reviewing a range of existing information-seeking models and examining their potential to inform the design and evaluation of electronic legal resources.

3.1 Overview

One way of informing the design of electronic resources, whether *directly* through the design of new resources (or re-design of existing resources) or *indirectly* through the evaluation and subsequent re-design of existing resources, is by using models of information-seeking and behaviour as theoretical lenses through which to examine information behaviour. However, according to Colbert et al. (1997), *"their fitness for purpose for design has rarely been assessed"* (p. 73). In this chapter, we review several highly cited information-seeking models that split information-seeking into a number of behaviours or processes (whether those processes be physical or cognitive). This is with the aim of asking the question 'what leverage can these models provide in order to inform the design or evaluation of electronic resources?' Or put another way, how useful might these models be to a systems designer for answering any of these three questions:

1. How can we *design a new electronic resource* that supports the behaviours or processes identified by the model?
2. How can we *re-design an existing resource* in order to support or better support the behaviours or processes identified by the model?
3. How can we *evaluate an existing electronic resource* and *subsequently re-design it* so that it supports or better supports the behaviours or processes identified by the model?

We suggest that in order to answer any of these questions, it is necessary to classify these models along a series of related dimensions. In doing so, we must ask some further questions of the model: If the model is to inform the design or evaluation of systems in a different domain or work context to that which it was derived from, how generalisable is it and how broad is the domain/work context scope of the model? If the model is to inform the design or evaluation of systems aimed at supporting a wide range of users' information-seeking activities, how wide is the coverage of the model (i.e. how much of an information-seeking and use episode is it likely to cover?). Finally, if the model is to generate concrete design requirements without requiring a sizeable creative leap by the system designer to turn observations from the model into design suggestions, how concrete and analytical is the model (i.e. how far does it analyse rather than summarise information-seeking and at what level of abstraction/granularity does it operate?) and how far is it explicitly modelled on a set of stages or behaviours?

In the remainder of this chapter, we ask the above questions of information-seeking models by Kuhlthau (1988, 1991), Sutcliffe and Ennis (1998), Marchionini (1995), Ellis and his colleagues (Ellis 1989; Ellis et al. 1993; Ellis and Haugan 1997) and Meho and Tibbo (2003), who propose a model which subsumes aspects of Ellis's model.

To illustrate our argument, we analyse an example narrative using each of the above models as a theoretical 'lens' on the data. The narrative is based on an observation of a Trainee Solicitor (who we fictionally call

'Thomas'), who works for the London office of a large multinational law firm. This observation was conducted as part of the naturalistic study presented in chapters 4 and 5 of this thesis ('Thomas' was participant P10-T in this study). The Trainee was asked to step through how he recently went about looking for electronic legal information as part of his work and the resultant think-aloud data was transcribed, then re-written as third-person narrative in order to make it more concise (but whilst preserving almost all of the detail). It is not necessary to read ahead to chapters 4 and 5 in order to understand the narrative.

The narrative is based on the Trainee's entire information-seeking and use episode, not just part of it. This particular episode was chosen because it illustrates a wide range of information behaviour from recognising that there is a need for information, to drawing the material together, to disseminating it. However, there is no such thing as a 'typical' information episode and therefore this example cannot and should not be used to make strict coverage comparisons between the models, as this would be akin to deciding which off-the-shelf suit vendor provides the overall best-fitting suits by asking one person to try on different suits from several vendors on. The example episode does, however, provide a good basis for illustrating how widely and deeply the models in this chapter capture the information behaviour illustrated by the participant and we believe this is important for illustrating the scope of each model. As we analyse the example narrative multiple times in this chapter (using each of the above models as a theoretical 'lens,') the narrative text is repeated a number of times – in figures 3, 4, 5 and 6. However, in each of these figures, the text in the margins relates to activities, processes or behaviours related to the model being used to analyse the narrative.

Through a discussion of the fit of a variety of information-seeking models to the data provided by this example narrative, we illustrate that although it may be possible to inform design and evaluation by using a number of these models, a particularly useful model (at least when taking the above questions into account) is Ellis's. We suggest that this model in particular is broad in scope, generalisable, covers a reasonably wide range of information-seeking behaviours and (in our opinion most importantly) is highly concrete and analytical and therefore likely to minimise the leap that the person applying the model has to make between parts of the model, resultant user behaviour and interface elements designed to support that behaviour.

3.2   Introduction to information-seeking models

Studies in the Information Science domain have yielded a number of models to help us to understand users' information-seeking behaviour. These have ranged from models that view information-seeking as a series of stages, such as Marchionini's (1995) model, to those that view it as a cognitive and/or affective process (e.g. Sutcliffe and Ennis, 1998; Kuhlthau, 1988) to those that regard information-seeking as a set of interrelated behaviours (e.g. Ellis, 1989; Ellis et al., 1993; Ellis and Haugan, 1997). Other models and approaches have been presented that allow us to conceptualise information-seeking in more ways still, such as a problem-solving activity (Wilson, 1999), an evolutionary process (Bates, 1989), a 'foraging' activity (Pirolli and Card, 1999), a gap in knowledge to be filled (Belkin, 1980) and a sense-making activity (Dervin, 1983; 1992). All of these models and approaches illustrate different possible ways of understanding information-seeking, and are each likely to yield slightly different insights.

One way that researchers have sought to themselves 'make sense' of these models and the different angles on information-seeking that they provide is by comparing information-seeking processes, activities or behaviours across several models and sometimes subsuming either one or several of other peoples' models under broader processes, activities or behaviours (see Wilson, 1999; Meho and Tibbo, 2003, O'Brien and Buckley, 2005). A different way of making sense of information-seeking models is illustrated by Ingwersen and Järvelin (2005), who present several dimensions along which they claim these models can be characterised:

1.   Broad vs. narrow scope (i.e. the degree to which the model encompasses several domains, aspects of behaviour and work contexts).
2.   Process vs. static (i.e. the degree to which the model is explicitly based on a set of stages).

3. Abstract vs. concrete (i.e. the degree to which the model is based on abstract theory or interpretations as opposed to behavioural observations).
4. Summary vs. analytical (i.e. the degree to which the model seeks to either summarise or analyse the important objects and relationships in the information-seeking process).
5. General vs. specific (i.e. the degree to which the model can be applied or generalised over a range of empirical domains).

Whilst Ingwersen and Järvelin (2005) suggest where some of the more highly cited models might be located along several of these dimensions, they do not rigorously categorise all of the models that they review. Indeed, the categorisation process can be regarded as reasonably subjective and made more complicated by the fact that the above dimensions, as recognised by Ingwersen and Järvelin 'interact and overlap.' This does not mean, however, that the above categories cannot be utilised to help provide focus to this review (just that it should be recognised that the dimensions are not orthogonal and that decisions about where a particular information-seeking model might be categorised along a particular dimension is somewhat subjective). In order to decide which information-seeking models to review, we chose to pose design and evaluation-focused questions to each of the candidate models which were aimed at helping us to answer the broader question of 'what leverage can these models provide in order to inform the design and evaluation of electronic resources?' These questions were based around Ingwersen and Järvelin's dimensions, but framed with a specific design and evaluation focus. The questions we asked about each candidate model were:

1. If the model is to inform the design and evaluation of systems in a different domain or work context to that which it was derived from, **how generalisable** is it and **how broad is the domain/work context scope** of the model?
2. If the model is to inform the design and evaluation of a system aimed at supporting a wide range of users' information-related activities, **how wide is the coverage** provided by the model (i.e. how much of an information-seeking and use episode is it likely to cover?).
3. If the model is to yield concrete design suggestions without requiring a sizeable creative leap by the system designer to turn observations from the model into design suggestions, **how concrete and analytical** is the model (i.e. how far does it analyse rather than summarise information behaviour at what level of abstraction/granularity does it operate at?) and **how far is it explicitly modelled on a set of stages or behaviours?** This question could also be asked as 'how deep is the coverage provided by the model?'

Of course, it would be unrealistic to ask the above questions about every information-seeking model that has been published in the Information Science domain. Therefore we have restricted our review to those which upon a cursory evaluation appear to be fairly generalisable with a broad scope, fairly wide coverage of behaviour, reasonably concrete and analytical and behaviour or process-based (whether those processes be physical or cognitive). In effect, we have used Ingwersen and Järvelin's dimensions to narrow down our choice of models to review, ending up with a group of candidate models that have reasonably similar characteristics according to Ingwersen and Järvelin's dimensions (but nonetheless differ somewhat across these dimensions). We use the discussion in the rest of this chapter as a basis of highlighting how we perceive these candidate models to differ across the dimensions. This discussion is supported through reference to an example electronic information-seeking and use episode of a Trainee Solicitor working for a large multinational law firm, which has been analysed using each of the candidate models as a theoretical 'lens' on the data.

We begin, in section 3.3.1, by presenting and discussing the cognitive aspects of Carole Kuhlthau's (1988, 1991) Information Search Process (ISP) model, followed by Alistair Sutcliffe and Mark Ennis's (1998) cognitive process model in section 3.3.2, Gary Marchionini's (1995) process model in section 3.3.3, and the behavioural model by David Ellis and his colleagues in section 3.3.4 (which also includes a discussion of Meho and Tibbo's (2003) model, which subsumes many of Ellis's behaviours). This is followed by a summary and synthesis of our discussion in section 3.3.5.

3.3 Review of selected information-seeking models

3.3.1 Kuhlthau's Information Search Process (ISP) model

Kuhlthau's Information Search Process (ISP) model is based on constructivist theories of system-related skills development and focuses on the cognitive and affective aspects of information-seeking. The ISP model was formulated through a series of user-centred information-seeking studies, making it one of the few models to have been empirically validated. The participants of the first study (described in Kuhlthau, 1988) were twenty-six 'academically capable' college-bound high school seniors in two classes of advanced placement English. Participants were asked to keep diaries whilst they completed a research essay, noting their feelings, thoughts and actions related to their search for information. Participants were also set a second assignment and asked to keep search logs whilst completing it, noting the sources used, the procedure used for finding them and whether they were useful, highly useful or not useful. In addition, participants were asked to write a paragraph about their topic two weeks after each assignment was set, and again after it had been submitted. Finally, six students were interviewed on six separate occasions during the course of the two assignments in order to verify and explain the data collected in the diary study, search logs and writings.

Kuhlthau also conducted several verification studies (reviewed in Kuhlthau 1991), which include a case study addressing how the English students' perceptions of the Information Search Process had changed over four years of college and a longitudinal study involving the same case study participants. According to Kuhlthau 1991, these studies showed that her ISP model held over time for the English students. The remaining studies examined the information-seeking process of high, middle and low-achieving high school seniors and validated the model with a large sample of academic, public and school library users. In subsequent years, Kuhlthau illustrated that her model could be generalised outside of an academic domain by relating it to the information-seeking behaviour of a Securities Analyst (see Kuhlthau, 1997, 1999). Although no specific reference was made to the above stages in Kuhlthau and Tama (2001), a study in which the authors examined lawyers' information-seeking, the authors noted that the lawyers *"described a process similar to that of the ISP model"* (p. 40).

The Information Search Process model that was empirically derived from these studies describes the information-seeking process as a series of cognitive stages, with each stage leading to the next. The stages are as follows (and are described in Kuhlthau, 1991):

- **Initiation** - Which involves becoming aware of the need for information when facing a problem. During this stage, information-seekers might frequently discuss possible topics and approaches.
- **Selection** – Which involves identifying and choosing a general topic for seeking information. During this stage, a typical action involves conferring with others.
- **Exploration** – Which involves seeking and investigating information on the general topic. Actions involve locating information about the general topic, reading to become informed and relating new information to what is already known.
- **Focus formulation** - Which involves fixing and structuring the problem to be solved.
- **Collection** – Which involves gathering pertinent information for the focused topic. Actions involve selecting information relevant to the focused perspective of the topic and making detailed notes on that which pertains specifically to the focus (as general information on the topic is no longer relevant after formulation).
- **Presentation** – Which involves completing information-seeking, reporting and using the result of the task. Actions involve conducting summary searches, usually involving in information with decreased relevance and increased redundancy.

Fit of Kuhlthau's ISP model to our summary information-seeking and use episode data

If we examine the narrative of Thomas the Trainee's information-seeking and use episode in figure 3, we can note that most of the information-seeking and use episode is covered by some of Kuhlthau's process stages (see figure 3). However, this coverage is reasonably broad-focused due to the nature of the model. In addition, only four of Kuhlthau's process stages were identified in the example data; initiation, focus formulation, collection and presentation. This may be to due to the fact that Kuhlthau's model was not derived from separate, single information-seeking episodes but from data of students working on assignments over a period of time – Kuhlthau's process stages *can* be identified in single information-seeking episodes but perhaps a broader range of stages might be identified in a longitudinal study.

Potential for Kuhlthau's ISP model to inform the design and evaluation of electronic resources

In relation to the design of electronic resources, Kuhlthau suggests that *"these systems need to be made more proficient at accommodating a range of tasks in response to the users' articulation of the problem at all stages in the ISP, such as offering preliminary, exploratory, comprehensive, or summary searches according to the state of the user's problem"* (p. 370). In practice, however, it may be difficult (although not by any means impossible) to use the model to inform electronic resource design and evaluation. We illustrate this point of view by asking ourselves the design and evaluation-focused questions (explained in more detail in section 3.2 and reproduced below):

1. How generalisable is the model and how broad is its domain/work context scope?
2. How far is the model explicitly based on a set of stages or behaviours?
3. How wide is the coverage of the model?
4. How concrete and analytical is the model (i.e. how deep is the coverage of the model)?

The first two questions are reasonably straightforward to answer. As with all of the models reviewed in this chapter, Kuhlthau's ISP model is broad in scope and intended to be applied to a wide range of domains. As we have noted, Kuhlthau has not only checked the validity of the model with a wider sample and that it holds over time for a given set of students, but has empirically validated the model in a non-academic domain. This makes the model highly generalisable and broad-scoped in terms of domain and work context. Whilst the model is explicitly based on a set of stages, these stages are fairly high-level (probably due to the fact that they refer to cognitive actions as opposed to physical actions). This places the ISP model somewhere in the middle of the concrete-abstract and summary-analytical dimensions proposed by Ingwersen and Järvelin (2005) (i.e. the model is certainly not as 'concrete' and 'analytical' as they get, but neither is it especially 'summary' or 'abstract'). The cognitive aspects of the model do mean, however, that the model not only covers cognitive processes that can be associated with physical behaviours, but also processes that are mainly internal (such as 'initiation' and 'selection'). This gives the model a fairly wide coverage.

In the summarised information-seeking and use episode below in figure 3, Kuhlthau's cognitive process stages are presented in the margin next to the part of the text that has been deemed to illustrate that particular process. The wide breadth of coverage by Kuhlthau's model is illustrated by the fact that sentences at the extremes of the text have been covered by the model (our lawyer Thomas's recognition of a need for information-seeking was deemed to be an example of 'initiation' and his write-up of the research note was deemed to be 'presentation'). However it is debatable how far electronic resources should be expected to support cognitive actions that come early in the information-seeking process (such as deciding to look for information and choosing a general topic to look for information on). Therefore breadth of coverage in the direction of mostly internal processes (i.e. those processes traditionally carried out at the very beginning of information-seeking) is by no means essential in order to inform the design and evaluation of electronic resources. We argue that more essential than a model's breadth of coverage is its *depth* of coverage, which in Kuhlthau's case can be noted in figure 3 by the fact that much behavioural detail is subsumed within each of the process stages. Indeed the large middle portion of the text is all concerned with 'gathering pertinent information for the focused topic' (and hence was deemed to

illustrate 'collection'). However, much of the detail would be lost if we simply summarised the middle portion of the text as 'collection.' The problem, we therefore argue, lies with the reasonably high level of abstraction that the model is presented at. If a designer asks himself how can we design to support or better support Kuhlthau's collection stage, the answer is not immediately obvious. This is because it is necessary for the designer to make an inference (or a more detailed analysis) of what lower-level behaviour 'collection' entails and then to use this inference or analysis to determine how the system can be designed to support this lower-level behaviour. This is by no means impossible but requires a sizable creative leap to be made by the designer of the support tool, who has to 'jump' between the cognitive process, the interface-level behaviour associated with that process and the design intervention aimed at supporting or better supporting that process. This is a leap that is an inherent problem when examining the potential leverage of all process-based models, as we will see in the next two sections.
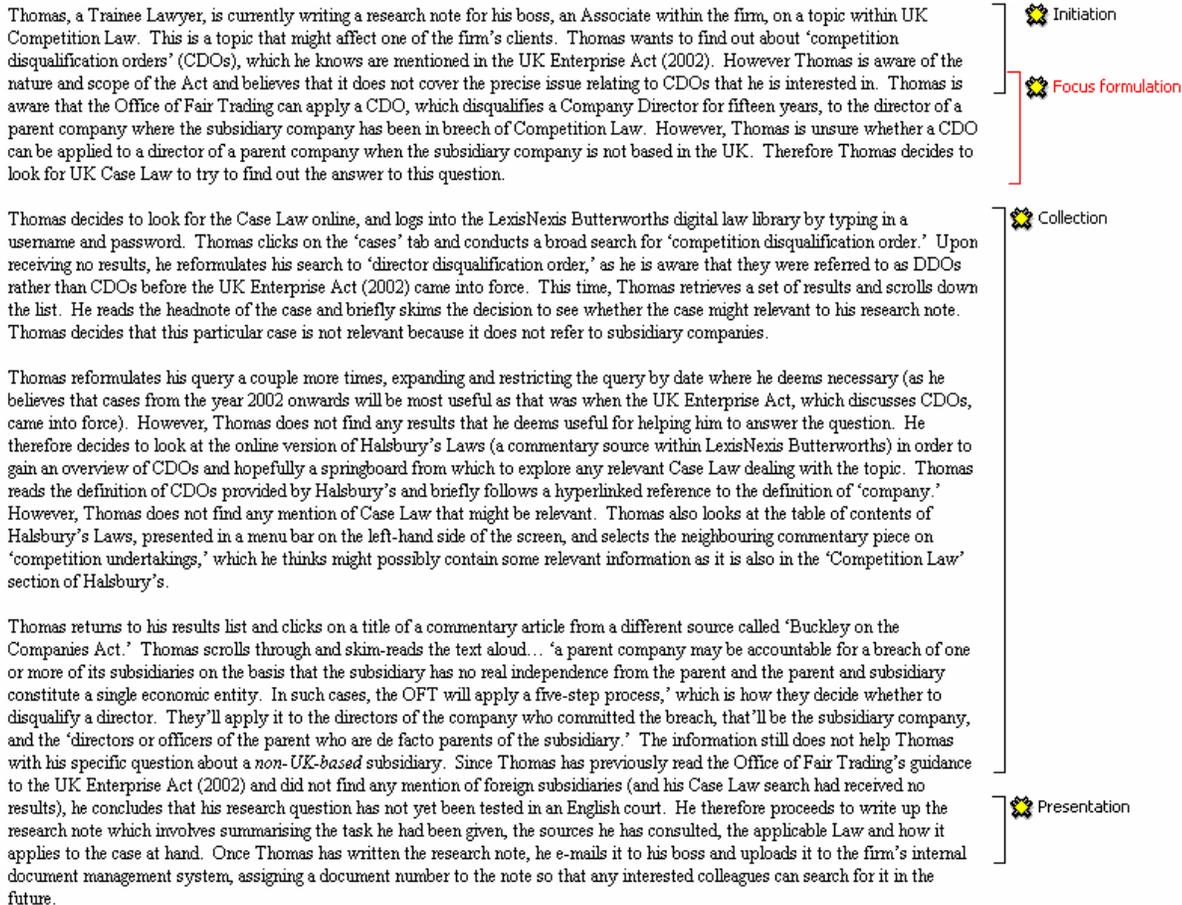
Figure 3: Summary of participant P10-T'S legal information-seeking and use episode (observed during the think-aloud part of our Contextual Inquiry) encoded with the physical process stages from Kuhlthau's ISP model.

3.3.2   Sutcliffe and Ennis's cognitive process model

Sutcliffe and Ennis (1998) propose a cognitive process model of information searching activities. Although this model has not been empirically validated, the development of the model itself motivated an empirical study of 'novice' and 'expert' MEDLINE users, who were all final year medical students (see Sutcliffe, Ennis and Watkinson, 2000). As the basis of their model, Sutcliffe and Ennis suggest that searching for information falls into four major cognitive activities:

1. **Problem identification** – Identifying the initial goal or information need and, if the problem is complex, decomposing it into smaller components and prioritising those components.
2. **Needs articulation** – Expressing the information need or initial goal as concepts or high level semantics. These are refined into lower level terms which are utilised in queries.
3. **Query formulation** – Identifying search terms and transforming them into the query language supported by the search system.
4. **Evaluating results** – Scanning the results set or examining the contents in detail in order to decide whether to accept the retrieved results or continue searching.

Sutcliffe and Ennis (1998) also present behavioural strategies for each of the four activities. According to Sutcliffe and Ennis, when identifying the problem, users are likely to employ general problem solving strategies. For example, 'divide and conquer' to horizontally partition the problem into separate areas or top-down decomposition to split the problem into sub-components. Another problem identification strategy Sutcliffe and Ennis describe is deciding whether to begin with a detailed query, in which lots of effort is required in order to try to get the query 'right first time' or to focus on a sensible, basic query with the expectation of re-formulating query terms based on results feedback. Sutcliffe and Ennis highlight that specific problem identification strategies depend on the user's goal. They point out that an exploratory goal implies a browsing style of search, whereas a more specific need suggests querying.

When articulating needs, according to Sutcliffe and Ennis (1998), users refine their initial task concepts into queries. If their domain knowledge is poor, Sutcliffe and Ennis suggest that users will have to acquire search terms from the environment (such as through indexing terms) or by finding similar cases that relate to their information goal and trying to extract search terms from them. During the stage of query formulation, Sutcliffe and Ennis highlight that users' activities are highly constrained by the user's device knowledge and knowledge of query languages. Therefore, they only suggest general strategies such as forming complex queries with Boolean logic or using simple queries and iterating the search to find appropriate results. When evaluating results, users may either decide to browse a results set or view the content of retrieved items. A variety of different sampling strategies can be used, depending on how the system presents the results and the quantity of retrieved items. Evaluation decisions feed back into query reformulation strategies, such as broadening or narrowing queries by adding or reducing search terms and/or synonyms, by stemming or de-stemming queries to add or remove suffixes, by adding or removing constraints (such as journals to be searched) or by converting conjunctions into disjunctions and vice versa (i.e. ORs into ANDs to decrease the volume of results and ANDs into ORs to increase the volume of results).

Fit of Sutcliffe and Ennis's model to our summary information-seeking and use episode data

All of Sutcliffe and Ennis's cognitive activities are observable in our summary information-seeking and use episode (see figure 4). The dense repetition of the query formulation and results evaluation activities in the middle of the text help to highlight the often cyclic nature of the information search process (i.e. that queries are often edited several times during an information-seeking episode depending on the perceived usefulness of the returned results). However, activities outside of the traditional search process (and hence some of our data, such as actions involving browsing rather than searching) are not covered by the model. This is discussed in more detail below, with reference to the potential for Sutcliffe and Ennis's model to inform design and evaluation.

Potential for Sutcliffe and Ennis's model to inform the design and evaluation of electronic resources

In relation to the design and evaluation-focused questions posed at the beginning of the chapter, Sutcliffe and Ennis's model is, like Kuhlthau's model, broad in scope. However a lack of empirical validation leaves its generalisability still to be determined. Also, like Kuhlthau's ISP model, Sutliffe and Ennis's model is based on fairly high-level cognitive processes (arguably even more high-level than Kuhlthau's). This would normally orientate the model slightly more towards 'abstract' on the concrete-abstract dimension and slightly more

towards 'summary' on the summary-analytical dimension. However, the inclusion of lower-level behavioural strategies for achieving each cognitive process makes the model slightly more concrete and analytical than Kuhlthau's model. With reference to the information-seeking example in figure 4 encoded with Sutcliffe and Ennis's process stages, we note a similarly broad coverage as Kuhlthau's model (although this model stops at the stage of evaluating results rather than the stage of presenting them, as with Kuhlthau's ISP model). Regarding depth of coverage, the model does not face the same lack-of-depth problems associated with a high-level of abstraction as Kuhlthau's model. However the cognitive nature of the model still necessitates, as with Kuhlthau's model, a leap between cognitive processes, interface-level behaviour and design.

With the above in mind, we argue that Sutcliffe and Ennis's model does not lend itself well to informing the design or evaluation of electronic resources. However, it does lend itself towards providing generic searching support, focusing primarily on query formulation and reformulation as opposed to providing requirements for the design of electronic resources. However, there is not a clear distinction to be made between electronic resources and systems designed to support users with information retrieval. This is because it is just as feasible for the designers of electronic resources to integrate search support within the system than to design a separate tool to support users when searching. The latter approach has indeed been adopted by both Sutcliffe and Ennis (2000) and Stelmaszewska, Blandford and Buchanan (2005). Sutcliffe and Ennis (2000) provide user-system dialogues to encourage more complete articulation of queries and both Sutcliffe and Ennis and Stelmaszewska et al. provide a form of context-sensitive search guidance. Although the guidance and support given by these systems often relates to system-specific syntax for formulating queries, both of these systems concentrate on supporting users with generic searching behaviour rather than with wider information-seeking behaviour.

Figure 4: Summary of participant P10-T's legal information-seeking and use episode (observed during the think-aloud part of our Contextual Inquiry) encoded with the cognitive process stages from Sutcliffe and Ennis's model.

3.3.3 Marchionini's information-seeking process model

Marchionini (1995) presents a part-cognitive part-physical process model of the information-seeking process, which he regards to be both systematic and opportunistic, in the sense that the user's information-seeking direction can change and evolve during the information-seeking process itself. Therefore although Marchionini presents an essentially linear process, the model also takes into account that information-seeking can also be a highly iterative and repetitive process.

Marchionini (2007) highlights that *"although information seeking is driven by human needs and behaviours and thus highly variable, there are several common sub-activities that may be supported by good technical design"* (p. 207). These sub-activities, discussed both in Marchionini (1995) and Marchionini (2007), are listed below. Marchionini (2007) also gives examples of interactive systems designed to support each of these sub-activities.

1. **Recognising and accepting an information problem** – which depends on the user becoming 'aware' of the problem. Once a user is aware of a problem, the problem must be accepted rather than suppressed. Marchionini suggests that systems that invite interaction and support engagement lead users to accept information more readily.
2. **Defining and understanding the problem** – which depends on the users' knowledge of the task domain and may be influenced by the setting (i.e. the physical, social and psychological context of the task). According to Marchionini, to understand and define a problem, it must be limited, labelled and a form or frame for the answer determined.
3. **Choosing a search system** – which depends on the user's domain knowledge, the expectations about the answer that may have been formed when defining the problem and task, and the scope of their personal information infrastructure (which is in turn dependent on past experience with information problems in general, their general cognitive abilities and experience with particular systems). Marchionini gives the example of lawyers being able to readily determine whether information in their assigned searches would be found in case law, statutes or treatise.
4. **Formulating a query** – which depends on how the user manages to map their vocabulary onto the systems' vocabulary (semantic mapping) and how they map the strategies and tactics that that they deem to be useful for the search onto the rules and features that the system interface allows (action mapping).
5. **Executing search** – which is the execution of the physical actions to query an information source.
6. **Examining results** – which depends on how the user chooses to rate the relevance of the retrieved search results or documents and whether and how to proceed with the search (e.g. by reformulating the query).
7. **Extracting information** – which depends on how relevance judgements made when examining results have been influenced by the previous manipulation and integration of information into the user's domain knowledge. A user is likely to judge more items as relevant and later discover that they are not so relevant when they have little domain knowledge. As this knowledge increases, users are likely to extract more relevant information and hone their domain or task specific rules and strategies for doing so.
8. **Reflecting/iterating/stopping –** which involves deciding when and how to iterate the search and therefore depends on the information-seeking process itself.

Fit of Marchionini's model to our summary information-seeking and use episode data

In a similar way to Sutcliffe and Ennis's model, Marchionini's model is strongly search focused. This can be noted by the dense coding of information-search-related activities, mostly in the middle two paragraphs of the text of our summary information-seeking and use episode (see figure 5 below). The model covers more of our data than the previous two models by Kuhlthau and Sutcliffe and Ennis, but still with a distinct cognitive focus that has implications for how useful the model might be for informing the design or evaluation of electronic resources. This is discussed in further detail below.

Potential for Marchionini's model to inform the design and evaluation of electronic resources

Like Sutcliffe and Ennis's model, Marchionini's process model is broad-scoped, but has not yet been empirically validated. As with both Kuhlthau's and Sutcliffe and Ennis's models, the model has a wide coverage that spans recognising and accepting the need to find information to stopping after retrieving the desired results (or indeed not retrieving the required results). In addition, as with the other models that have cognitive elements, Marchionini's model also covers internal processes such as recognising and accepting an information problem and defining and understanding it. These internal processes are particularly difficult to design for. Indeed, Marchionini (1995) suggests that *"recognition and acceptance are typically ignored by system designers as they are viewed as user specific and thus uncontrollable"* (p. 51). He goes on to say, however, that *"systems that invite interaction and support satisfying engagement lead users to accept information problems more readily"* (p. 51). While this may well be the case, understanding these processes does not provide any direct scope for informing the design or evaluation of electronic resources.

Regarding depth of coverage, Marchionini's model is based on more, and finer-grained, stages than Kuhlthau's and Sutcliffe and Ennis's models (as can be noted by the densely-spaced process categories assigned to the text in figure 5 to denote examples of Marchionini's processes). However although the model is somewhat concrete and analytical, the model does not have as deep coverage as it could do. For example, whilst it is certainly feasible for designers to ask how they can design to support or better support 'query formulation' or 'results examination' (as illustrated by Colbert et al., 1997), these processes are still fairly abstract and encompass a number of behaviours. For example, the stage 'extract information' has been assigned to four parts of the text in figure 5, three of which illustrate a slightly different aspect of information extraction (reading structured parts of the document), browsing through a table of contents and skim-reading the full-text of a document). In addition, like Sutcliffe and Ennis's model, Marchionini's model has a heavy focus on searching behaviour (noted by several cycles of query formulation and execution, results examination and reflecting in the text in figure 5). This makes the model more suitable for informing the design and evaluation of systems to provide search support than the design of electronic resources in general. Indeed, Marchionini (2007) details the evaluation of a novel search support system that asks users to enter query terms, then suggests additional or alternative terms based on the results that would have been returned if the current search had been executed.

Thomas, a Trainee Lawyer, is currently writing a research note for his boss, an Associate within the firm, on a topic within UK Competition Law. This is a topic that might affect one of the firm's clients. Thomas wants to find out about 'competition disqualification orders' (CDOs), which he knows are mentioned in the UK Enterprise Act (2002). However Thomas is aware of the nature and scope of the Act and believes that it does not cover the precise issue relating to CDOs that he is interested in. Thomas is aware that the Office of Fair Trading can apply a CDO, which disqualifies a Company Director for fifteen years, to the director of a parent company where the subsidiary company has been in breech of Competition Law. However, Thomas is unsure whether a CDO can be applied to a director of a parent company when the subsidiary company is not based in the UK. Therefore Thomas decides to look for UK Case Law to try to find out the answer to this question.

Recognise/accept
Define/understand problem

Thomas decides to look for the Case Law online, and logs into the LexisNexis Butterworths digital law library by typing in a username and password. Thomas clicks on the 'cases' tab and conducts a broad search for 'competition disqualification order.' Upon receiving no results, he reformulates his search to 'director disqualification order,' as he is aware that they were referred to as DDOs rather than CDOs before the UK Enterprise Act (2002) came into force. This time, Thomas retrieves a set of results and scrolls down the list. He reads the headnote of the case and briefly skims the decision to see whether the case might be relevant to his research note. Thomas decides that this particular case is not relevant because it does not refer to subsidiary companies.

Choose search system
Examine results *Formulate/execute query*
Formulate/execute query
Examine results
Extract information
Reflect/iterate/stop

Thomas reformulates his query a couple more times, expanding and restricting the query by date where he deems necessary (as he believes that cases from the year 2002 onwards will be most useful as that was when the UK Enterprise Act, which discusses CDOs, came into force). However, Thomas does not find any results that he deems useful for helping him to answer the question. He therefore decides to look at the online version of Halsbury's Laws (a commentary source within LexisNexis Butterworths) in order to gain an overview of CDOs and hopefully use it as a springboard from which to explore any relevant Case Law dealing with the topic. Thomas reads the definition of CDOs provided by Halsbury's and briefly follows a hyperlinked reference to the definition of 'company.' However, Thomas does not find any mention of Case Law that might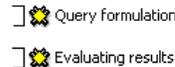 be relevant. Thomas also looks at the table of contents of Halsbury's Laws, presented in a menu bar on the left-hand side of the screen, and selects the neighbouring commentary piece on 'competition undertakings,' which he thinks might possibly contain some relevant information as it is also in the 'Competition Law' section of Halsbury's.

Formulate/execute query
Choose search system *Examine results*
Define/understand problem *Reflect/iterate/stop*
Extract information

Extract information *Reflect/iterate/stop*

Thomas returns to his results list and clicks on a title of a commentary article from a different source called 'Buckley on the Companies Act.' Thomas scrolls through and skim-reads the text aloud… 'a parent company may be accountable for a breach of one or more of its subsidiaries on the basis that the subsidiary has no real independence from the parent and the parent and subsidiary constitute a single economic entity. In such cases, the OFT will apply a five-step process,' which is how they decide whether to disqualify a director. They'll apply it to the directors of the company who committed the breach, that'll be the subsidiary company, and the 'directors or officers of the parent who are de facto parents of the subsidiary.' The information still does not help Thomas with his specific question about a *non-UK-based* subsidiary. Since Thomas has previously read the Office of Fair Trading's guidance to the UK Enterprise Act (2002) and did not find any mention of foreign subsidiaries (and his Case Law search had received no results), he concludes that his research question has not yet been tested in an English court. He therefore proceeds to write up the research note which involves summarising the task he had been given, the sources he has consulted, the applicable Law and how it applies to the case at hand. Once Thomas has written the research note, he e-mails it to his boss and uploads it to the firm's internal document management system, assigning a document number to the note so that any interested colleagues can search for it in the future.

Examine results
Extract information

Reflect/iterate/stop

Figure 5: Summary of participant P10-T's legal information-seeking and use episode (observed during the think-aloud part of our Contextual Inquiry) encoded with the part-cognitive-part-physical process stages from Marchionini's model.

3.3.4 Ellis's behavioural model

Ellis's (1989) model is based on interviews with academics and commercial researchers from a number of scientific disciplines: the social sciences (Ellis, 1989), physical sciences (Ellis et al. 1993) and engineers and research scientists (Ellis and Haugan, 1997). The latter study served to validate the model outside an academic domain, whilst interview data by Smith (1988) was analysed by David Ellis to validate the model in the non-scientific discipline of English Literature. In addition, a more recent study has been conducted by Meho and Tibbo (2003) on another group of social scientists (researchers on the topic of 'stateless nations'). This study re-examined Ellis's findings in order to see if they were still applicable to social scientists in 2003, where electronic information-seeking had further increased in popularity.

The original model, described in Ellis (1989), was empirically derived from interviews of academic social scientists at the University of Sheffield. In this and each of his subsequent studies, Ellis chose participants based on the principles of theoretical sampling. In his study with social scientists, this involved interviewing members of every department in the faculties of Social Science and Education (both researchers and students) and including a sample of social scientists who conduct electronic as well as paper-based information-seeking. In the interviews, Ellis asked participants to describe their work and the information-seeking activities involved in their work. As reported in Ellis (1993), he was supported by an interview guide which included questions on the researchers' research and teaching interests and characteristics of information use, as well as more specific questions on the use of indexing, abstracting and online information services. The interviews were transcribed and analysed using Glaser and Strauss's Grounded Theory approach (see Strauss and Corbin, 1998 for a discussion of the origins of the approach). Ellis revised and refined his categories, noting that *"as the analysis progressed it became clear that certain categories were more or less synonymous with others and that some*

candidate categories were better treated as properties of broader categories" (p. 174). In the end, Ellis ended up with six major behavioural characteristics which, he argued, *"seemed to subsume satisfactorily the important characteristics of the information-seeking patterns"* (p. 174) displayed by the social science researchers.

To see whether similar behavioural characteristics might be identified in other scientific domains, two of Ellis's students, Cox and Hall, conducted separate studies on research physicists and research chemists. In addition, Ellis and Haugan (1997) conducted a study on engineers and research scientists in the Research and Development department of an oil and gas company. To see whether similar characteristics might be identified in a non-scientific domain (i.e. the domain of humanities), another of Ellis's students conducted a study on English Literature researchers (which is briefly reported in Ellis, 1993). These studies also involved interviews that were transcribed and analysed using the Grounded Theory approach. As with Ellis's original study, theoretical sampling was used to decide on who to interview. Cox (1991), reported in Ellis et al. (1993) took a vertical slice across the doctoral and post-doctorate level researchers within the Atomic, Molecular and Polymer Physics group at UMIST. Hall (1991), also reported in Ellis et al. (1993), interviewed staff from the Department of Chemistry at the University of Sheffield that worked across four fields of Chemistry - Organic, Inorganic, Physical and Theoretical. Ellis and Haugan (1997) selected a sample of participants to represent a mix of scientists and engineers, a vertical slice across job roles, the different phases in R&D projects, the three basic types of R&D project (incremental, radical and fundamental) and the different research areas within the company.

Ellis compiled an interview guide, which comprised a list of questions that were used in his semi-structured interviews. As with the information-seeking behaviours he identified, Ellis's interview questions also emerged from the data (Ellis 2007, personal communication). This meant that interviews became progressively more focused (and the interview questions more structured) as more participants were interviewed. Ellis and his students all included the same interview guide in their dissertations. The questions in the guide first probed the scientists' research and teaching interests by asking questions such as 'how did you commence work on [your research] topics?', 'how do you keep up-to-date with developments relating to this topic?' and 'what criteria do you employ when assessing whether to follow up material?' Then the questions focused on characteristics of the scientists' information use, such as 'are there any distinctions between the sources or the material which are of particular importance to you?', 'which are the principal ways you have employed, or intend to employ to publish your own results?' and 'do you follow up references cited in material consulted?' The final set of questions was more general, focusing on the scientists' use of computer-based abstracting, indexing and citation index services.

The information-seeking characteristics identified in these studies are, according to Ellis (1989), non-sequential and it is possible to display more than one characteristic at any given time. Ellis highlights that although the characteristics are likely to interrelate or interact with each other, this *"will depend on the unique circumstances of the information-seeking activities of the person concerned at that particular point in time"* (p.178) and therefore it is only possible to describe the categories in abstract form when they are divorced from a particular information-seeking episode. As Ellis asserts, *"the model does not, therefore constitute a hierarchic sequence for classifying the information-seeking patterns, nor a prescriptive set of search heuristics, but rather a set of related categories which, taken together, can be used to describe individual information-seeking patterns, and perhaps help to explain details of their topography"* (p. 179). In a similar vein, Ellis (2006) highlights that *"the relation between the different characteristics can only be described in the most abstract and general terms unless there is reference to a particular information-seeking pattern at a particular time"* (p. 139).

The characteristics identified are listed below, along with examples of each characteristic, a brief mention of design considerations for electronic resources related to most of the characteristics and how current electronic resources already or might in the future support these characteristics:

- **Starting/surveying** – According to Ellis et al. (1993), starting involves *"activities characteristic of the initial search for information"* (p. 359). Ellis and Haugan (1997) elaborate on this definition, suggesting that this behaviour (which they re-named 'surveying') is *"characteristic of the initial search for information to obtain an overview of the literature within a new subject field, or to locate key people*

operating in this field" (p. 395). The social scientists in Ellis's (1989) study achieved this behaviour by seeking out people that knew key references and authors in the area, by reading reviews of materials, by consulting bibliographies, abstracts, indexes and library catalogues and by using previously collected or newly recommended starter references. Many of the physicists interviewed by Ellis et al. (1993) were provided with starter references by their PhD supervisors. Librarian intermediaries generally carried out surveying for the engineers and research scientists in Ellis and Haugan's (1997) study, who surveyed the literature not only to find background information for new parts of an R&D project, but also to generate ideas for new projects or to carry out pre-studies prior to a project. In addition, one-third of Meho and Tibbo's social scientists started their research from their own personal collection first, which would include both primary and secondary sources (see Meho and Tibbo, 2003). Ellis (1989) highlights that 'starting' behaviour may be supported in the design of electronic resources by alerting users to key ideas or documents and providing them with overviews of research areas to facilitate later chaining. Ellis also states that this behaviour can be supported by helping users to identify review-type or heavily-cited materials and by providing an indication of the sources that publish material in the required area. These design recommendations are sometimes, but not often, supported in current electronic resources.

- **Chaining** - *"Following chains of citations or other forms of referential connections between material"* (Ellis, 1989, p.179). Ellis (1989) highlights that there are two types of chaining; forwards chaining (which involves identifying and accessing documents which have subsequently cited the current document) and backwards chaining (which involves following references to documents that have been cited in the current document). Ellis suggests that electronic resources should support both types of chaining and can provide enhanced chaining support by indicating material on the database by a particular author, by providing a facility for identifying all types of (forward and backwards) referential connections to/from the material of interest and by supporting more advanced forms of citation chaining (such as bibliographic coupling – identifying common works that particular documents cite and co-citation searching – identifying pairs of highly cited papers). Basic support for forwards and backwards chaining is commonly supported by current electronic resources, however the type of advanced chaining support described by Ellis is less commonly supported.

- **Browsing** - *"Semi-directed searching in an area of potential interest"* (Ellis, 1989, p.179). The social scientists in Ellis's (1989) study achieved this behaviour by reading journal contents pages in a broad subject area, checking library periodicals and by browsing along the shelves of books or journals, whilst the social scientists in Meho and Tibbo's (2003) study also tended to browse online library catalogues, indexes and abstracts, references of materials that had been preciously found or read and web resources (although it is unclear what exactly these resources were). The physicists in Ellis et al.'s (1993) study tended to browse abstracts, whilst the chemists mentioned browsing bookshops, particularly to see what new textbooks had been published. Ellis suggests that support for browsing all types of information held on an electronic resource should be provided by Electronic resources. Support for browsing material in electronic resources is increasing, although we are yet to see many electronic resources that allow all of their content to be browsed.

- **Differentiating** - *"An activity which uses differences between sources as a filter on the nature and quality of the material examined"* (Ellis et al. 1993, p.179). Ellis (1989) explains that *"differentiating is effected by the researcher identifying different sets of sources in terms of the differing probability of their containing useful material"* (p. 190). The social scientists in Ellis's (1989) study employed three main criteria for differentiating material:
1. The substantive topic of study.
2. The approach or perspective adopted (for example differences in 'schools of thought' or methodological approach).
3. The quality, level or type of treatment (for example the reputation of a particular publication, the generality or technicality of how the publication deals with the topic and whether the publication is aimed at academics or practitioners).

Ellis (1989) suggests that *"the same criterion can be used in different ways, for example, the level of treatment criterion could be employed to exclude material too specialised, detailed or technical for the*

requirements of one person, or too general, popular, journalistic, or lacking in rigour for the requirements of another" (p. 191).

Like the social scientists, the physicists (see Ellis et al. 1993) also differentiated material by substantive topic, whilst the chemists differentiated not only by the three criteria above as used by the social scientists, but also by the author and language of the material. Meho and Tibbo's (2003) social scientists studying the topic of 'stateless nations' also differentiated material based on the identity or origin of the material. In Ellis and Haugan (1997) the similar behaviour of '**distinguishing**' was identified instead of differentiating. Distinguishing involves *"ranking information sources according to their relative importance based on own perceptions"* (p. 399). The engineers and research scientists in Ellis and Haugan's (1997) study distinguished between materials either informally through discussion and conversation or more formally by looking through tables of contents and abstracts of material. The engineers and research scientists deemed material published in recognised outlets to be particularly relevant. In a recent paper co-authored by David Ellis (see Ford et al., 2002), differentiating was defined as *"distinguishing between different sources of information on the basis of the nature and quality of the material examined"* (p. 732). As this recent definition includes elements of both the original definition of 'differentiating' from Ellis (1989) and the word 'distinguishing,' we can assume that 'differentiating' and 'distinguishing' are essentially the same behaviour, an assumption confirmed by Ellis himself (Ellis 2007, personal communication). Ellis suggests that Electronic resources can support differentiating by allowing users to specify preferences for sources that they think are most likely to contain material of interest and then using these preferences to restrict the search, to exclude certain sources or types of source from the search or to display material in order of source or source type. This, Ellis (1989) asserts, *"is far easier than attempting to deal with differences of approach and level of treatment directly"* (p. 193).

- **Filtering -** The *"use of certain criteria or mechanisms when searching for information to make the information as relevant and as precise as possible"* (p. 399). The engineers and research scientists in Ellis and Haugan's (1997) study displayed 'filtering' behaviour by restricting searches (for example through the use of keywords or date restrictions). Current electronic resources support filtering in a number of ways, all allowing users to refine or reformulate their searches and some allowing users to search within the current search results, perform field-restricted searches (including date restriction). Some electronic resources provide filtering at the document level by highlighting users' search terms in the document text, allowing them to restrict their searches to particular parts of the document or by providing advanced Boolean connectors to allow users to specify exactly how close together the search terms must be in the document text.

- **Monitoring** - *"Maintaining awareness of developments and technologies in a field through regularly following particular sources"* (Ellis and Haugan, 1997, p. 396). The social scientists interviewed by Ellis (1989) frequently used informal contacts to keep them abreast with developments. They also achieved 'monitoring' behaviour by regularly consulting sets of journals which are deemed to publish material of interest often and by regularly scanning book publishers' lists, reading reviews and checking new library acquisitions. The engineers and research scientists in Ellis and Haugan's (1997) study also achieved monitoring behaviour through participation in conferences and other international forums. Like the social scientists, the physicists, chemists, research scientists and engineers in the studies by Ellis et al. (1993) and Ellis and Haugan (1997) also used formal and informal information channels for monitoring, as did the social scientists in Meho and Tibbo's (2003) study. Ellis (1989) suggests that 'monitoring' can be supported in Electronic resources by allowing users to specify sources to monitor and automatically searching these sources when the user next logs in to the system or when the sources are next updated. Ellis also suggests that monitoring may be enhanced through the provision of an 'alerting' function, where sources of interest are brought to the attention of the user. E-Mail alerts are becoming increasingly common in current day electronic resources, with the systems providing the facility for users to save searches and have them automatically and periodically run by the system on their behalf (and any new results sent to them in an e-mail).

- **Extracting** - *"Systematically working though a particular source to identify material of interest"* (Ellis et al. 1993, p. 364). Ellis (1989) explains that for social scientists, *"the source may consist of a run of a periodical, a set of conference proceedings, a series of monographs, the contents of an archive, a*

collection of publishers' catalogues, or bibliographies, indexes or abstracts" (p. 198). Therefore 'material of interest' in this instance can be regarded as a relevant document, reference or abstract as opposed to the textual content within it. Extracting was also identified as a behaviour displayed by both physicists and chemists (see Ellis et al., 1993). However, it was not found to be a particularly significant activity for either group. In addition, 'extracting' was identified as a behaviour by Meho and Tibbo (2003), who found that social scientists could extract from either direct sources (such as books or journal articles) or indirect sources (such as bibliographies, indexes, abstracts, online catalogues). Ellis suggests that 'extracting' can be supported in Electronic resources by facilitating *"continuous movement through different source streams,"* (p. 200) ensuring that material in the electronic resource is *"recomposed more or less in their original form for searching purposes"* (p. 200). Many current day electronic resources store and display documents under hierarchies that mirror the paper-based form of the information (for example all the papers from a particular conference proceedings or journal issue are stored together to facilitate browsing as well as searching).

- **Verifying** - *"Checking the information and sources found for accuracy and errors"* (Ellis et al., 1993, p.364). This characteristic was identified in the study of physical scientists and Ellis et al. (1993) note that *"although similar activities were mentioned by some social scientists it was a very minor part of those activities and would have been subsumed under chaining"* (p. 364). For the chemists, verifying involved looking for errors – particularly typographical ones, such as errors in numerical data. One of these chemists chose to verify material only from sources that he deemed unreliable, for example textbooks and reviews. The engineers and research scientists in Ellis and Haugan's (1997) study also displayed 'verifying' behaviour, but rarely consulted original sources to check information for correctness unless they are using the information with a view to publishing a journal article or conference paper. Interestingly, although verifying was not identified amongst Ellis's social scientists, it *was* amongst the social scientists in the study by Meho and Tibbo (2003). According to the authors *"many participants wrote [in the e-mail interviews] about 'bias,' 'disinformation,' and 'lack of reliability and accuracy'"* (p. 582) of information sources. Electronic resources rarely support verifying behaviour.

- **Ending** - *"The assembly and dissemination of information or the drawing together of material for publication"* (Ellis et al, 1993, p. 365). This behaviour was identified in the information behaviour of the chemists in the study by Ellis et al. (1993), who sometimes needed to find further information in order to discuss their findings in light of published work. Similarly the engineers and research scientists in Ellis and Haugan's (1997) study sometimes carried out small-scale searches targeted toward specific unsolved questions or looked for newly published literature that might be useful when writing a conclusion. A similar behaviour, 'assembly and dissemination' was also identified in the study of English Literature researchers by Smith (1988) and reported in Ellis et al. (1993). It is important to note that all of the above studies give examples of 'ending' behaviour that involve *finding information* as part of the process of assembly and dissemination, which suggests that the scope of this behaviour as defined by Ellis is limited to activities associated with finding information towards the end of the writing process.

All of the studies reviewed in this section identified similar behavioural characteristics (although reports of the studies of Physicists and English Literature students in particular used different terminology to describe what Ellis et al., 1993, consider to be similar behaviours). Ellis (1993) attributes this difference in terminology to the fact that *"the later studies were not simple verification studies as, in each case, the models were developed from the data and then compared to the original model"* (p. 483). Ellis also notes that the models in each of the studies differed in detail somewhat, partly due to minor differences in focus between studies and partly from subject differences between groups. Overall, Ellis and Haugan (1997) deem the model to be *"quite robust in relation to the information-seeking patterns of scientists, engineers and social scientists in both an academic and an industrial research environment and over a period of time which has seen accelerating changes in the information environment itself"* (p. 402). Ellis does not make any stronger claims about how far his model can be generalised to other groups. Indeed Ellis (2006) suggests that *"the basic research design and methodology can be replicated for studies of other groups without presupposition as to the outcome"* (p. 140).

In 2003, Meho and Tibbo conducted a semi-structured e-mail and face-to-face interview study with another group of social scientists. They identified potential participants by extracting e-mail addresses from four

bibliographic databases, interviewing sixty participants by e-mail and five of them in person. Unlike with the studies by Ellis and his colleagues, the participants interviewed by Meho and Tibbo *were broadly diverse in terms of gender, rank, discipline, geographical location and research topics"* (p. 576). Meho and Tibbo found similar behavioural characteristics to those found by Ellis and identified three new behaviours displayed by their group of social scientists. Although Meho and Tibbo do not provide definitions of some of these behaviours, they describe their findings related to these behaviours in sufficient detail for us to infer the scope of these behaviours (and to briefly discuss how electronic resources currently or might in future support these behaviours):

- **Accessing -** Meho and Tibbo found that many of the social scientists that they interviewed faced access problems when attempting to access the material that they required. All of the quotations relating to accessing are related to physical problems accessing paper-based information, such as long travel distances and difficulty getting hold of published materials within certain countries. No mention was made of problems surrounding accessing electronic information or resources. It might be argued that 'accessing' behaviour was not identified in any of Ellis's studies due to the fact that his participants did not face such access problems. Current day electronic resources support accessing through two main methods; some offer automatic (and invisible) IP recognition to recognise the company or institution that the user is accessing the system from and grant them access to the system or materials within the system. Others rely on users logging in with a username and password. Others still rely on a mixture of the two, mainly in order to provide personalisation services to users.

- **Information managing** – Meho and Tibbo found that many social scientists spoke about *"filing, archiving, and organizing information collected or used in facilitating their research"* (p. 582). The quotations that the authors provide relating to information managing suggest that participants sometimes devise their own electronic or paper-based archiving systems of information. It might be argued that 'information managing' behaviour was not identified in any of Ellis's studies as, conceptually, it can be thought of as lying on the boundary between information-seeking and information use (and Ellis's studies excluded information use behaviours). This argument is supported by the fact that Ellis et al. (1993) identified a borderline information-seeking/use behaviour, 'information diffusing,' amongst physicists but excluded it from analysis. We believe that, since this behaviour has design implications for electronic resources, it is not necessary to draw a firm distinction (and cut-off) between information-seeking behaviour and information use behaviour. Current day electronic resources do not provide much support for information managing. Some, however, like the electronic legal resource LexisNexis Butterworths, allow users to save a customised list of sources within the library that they frequently use and keep a research trail in the form of a time-stamped record of their search activity.

- **Networking** – *"Characterised by activities associated with communicating and maintaining a close relationship with a broad range of people such as friend, colleagues and intellectuals working on similar topics, members of ethnic organisations, government officials and booksellers"* (p. 582). As highlighted earlier, many of the participants of Ellis's studies also displayed informal information-seeking patterns. Therefore this behaviour may have been subsumed under other behaviours in previous studies. Or put another way, networking may be a more prominent behaviour it its own right for social scientists studying the topic of 'stateless nations' and hence was identified in Meho and Tibbo's study. Meho and Tibbo assert that *"both networking and verifying could be facilitated by including the corporate source field and full contact information of authors in both indexing and abstracting services and in online catalogues"* (p. 586). Many electronic resources already include contact details for authors, however in our opinion this does not provide much direct support for either behaviour as it would only allow users to view the contact details of authors. Electronic resources could further support networking by fostering informal communications within the library (for example, by facilitating blog-type discussions on papers or allowing the secure messaging of authors and researchers that are members of the library system).

Although Meho and Tibbo (2003) identified mostly the same behavioural characteristics as Ellis and his colleagues, they also present a summary model, placing each of the behavioural characteristics they identified under the four inter-related headings of 'accessing,' 'searching,' 'processing' and 'ending.' Whilst the behavioural characteristics that Meho and Tibbo identified emerged from an analysis of their data, it is not clear whether their subsumption model of Ellis's behaviours is also empirically grounded.

They describe the searching stage as *"the period where identifying relevant and potentially relevant materials is initiated"* (p. 584), the accessing stage as "the bridge between the searching stage and the processing stage, especially when indirect sources of information are used (e.g. online catalogues, indexes and abstracts and bibliographies, p. 586) and processing as *"the stage where the synthesising and analysing of the information gathered takes place"* (p. 586). Indeed, although Meho and Tibbo identify 'analysing' and 'synthesising' as separate behaviours, they do not make reference to these behaviours other than to stipulate that they occur during the 'processing' stage of information-seeking. Similarly, they make no reference to 'decision making' behaviour other than to suggest that it involves deciding whether to proceed to the process stage, or return to the searching stage.

The information-seeking behaviours that Meho and Tibbo (2003) claim each stage of their model subsumes are displayed in table 1. Interestingly, the behaviours of 'chaining,' 'extracting' and 'differentiating' appear in both the 'searching' and 'processing' stages of Meho and Tibbo's model, indicating some behavioural overlap. The authors do not, however, give any reason for their attribution of behaviours to each process stage and therefore no explanation for the overlap. A potential reason for the overlap might be that searching for information and processing the information found are two highly-related activities and therefore it would be difficult to classify behaviours such as 'chaining,' 'extracting' and 'differentiating' as behaviours that occur during searching but not processing or vice versa.

Meho and Tibbo's (2003) Process Stage	Behaviours Meho and Tibbo claim are subsumed in process stage
Accessing	**Decision-making**
Searching	Starting
	Chaining
	Browsing
	Monitoring
	Differentiating
	Extracting
	Networking
	Chaining
Processing	Extracting
	Differentiating
	Verifying
	Information managing
	Synthesising
	Analysing
	Writing
Ending	

Table 1: Meho and Tibbo's (2003) process model of information-seeking and claimed subsumed behaviours. Behaviours in bold were identified in Meho and Tibbo's model, but not discussed in detail by the authors.

A study of technical support workers by Cunningham, Knowles and Reeves (2001) also resulted in the identification of a number of information behaviours that *"correspond well with the framework described by Ellis,"* (p. 196). These behaviours included browsing, verifying (which they also refer to as cross-checking), filtering, monitoring and extracting. The authors also mentioned (but did not elaborate on) the behaviours of 're-presenting and re-structuring information,' which they suggested could be better supported in digital libraries through the introduction of annotation tools. The authors also suggested that existing digital libraries did not provide good support for browsing between and cross-checking sources.

Fit of Ellis's model (and subsequent refinements) to our summary information-seeking and use episode data

Due to the behavioural nature of Ellis's model, it does not cover many of the cognitive activities that *are* covered by the other models (this can be noted by the fact that the first one-and-a-half paragraphs of the text in the summary information-seeking and use episode in figure 6 have not been encoded under the model's categories). This is acknowledged by Ellis (2006), who highlights that this behavioural approach does not address cognitive or affective aspects of information-seeking. Aside from this, most (but not all) of Ellis's behavioural characteristics can be found in the remaining paragraphs. These missing characteristics were either behaviours that have not been identified in our study of lawyers (such as 'verifying') or behaviours that *have* been identified in our study, but are displayed in other information-seeking and use episodes and not this one (such as 'monitoring').

Potential for Ellis's model (and subsequent refinements) to inform the design and evaluation of electronic resources

As noted above, Ellis's model has been validated in a number of academic scientific disciplines, along with the non-scientific humanities discipline. It has also been validated outside of an academic discipline with a group of engineers and research scientists and been shown to remain applicable for helping to analyse the information-seeking behaviour of another group of social scientists, fourteen years after Ellis's original study of social scientists was published. This makes the model broad-scoped and highly generalisable. In addition, Ellis's model has a wide breadth of coverage (although interestingly the model does not directly cover searching activities, which can be noted by the lack of coverage of any of the search formulation and re-formulation activities in the first half of the second paragraph of text in figure 6). Whilst Ellis's model also does not capture cognitive behaviour (and therefore sacrifices coverage of the internal processes that the other models reviewed in this chapter cover), this is likely to have little impact on the model's potential to inform design and evaluation throughout the *observable* parts of the information-seeking and use process.

Ellis's model is presented at a lower level of abstraction than the other models we have discussed in this chapter (indeed Wilson, 1999, highlights this fact, stating that Ellis's model can be regarded as a deeper specification of broader information-seeking processes). Ellis (1989) asserts that *"it was this type of micro-level information about the activities and perceptions of the academic social scientists which was considered necessary for a detailed analysis of electronic resource requirements to be possible"* (p. 172). As the purpose of the model was to help define requirements for the design of interactive systems, it can be argued that an analysis of the finer-grained detail of the social scientists' behaviour is more likely to yield more concrete design suggestions and since the model is based on reasonably low-level information *behaviours* and not processes, designers only need to make a leap from the observed behaviour to the ways in which an electronic resource might support or better support the behaviour. Therefore the behavioural focus of the model minimises the potentially problematic creative leap that must be made in order to inform systems design and evaluation. Whilst it would be possible for an alternative model to analyse information behaviour at an even lower-level of abstraction (and thus finer grain of detail), it is nonetheless relatively easy for an electronic resource designer to ask himself how he can design or better design to support any of Ellis's behaviours (as we have illustrated in our discussion of each behaviour, which included Ellis's (1989) design recommendations based on many of his behaviours and our brief comment on how electronic resources currently support each behaviour). We therefore deem Ellis's model to be highly suitable for informing the design or evaluation of electronic resources in general and now briefly discuss the model's potential for helping us to inform the design and evaluation of resources in the specific domain of law.

Thomas, a Trainee Lawyer, is currently writing a research note for his boss, an Associate within the firm, on a topic within UK Competition Law. This is a topic that might affect one of the firm's clients. Thomas wants to find out about 'competition disqualification orders' (CDOs), which he knows are mentioned in the UK Enterprise Act (2002). However Thomas is aware of the nature and scope of the Act and believes that it does not cover the precise issue relating to CDOs that he is interested in. Thomas is aware that the Office of Fair Trading can apply a CDO, which disqualifies a Company Director for fifteen years, to the director of a parent company where the subsidiary company has been in breech of Competition Law. However, Thomas is unsure whether a CDO can be applied to a director of a parent company when the subsidiary company is not based in the UK. Therefore Thomas decides to look for UK Case Law to try to find out the answer to this question.

Thomas decides to look for the Case Law online, and logs into the LexisNexis Butterworths digital law library by typing in a username and password. Thomas clicks on the 'cases' tab and conducts a broad search for 'competition disqualification order.' Upon receiving no results, he reformulates his search to 'director disqualification order,' as he is aware that they were referred to as DDOs rather than CDOs before the UK Enterprise Act (2002) came into force. This time, Thomas retrieves a set of results and scrolls down the list. He reads the headnote of the case and briefly skims the decision to see whether the case might relevant to his research note. Thomas decides that this particular case is not relevant because it does not refer to subsidiary companies.

Thomas reformulates his query a couple more times, expanding and restricting the query by date where he deems necessary (as he believes that cases from the year 2002 onwards will be most useful as that was when the UK Enterprise Act, which discusses CDOs, came into force). However, Thomas does not find any results that he deems useful for helping him to answer the question. He therefore decides to look at the online version of Halsbury's Laws (a commentary source within LexisNexis Butterworths) in order to gain an overview of CDOs and hopefully a springboard from which to explore any relevant Case Law dealing with the topic. Thomas reads the definition of CDOs provided by Halsbury's and briefly follows a hyperlinked reference to the definition of 'company.' However, Thomas does not find any mention of Case Law that might be relevant. Thomas also looks at the table of contents of Halsbury's Laws, presented in a menu bar on the left-hand side of the screen, and selects the neighbouring commentary piece on 'competition undertakings,' which he thinks might possibly contain some relevant information as it is also in the 'Competition Law' section of Halsbury's.

Thomas returns to his results list and clicks on a title of a commentary article from a different source called 'Buckley on the Companies Act.' Thomas scrolls through and skim-reads the text aloud… 'a parent company may be accountable for a breach of one or more of its subsidiaries on the basis that the subsidiary has no real independence from the parent and the parent and subsidiary constitute a single economic entity. In such cases, the OFT will apply a five-step process,' which is how they decide whether to disqualify a director. They'll apply it to the directors of the company who committed the breach, that'll be the subsidiary company, and the 'directors or officers of the parent who are de facto parents of the subsidiary.' The information still does not help Thomas with his specific question about a *non-UK-based* subsidiary. Since Thomas has previously read the Office of Fair Trading's guidance to the UK Enterprise Act (2002) and did not find any mention of foreign subsidiaries (and his Case Law search had received no results), he concludes that his research question has not yet been tested in an English court. He therefore proceeds to write up the research note which involves summarising the task he had been given, the sources he has consulted, the applicable Law and how it applies to the case at hand. Once Thomas has written the research note, he e-mails it to his boss and uploads it to the firm's internal document management system, assigning a document number to the note so that any interested colleagues can search for it in the future.

Differentiating/distinguishing

Starting/surveying   Differentiating/distinguishing

Chaining
Browsing
Extracting

Differentiating/distinguishing

Ending

Figure 6: Summary of participant P10-T's legal information-seeking and use episode (observed during the think-aloud part of our Contextual Inquiry) encoded with the physical process stages from Ellis's behavioural model.

Potential for using Ellis's model to analyse the information-seeking behaviour of lawyers with the aim of informing electronic legal resource design and evaluation

Although no previous studies of lawyers have analysed observational data according to Ellis's model, the potential for doing so (and for using the model to provide design insights for electronic legal resources) has been noted by Sutton (1994), who suggests that *"both Lexis and Westlaw were designed with no apparent attention being paid to the information-seeking behaviour of attorneys"* (p. 198). Sutton highlights that a lawyer might use particular information sources as springboards to gain a basic outline of a particular legal area before moving on to more narrowly-focused sources. This is noted by Sutton to be equivalent to Ellis's 'starting' behaviour. Sutton also suggests that in order to make competent predictions and give informed advice, *"lawyers must engage in a focused, context sensitive, exploration of this legal area by following contextual clues from cases found when 'starting' and tracking the citations of useful cases"* (p. 194). This, as Sutton points out, is an example of 'chaining.' Using the behavioural model to inform design is, indeed, an approach that Ellis (1989) strongly advocates. He suggests that *"the general principle of using the behavioural aspects of users' information-seeking activities to inform the design of information retrieval systems…could play a more prominent role in the design of computer based information retrieval systems than, at present, it does"* (p. 202). Ellis also recognised that his behavioural model might be used as the basis for *evaluating* electronic resources, *"as an evaluatory tool to identify the existence and ease of implementation of features of the model in existing systems"* (Ellis 1987, p. 242).

In the remainder of this thesis, we move beyond Sutton's anecdotal examples by using Ellis's model to analyse the observed information-seeking behaviour of academic lawyers with the purpose of informing the evaluation and subsequent re-design of electronic legal resources. This serves to validate Ellis's model in a domain in which it has not previously been applied and to provide a much needed user-centred focus to the design of

electronic legal resources, by helping us to understand how lawyers work with existing electronic legal resources.

3.3.5 Summary

All of the models reviewed in this chapter are broad-scoped in terms of domains or work contexts that they are applicable to. Whilst some models (like Kuhlthau's ISP model and Ellis's behavioural model) had been empirically validated in various academic and practical domains and shown to remain applicable over time, others (such as those by Marchionini and Sutcliffe and Ennis) have not yet been empirically validated. Whilst this brings the generalisability of these models into question, it also highlights the potential for seeking to validate them empirically.

The models reviewed in this chapter differ somewhat regarding breadth of coverage. The main differences in breadth of coverage could be identified:

1.   In the coverage of cognitive processes.
2.   In the coverage of processes/behaviours that border information-seeking and information use.

As the models by Marchionini, Sutcliffe and Ennis and Kuhlthau all have a strong cognitive element, all three cover internal processes common to the early stages of information-seeking (such as recognising the need for information, stating their information needs and selecting a broad topic aimed at satisfying those needs). These cognitive processes are shown to the left of the vertical dividing line in figure 7, which is a temporal representation (i.e. linear comparison) of the models reviewed in this chapter. Figure 7 also summarises the fact that only the models devised by Kuhlthau and Ellis (and the subsequent subsumption model by Meho and Tibbo) include processes or behaviours that involve information *use* rather than information-*seeking*. This can be noted in the right-hand-most column of figure 7 which shows the similar behaviours of 'presenting' and 'ending.' It should be noted, however, that these behaviours and processes are only broadly comparable and, although they are presented linearly in figure 7, are *not* strictly linear. This is particularly important when examining some of Ellis's behaviours which are absent from figure 7 (such as 'browsing' and 'chaining'). These are behaviours that can occur at any time during information-seeking (see Ellis, 1989) and therefore cannot be shown in a linear representation.

The models reviewed in this chapter also differed somewhat in their depth of coverage (i.e. how concrete or abstract they were and far they sought to analyse as opposed to summarise information-seeking). The main differences in depth of coverage could be identified in whether the model:

1.   Chose to focus more on information-searching as opposed to information-seeking. As Wilson (2000) explains, *"information-seeking behaviour is more broad-grained than information searching behaviour, which focuses on the 'micro-level' of behaviour employed by the searcher in interacting with information systems of all kinds"* (p. 49).
2.   Presented higher-level processes, lower-level processes or behaviours.
3.   Presented cognitive or physical processes.

Using Wilson's (2000) distinction, the models by Marchionini and Sutcliffe and Ennis had a stronger focus on information-*searching* behaviour as opposed to information-*seeking* behaviour. This focus can be noted in figure 7 by the fact that the two models form two small clusters of processes focusing on the core searching activities of query formulation and execution and result evaluation/examination (and, in Marchionini's model, information extraction as well). These two models therefore can be deemed to cover information-searching activities in detail, but at the expense of other information-seeking activities (note the white gaps around the two clusters). Figure 7 also illustrates differences in level of abstraction. For example, Sutcliffe and Ennis's process of 'evaluating results' subsumes Marchionini's processes of 'examining results' and 'extracting information' whilst Meho and Tibbo's process of 'searching' subsumes Marchionini's processes of 'query formulation' and

'query execution.' Finally, differences in depth of coverage due to the coverage of cognitive as opposed to physical processes and behaviours can also be noted in figure 7. Although Kuhlthau's processes and Ellis's behaviours are not directly comparable, Kuhlthau's process of 'collection' is parallel to Ellis's 'extracting' behaviour. Indeed, it is reasonable to assume that while 'gathering pertinent information for the focused topic' (Kuhlthau's definition of collection), a possible behavioural activity might be to 'systematically work through a particular source to identify material of interest.' (Ellis's definition of extracting). However, as Kuhlthau's model is cognitively based, in order to design to support or better support 'collection,' it would be necessary to observe or otherwise probe details of the behaviours that underlie this cognitive process (and, of course, one of these underlying behaviours might well prove to be 'extracting'). The cognitive nature of Kuhlthau's model, combined with the relative high-level of abstraction that the model is presented at might help explain why the stages in the model are not directly comparable to those of the other models.
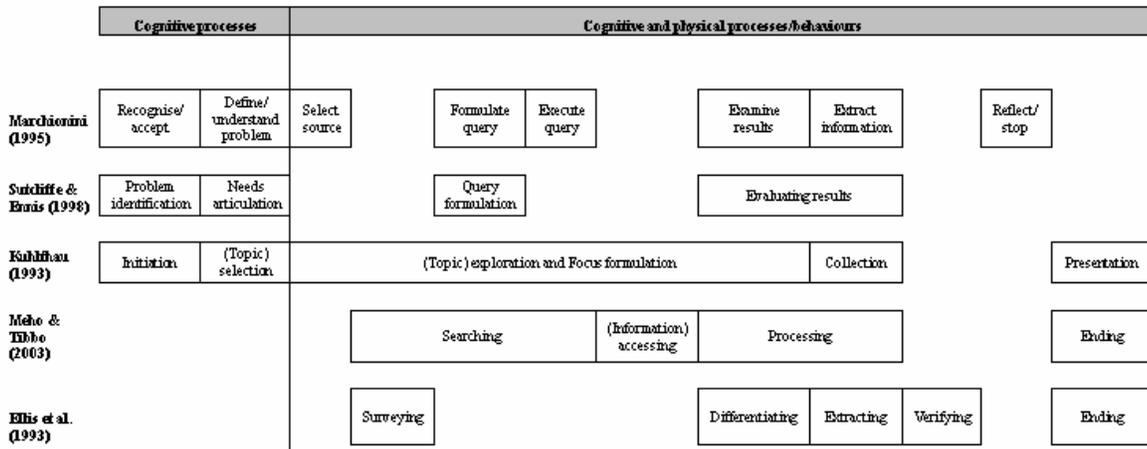
Figure 7: A temporal representation/linear comparison of the information-seeking models reviewed in this chapter. Note that processes/behaviours that did not correspond to those from other models (e.g. Ellis's 'monitoring' behaviour) are not shown. Also note that these behaviours and processes are only broadly comparable and although they are presented linearly, are not intended to illustrate a process that is strictly linear.

In summary, whilst we are not prepared to make the strong claim that any of the information-seeking models reviewed in this chapter are unsuitable for informing the design and evaluation of electronic resources, we certainly deem Ellis's model to be highly suitable for this purpose. This is for the most part due to the fact that it is based on concrete and observable behaviours as opposed to cognitive or physical processes that demand a further level of filtration to the underlying behaviour level so that they may truly inform design and evaluation. We also deem Ellis's model to provide good potential leverage for informing resource design and evaluation due to its wide breadth of coverage and due to the generalisability of the model, which has been empirically validated in several different domains. Through posing several design and evaluation-focused questions related to Ingwersen and Järvelin's (2005) dimensions of information-seeking models, we have selected a strong candidate model that we believe is likely to minimise the leap that a system designer will have to make between parts of the model, resultant user behaviour and interface elements aimed at supporting or better supporting that behaviour. Or to put it another way, we believe that if we were to ask system designers how we could design an electronic resource based on each of Ellis's behavioural characteristics, we would get a meaningful response in the form of a potentially useful set of design recommendations. Similarly if we were to ask them to evaluate the functionality and/or usability of the resource using the behaviours as a framework, we would expect to get some useful results. Indeed, in the remainder of this thesis we describe an empirical study that resulted in the emergence of a number of information behaviours (many of which were previously identified in other domains by Ellis and his colleagues) and the feed-in of these behaviours into the development of two novel methods for evaluating the functionality and usability of electronic legal resources.

Chapter 4: Methodology

This chapter at a glance…

In this chapter we:

- Discuss theoretical influences on our naturalistic study of academic and practicing lawyers' information behaviour.
- Discuss our data collection and analysis approach in detail.

4.1 Overview

In this chapter, we discuss the data collection and analysis relating to our study of academic and practicing lawyers' electronic information behaviour. In order to capture lawyers' information behaviour, we conducted in-depth naturalistic interviews and observations, where the lawyers were asked to think aloud whilst using existing electronic legal resources to satisfy a current or recent information need. Probing questions asked before, during and after the observation provided further insight into their behaviour, helping us to gain an understanding of *what* the lawyers were doing when using the existing resources and *why*. The interviews and think-aloud observations were transcribed and analysed using an approach based on the open and axial coding elements of Glaser and Strauss's (1967) Grounded Theory, not with the aim of generating theory per se, but with the aim of identifying a set of behaviours that might inform the design and evaluation of electronic legal resources.

We begin the chapter by highlighting the aims and objectives of the study, followed by a discussion on the methods and approaches that have influenced our choice of methodology. We discuss each of these methods and approaches in detail, explaining why they were chosen to guide our methodology and how we adapted them to fulfil our aim of examining lawyers' information behaviour in order to inform design and evaluation. The final section of the chapter includes a detailed discussion of our data collection and analysis approach. This serves as a practical discussion and justification of the methodological choices made and includes a detailed explanation of our choice of sample, sampling technique and setting, our recruitment approach, how we carried out the in-depth interview and think-aloud parts of our study and how we analysed and transcribed the resultant data. In this section we also briefly discuss ethical considerations surrounding the study, focusing in particular on privacy-related issues.

4.2 Study aims and objectives

This study aimed to examine the electronic information behaviour displayed by lawyers. This aim was driven by the motivation that in order to ensure that interactive systems truly support their users, it is necessary to gain a detailed understanding of how people use existing systems as part of their everyday work. It is possible to either feed this understanding into the design of new systems (or re-design of existing systems) or into the evaluation and subsequent re-design of systems so that they can better support users and their work. In this thesis, we do the latter – feeding the findings arising from the study described in this chapter into two novel methods for evaluating electronic legal resources.

In order to gain an understanding of how lawyers used existing electronic legal resources, we observed their behaviour when using these resources and asked them questions to probe further details about the behaviour displayed. We did not have a prior aim of identifying a set of information behaviours per se. (our aim was to investigate lawyers' information behaviour in general). Indeed, our study resulted in rich qualitative data that

had the potential to be analysed with a number of theoretical 'lenses.' As our data collection and analysis progressed, it became clear that framing our findings in terms of a set of information behaviours had the potential to provide a useful lens on the data for informing the design and evaluation or electronic resources.

4.3   Theoretical influences on our approach

4.3.1   Overview

Our data collection and analysis was informed by three methodological approaches: Strauss and Corbin's (1998) Grounded Theory methodology, Beyer and Holtzblatt's (1998) Contextual Inquiry method and Ericsson and Simon's (1984) work on verbal protocol analysis. Our methodology encapsulates many of the core principles of these methods. However, we have also made several practical decisions that distinguish our approach from these existing methodologies. In this section, we discuss our adoption and adaptation of key aspects of the methods for the purpose of understanding information behaviour to inform design and evaluation. This involves a detailed discussion of which aspects of each method we have adopted, in which ways we have deviated from classic versions of the methods, and the safeguards we have put in place for avoiding potential bias of data when adopting and adapting aspects of each method.

4.3.2   The influence of Strauss and Corbin's Grounded Theory on our approach

As we have reported in the previous section, our theoretical basis for analysis was Glaser and Strauss's (1967) Grounded Theory. Grounded Theory (which split into the Glaserian and Straussian approaches after disagreements between the original authors, Glaser and Strauss) has been found to be useful in several studies that are similar in nature to our own. This includes empirical studies of information behaviour (see Ellis, 1993) and of the information-seeking process in general (see Kuhlthau; 1988, Adams and Blandford; 2005).

Strauss and Corbin (1998) describe grounded theory as *"theory that was derived from data, systematically gathered and analysed through the research process"* (p. 12). It is 'grounded' in the sense that the theory is heavily rooted in the data and emerges through the process of cyclic data-gathering and analysis (i.e. *"analysis begins with the first interview and observation, which leads to the next interview or observation, followed by more analysis, more interviews or fieldwork, and so on"* (p. 42). We adopted this cyclic approach to data collection and analysis, allowing us to test hypotheses as they were emerging from our data and fill any theoretical gaps, hence allowing us to develop increasing confidence in the 'story' that we had interpreted from our data through the process of analysis. This cyclic approach is known as the 'constant comparative method' and is a key tenet of Grounded Theory.

We used the Open and Axial coding elements of Grounded Theory as a high-level data collection and analysis approach, helping us to identify categories of behaviour and how they relate to one another. As explained earlier, the process of open coding involves assigning codes to parts of transcripts that seem to describe similar phenomena and the process of axial coding involves relating the codes that have been identified from the data to each other. However, according to Strauss and Corbin (1998), it is common for Grounded Theorists to undertake a third stage of coding, 'selective coding' – defined as *"the process of integrating and refining the theory"* (p. 143) and is achieved by relating all categories of code to a central 'core' category.

By only conducting the first two stages of Grounded theory, we effectively 'stopped short' of theory generation. This was a deliberate choice, made for practical reasons of focus. Whilst it was necessary to relate these behaviours to each other in a broad sense (i.e. to define the scope or theoretical boundary of each code and to ascertain how the behavioural characteristics might be theoretically related), it was not necessary to relate each category of behaviour to a *core* category. That is not to say that a different study, with a primary aim of generating a Grounded Theory of how information behaviours are related to one another, might not have pursued

the full Grounded Theory approach. However, our aim was not to generate theory per se, but to identify distinct types or categories of information behaviour that could be used as springboards to inform the design and evaluation of electronic resources. This focus on design dictated the need to stop short of the Selective coding stage. This is a decision that also appears to have been made by Ellis (1987) when identifying the information behavioural characteristics of different groups of scientists. Like Ellis, we chose not to 'overshoot' with relation to coding (i.e. to conduct the 'selective coding' stage of Grounded Theory). This was because we did not believe that selecting a 'core' behaviour that could act as a main focus for a story about information behaviours would add any value to the analysis process. In addition we had followed the open and axial coding procedures (and the other key parts of Grounded Theory) rigorously and also found that our findings related closely to those from previous research on other disciplines. Therefore we deemed that selective coding would not provide us with any additional confidence in our theoretical outcome.

Safeguards for avoiding data bias based on Grounded Theory recommendations

Strauss and Corbin (1998) describe six key techniques to avoid the introduction of bias during data analysis:

1. **Think comparatively** – compare instances of coding in the data with one another. Strauss and Corbin also argue that researchers can think comparatively by turning to the literature or experience to find similar examples. They stress that *"this does not mean that we use the literature or experience as data per se, [but instead] use the examples to stimulate our thinking about properties or dimensions that we can then use to examine the data in front of us"* (p. 44). We found this to be an important means of avoiding bias in our data (as discussed in the previous section).

2. **Obtain multiple viewpoints of an event (triangulation)** – this can be achieved by interviewing/observing people that might provide complementary theoretical angles on the same event or phenomenon. Strauss and Corbin also explain that multiple viewpoints can be gained by collecting data *"on the same event or phenomenon in different ways such as interviews, observations and written reports"* (p.44). In some respects, the interview and think-aloud parts of our methodology can be considered as multiple viewpoints, as both helped us to identify information behaviour but in slightly different and complementary ways. When originally deriving his behavioural model of information-seeking characteristics, Ellis (1989) considered using interviews and observations in a complementary fashion in order to triangulate his findings. However, he concluded that this was not feasible as *"academic researchers typically have to balance the demands of research with those of teaching and administration, and information-seeking for research or training represents only a part of those activities and is diffuse both spatially and temporally"* (Ellis, 1993). We found that by asking participants to step-though a recent information-seeking episode (if they did not currently need to perform any electronic information-seeking), we were able to use interviews and think-aloud observations as useful complementary approaches.

3. Strauss and Corbin also suggest the need to *"occasionally check out assumptions, and later hypotheses, with respondents and against incoming data; that is, simply explain to respondents what you think you are finding in the data and ask them whether your interpretation matches their experience with that phenomenon- and if not, then why"* (p. 45). As discussed in the previous section, feeding back data to participants (particularly to support staff) was found to be another important means of avoiding bias.

4. **Periodically step back from the data** and ask the questions 'what is going on here?' and 'does what I think I see fit with the reality of the data?' We found it necessary to refer back to the data on a regular basis which we regard as another cyclic process involved in a Grounded Theory approach. We often found that we could gain an increasing amount of confidence with theoretical findings by regularly stepping back and questioning our assumptions by making reference to the data, then feeding any discrepancies into the creation of a revised (and invariably richer) theoretical picture.

5. **Maintain an attitude of scepticism** – validate emerging findings against data from subsequent (or previous) interviews or observations. Strauss and Corbin explain that maintaining a sceptical attitude is especially important for researchers who use categories defined from the literature, because of the context-specific nature of code categories. They argue that concepts derived from previous studies might *"have some relevance or explanatory power for the present problem under investigation... however, their*

properties and how they are expressed might be quite different with a different set of data" (p. 46). We found this to be particularly relevant to our study, especially once we began to identify similarities between our coding categories and those of previous researchers in the area. For example, we found the boundaries as defined by Ellis's information behaviours to be somewhat fluid. This meant that it was all too easy to make subtly incorrect assumptions about the relationship of our data to previous literature (akin to comparing an orange with a clementine rather than another orange). An ongoing attitude of scepticism allowed us to re-visit assumptions we had made on an ongoing basis and, as with periodically stepping back from the data, turn any discrepancies from negative to positive by creating a revised and stronger theoretical picture.

6.   **Follow the research procedures** – Strauss and Corbin stress that *"although researchers may pick and choose among some of the analytic techniques that we offer, the procedures of making comparisons, asking questions, and sampling based on evolving theoretical concepts are essential features of the methodology"* (p. 46). Our study followed all of these essential procedures, constantly making comparisons between different parts of the data and asking questions surrounding the content and boundaries of our coding scheme. We also followed a theoretical sampling approach (although our theoretical sample was only partly evolving, as explained in section 4.3.2). We believe our mix of both non-evolving and evolving theoretical sampling can still be regarded as following the essential procedures of Grounded Theory as both sampling techniques aim to generate rich data that is both internally and externally valid.

4.3.3   The influence of Beyer and Holtzblatt's Contextual Inquiry on our approach

Beyer and Holtzblatt's (1998) Contextual Inquiry approach is a naturalistic methodology that involves the combined observation and interview of users and their everyday work, with the intention of feeding findings from these observations into the design of interactive systems. As Beyer and Hotzblatt describe, the core premise of the methodology is to *"go where the customer works, observe the customer as he or she works, and talk to the customer about their work"* (p.41) in order to gain a better understanding of the 'customer' (or user).

According to Beyer and Holtzblatt, observing participants carrying out their everyday work helps to reveal concrete details about the work, the structure implicit in the work and the important and less important parts of the work. It can also, according to the authors, help participants recall past instances of carrying out the work that are either similar or different to the current instance. This methodology has an inherent ethos that is similar to our own. It involves understanding users' work with the purpose of feeding this understanding into generic and specific recommendations for the design and evaluation of these resources. However, Contextual Inquiry is aimed at understanding users' *work*, not their interactive *behaviour*. This is a subtle, but important difference that is a symptom of the fact that Contextual Inquiries are intended to operate at a different level of abstraction to that required for our study (which aimed to examine lawyers' information *behaviour*, not their broader information *work*).

There are, however, many aspects of Contextual Inquiry which were valuable in shaping our approach. For example, Beyer and Holtzblatt (1998) suggest that in order to observe the 'common underlying structure' in participants' work, it is useful to interview and observe participants whose work is as different as possible. This was important for our study so that we could have confidence in the generalisability of the information behaviours we identified and this was one reason why we chose to interview and observe a broad cross-section of both academic and practicing lawyers. Also, like Strauss and Corbin, Beyer and Holtzblatt recommend an evolving and theoretical basis for choosing who to interview and observe next, allowing the changing focus to drive participant selection. This approach was useful in our study, although only to some extent. We found taking an evolving sample to be useful primarily for obtaining multiple viewpoints on our data. For example, in several cases, practicing lawyers made reference to gaining assistance from library support staff. This led us to include members of support staff in our sample. However, whilst it was important to ensure that we sampled a broad cross-section of electronic resource users so that we could identify as wide a range of information behaviours as possible, it was not as important to ensure that was an *evolving* sample.

Beyer and Holtzblatt explain that although every Contextual Inquiry interview is unique, each has four parts:

1. **The conventional interview** – Where the participant and researcher get to know each other by the researcher introducing his focus and asking background questions about the tools the participant uses and an overview of their role. As in a Contextual Inquiry, our introductory background and icebreaker questions were aimed at gaining an overview of the participant's place in the academic or legal practice chain as well as an idea of their use of electronic resources.

2. **The transition** – Where the researcher states the new rules for the contextual interview "*the [participant] will do her work while you watch, you will interrupt whenever you see something interesting, and the [participant] can tell you to hold off it's a bad time to be interrupted*" (p.65). The transition process from our interview questions to the think-aloud session was similar and involved informing participants of the think-aloud process and that the researcher would occasionally interrupt to ask questions.

3. **The contextual interview proper** – Where the participant starts working and the researcher observes and interprets. In a classic Contextual Inquiry, this involves the participant carrying out their day-to-day work *without* explaining what they are doing and why, and the researcher intervening to ask questions in order to check understanding or probe for more detail. Our approach included a similar form of researcher intervention to ask questions and also involved the participant carrying out their everyday work using electronic legal resources. However, we asked participants *to verbalise their thoughts and actions* during the observation. This was deemed necessary for the practical reason of needing to gain a detailed understanding of lawyers' information behaviour. Whilst it may have been possible to infer aspects of what the lawyers were doing and why when using electronic resources, this was not desirable as there was considerable room for mis-interpretation and error when making inferences. Similarly whilst it may have been possible to save all questions until the end of the think-aloud part of the session, this might have lead to answers that were abstract, distorted, over-rational or even knowingly false due to the difficulty involved in remembering interface actions.

4. **The wrap up** – Where the researcher summarises to the participant what has been learned and what has come across as important. In our interviews, we did not wrap up with a summary that the participant should correct if necessary, but with some final questions to fill in any gaps in our interpretation of the verbal or interface-level data that they provided and to fill any gaps in our overall theoretical picture. We chose this form of wrap-up session mainly due to time-related issues, as it seemed a higher priority to ask questions where the meaning of the participant's words or actions were unclear rather than to summarise and feed-back an interpretation of these words and actions to the participant. This was partly due to the fact that the researcher commonly checked his understanding of participants' comments and interface-level actions at several points during the study.

Safeguards for avoiding data bias based on Contextual Inquiry recommendations

Beyer and Holtzblatt describe four key principles of Contextual Inquiry, which we have adhered to in our study:

1. **Context** – "*Get as close as possible to the ideal situation of being physically present*" (p. 47) in order to observe ongoing rather than summary experience and concrete data rather than abstract data. The authors explain that summary data, which is particularly common during retrospective accounts, can be avoided by the interviewer asking questions to "*fill in the holes*" in the participant's account. (p. 49). This was regularly conducted as part of our in-depth interview approach . The authors also explain that abstractions can be avoided by steering the participant back towards the current work. Whilst it was not necessary to do this very often, we found it to be a useful way of ensuring that concrete data was obtained (particularly during think-aloud observations where the participant stepped through a recent information-seeking episode and slipped into an abstract account).

2. **Partnership** – Strike an equal balance of power between the participant and the researcher. According to the authors, this can be achieved by alternating between watching and probing. For example the researcher might interrupt the work when he/she sees something that does not fit or identifies structure

within the work in order to ask relevant questions. However, he should then allow the participant to return to their work after these questions have been answered. Our think-aloud sessions followed a similar format, with researcher interventions kept to a minimum, but still used to probe details on participants' information behaviour. The authors also explain the need to avoid other relationship models, such as interviewer/interviewee (which they claim can result in the participant and researcher acting *"as though there were a questionnaire to be filled out,"* p. 55) and expert/novice (which they claim can result in the researcher helping the participant use the system rather than learning from the problems that they face with it). We found that these two models could be avoided with reasonable ease. The interviewer/interviewee relationship was avoided by following a line of in-depth and reactive questioning, based on what the participant had just told the researcher and rarely having the need to consult a formal interview guide. The expert/novice relationship was avoided by the researcher presenting himself someone who was not an expert in either the use of electronic legal resources or in the domain of law itself. This model was only challenged on one occasion, when a participant faced technical difficulties asked the researcher what he thought the problem was. The researcher explained he was not particularly technical and, after the end of the session, returned to the problem with the participant and made potential suggestions as to the cause of the problem.

3. **Interpretation** – Assign meaning to the observation by determining what the participant's words and actions *"[imply] about work structure and about possible supporting systems"* (p. 56). Beyer and Holtzblatt argue that checking an interpretation with a participant will not bias the data, but allow the participant to fine-tune the researcher's interpretation as *"the statement that doesn't fit is like an itch, and [the participants] poke and fidget with it until they've rephrased it so it represents their thought well"* (p. 58). We found the same in our study. Often the participant would offer subtle corrections if our interpretation did not fall completely in line with his own. This occurred mostly when the researcher checked his understanding of the information-seeking context surrounding the think-aloud part of the session (i.e. what information the participant needed to find and why), and less commonly when we were feeding-back our understanding of behaviours that participants commented on during the interview part of the session.

4. **Focus** – Set a clear focus to steer the conversation towards the research questions of interest. The authors argue that a clear focus *"gives the interviewer a framework for making sense of the work"* (p. 62). However, the authors also argue for the need for the researcher to expand focus where necessary to avoid ignoring potentially useful data. Beyer and Holtzblatt (1998) suggest that Contextual Inquiry interviews can be kept on track by providing participants with a 'pithy focus statement' (p. 77) before the start of the interview. At the start of each naturalistic interview/observation, we informed participants that we were approaching the interview/observation from a user-centred perspective, aiming to understand how lawyers interact with the current electronic legal resources that they use in order to inform the design and evaluation of these resources. We found this approach to be useful, helping participants to become accustomed to the probing (or as Beyer and Holtzblatt refer to them, 'nosy') questions involved in understanding their use of the resources and to frame their responses to these questions according to their understanding of the research agenda.

4.3.4   The influence of Ericsson and Simon's Protocol Analysis on our approach

As highlighted by Boren and Ramey (2000), the most cited source for justifying the theoretical basis for think-aloud-based usability studies is Ericsson and Simon's (1984) seminal work on protocol analysis. Despite the high citation count, Boren and Ramey highlight that usability practitioners' work rarely conforms to Ericsson and Simon's theory and that many practitioners choose to relax the rules as stated by Ericsson and Simon.

Ericsson and Simon (1984) present a model of three different levels of verbalisations:
1. **Level 1 verbalisations -** Verbalisations that do not need to be transformed before being verbalised (e.g. sequences of numbers whilst solving a mathematics problem).
2. **Level 2 verbalisations -** Verbalisations that must be transformed (e.g. images or abstract concepts that must be transferred into words).

3. **Level 3 verbalisations -** Verbalisations that require additional cognitive processing beyond that required for task performance or verbalisation (e.g. due to filtering).

The authors argue that level 1 verbalisations cause minimal potential for biasing verbal protocol data, whilst level 3 verbalisations are not considered to be reliable data. They also argue that any researcher intervention in the verbalisation process (e.g. through asking questions) can turn subsequent verbalisations into non-reliable level 3 data.

Boren and Ramey (2000) highlight four key points from Ericsson and Simon's work which capture the spirit of their stance on protocol analysis:
1. **Collect and analyse only 'hard' verbal data** (i.e. do not elicit participant inference, introspection or opinion during the think-aloud).
2. **Give detailed initial instructions** for thinking aloud.
3. **Remind participants to think aloud** (in as short and non-directive a manner as possible).
4. **Avoid intervention** (such as asking questions or making comments, even if those questions or comments are 'neutral').

Our adoption of the think-aloud procedure differs significantly from Ericsson and Simon's work with regard to the issue of intervention, partly because it was only possible to gain a truly detailed understanding of participants' behaviour by asking opportunistic questions. In effect, this is an example of *"data that is not collectable under Ericsson and Simon's model"* (Boren and Ramey, 2000, p. 266). In our study, it was necessary to ask these questions either at or close to the time when the researcher wanted to understand *why* the participant performed certain actions at the interface level. Boren and Ramey suggest that this problem can be overcome by asking questions *after* the think-aloud session. However, this was not practical for the purposes of our study, as explained earlier. Boren and Ramey (2000) also highlight that designers often need to prompt for data about users' expectations, explanations etc. and that sometimes this Level 3 data can be valued over more procedural information, which we believe to also be the case when examining information behaviour. They imply the need for researchers to weigh-up their priorities with regard to deciding whether to stick to Ericsson and Simon's 'no intervention' rule. Boren and Ramey suggest that sticking to the rule may be wise for research questions concerning cognition. However we are not interested in 'the internal structures and processes that are involved in the acquisition and use of knowledge' (McGraw Hill Professional dictionary definition) per se, but the information behaviours performed by lawyers. In addition, Bainbridge and Sanderson (1995) highlight that there is no one accepted way of performing protocol analysis – no 'canon.' They suggest that, as we have done for our study, it is sometimes necessary to adapt existing methods to new situations and develop new methods entirely. In short, we consider our researcher interventions to be both necessary and practical in order to address our goal of understanding *behaviour* as opposed to cognition.

With regard to analysing verbal data, Ericsson and Simon (1984) describe a coding method that is in some ways similar and in other ways different to Strauss and Corbin's (1998) Grounded Theory methodology. For example both sets of authors highlight the need for the coding scheme that is generated to account for as much as the raw data as possible. In addition, the coding method described by Ericsson and Simon (1984) is similar to the 'open coding' process of Strauss and Corbin's (1998) Grounded Theory, whereby the protocol is segmented and codes assigned to each segment based on the content.

However, this is where the similarity ends. Although Ericsson and Simon discuss the identification of synonyms as part of the process of generating names of codes, they do not mention the refinement, combination or splitting of codes, as is the case when analysing data according to Strauss and Corbin's Grounded Theory. In addition, the two methodologies are fundamentally at odds with regard to code generation. Whilst Ericsson and Simon suggest scanning the protocol for a 'vocabulary' to be used as code names, they also suggest generating a vocabulary from a prior analysis of the task being observed (probably because they were aiming to understand cognition, not behaviour). This is against the spirit of Grounded Theory, which emphasises the emergence of code categories from the current *data* as opposed to from previously observed data or pre-conceived notions

about the task. As our study also holds the practice of 'listening to the data' in high regards, we decided to allow our coding categories to emerge from the data rather than pre-defining them.

Safeguards for avoiding data bias based on Protocol Analysis recommendations

Bainbridge and Sanderson (1995) highlight five types of potential distortion of verbal protocol data, which we have actively sought to safeguard against in our study:

1. Having to give a verbal protocol changes the task and may change the way the task is done.
2. Verbal protocols can be influenced by social biases, such as over-cooperation and trying to say what they think the experimenter wants to hear. If the person reporting feels the listener has a superior status, there may be pressures to present a non-habitual approach to the work, such as appearing rational, knowledgeable and correct or conversely being uncooperative and unforthcoming. To address these first two issues, we explained to participants at the outset of the need for naturalistic data that reflects their actual use of the electronic resources. Therefore participants were encouraged to approach the task in the normal way that they might. Although not always detectable, we did not notice any anomalies or inconsistencies that would suggest that any of our participants altered the way that they used the electronic resources for the benefit of the study.
3. There are time constraints when giving a verbal protocol. Many things may quickly pass through people's minds and can be forgotten before there is time to report them, there may not be time to mention everything that is relevant and they may not mention information which they collect while reporting other activities, which may lead to unexplained behaviour later. As the think-aloud tasks were not time constrained, it is unlikely that time constraints played a large role in our data collection. However, probing questioning by the researcher during the think-aloud part of the session addressed many of the issues.
4. Asking to give a verbal report when the person usually does the job in a non-verbal way may lead to distortion. This did not appear to be an issue for our participants, who all conducted the think-aloud process without difficulty.
5. Knowledge about the components, mechanisms, functions and causal relations in a machine (along with some other types of knowledge) will only be mentioned if the think-aloud task involves a problem that requires use of this knowledge. In the context of our study, we were faced with the similar issue that participants would only display particular information behaviour if the task they choose to perform facilitated the demonstration of this behaviour. In practice, this did not prove to be a serious issue. Even though participants were free to use any electronic legal resource or resources of their choice in order to step through their information-seeking episode, they usually chose to use a core set of resources that each allowed them to display a wide range of information behaviour.

4.3.5 Summary of the influence of Grounded Theory, Contextual Inquiry and Protocol Analysis on our approach

Our methodology was influenced by many of the core principles of Grounded Theory, Contextual Inquiry and Protocol Analysis. This is due to the fact that many of these principles are well-aligned to our goal of understanding lawyers' information behaviour to inform design and evaluation. The open and axial coding elements of Grounded Theory enabled us to 'listen to the data,' allowing the information behaviours we identified to emerge. In addition, adopting the essential procedures outlined by Strauss and Corbin (1998) of making constant comparisons, maintaining scepticism and theoretical sampling enabled us to question and revise the boundaries of the behaviours we identified and gain confidence in the internal validity of our data. The concepts of context, partnership, interpretation and focus from Contextual Inquiry were also important for helping us to achieve our goal of understanding behaviour to inform design and evaluation. These concepts helped us adopt an interview and observation approach that resulted in rich, concrete think-aloud data and a strong understanding of the information behaviour performed by participants. The concept of having

participants verbalise their thoughts and actions as part of our think-aloud observations also enabled us to gain an understanding of the information behaviours performed.

However, there were also aspects of these methodologies which we did not incorporate into our study in order to achieve our specific goal of understanding lawyer's information behaviour to inform design and evaluation. We felt that performing 'selective coding' would be of little value, since our aim was not to create a behavioural theory per se, but to understand the information behaviours performed by participants. This is why we chose to stop short of generating a theory by only performing 'open' and 'axial' coding. We also felt that observing participants performing electronic tasks without asking them to think-aloud would forfeit the chance to gain useful insights into the information behaviours displayed. This is why we asked participants to think-aloud whilst using electronic resources. We also decided that asking opportunistic, probing questions during the think-aloud sessions would help us gain further insights into participants' information behaviour. Finally, we decided it was more important to ask wrap-up interview questions in order to fill gaps in the behaviours displayed rather than to summarise and feed back our understanding to participants. A summary of how our study was informed by, and differs from, the classic approaches of Glaser and Stauss's (1967) Grounded Theory, Beyer and Holtzblatt's (1998) Contextual Inquiry and Ericsson and Simon's (1984) Protocol Analysis is presented in table 2. Table 2 also lists our justifications for the differences between our study and the classic approaches and the safeguards we have put in place to avoid bias based on recommendations from the authors of each of the methodologies.

Method and theorists responsible	How our study was informed by and differs from the classic method	Justification for differences	Safeguards for avoiding bias
Grounded Theory, Glaser and Strauss (1967)	Study followed the core approach of Grounded Theory, but stopped short of generating a 'theory' (i.e. open and axial, but not selective coding was undertaken).	The research motivation was primarily aimed at informing the design and evaluation of electronic resources, not to understand how aspects of participants' information behaviour relate to one another.	Lack of tie-in of categories to a single core-category requires more stringent internal validity in the early stages of Open and Axial coding. Our study steered towards saturating all categories of behaviours, by asking about the use of systems that facilitate certain behaviours.
Contextual Inquiry, Beyer and Holtzblatt (1998)	Participants asked to think aloud as they complete their work, in addition to answering specific researcher questions about the work. Wrap-up interview questions were used to fill gaps in the participants' account of their information-seeking episodes and in the overall theoretical picture, rather than to summarise and feed back our understanding to the participant.	An ongoing account of participants' interaction with existing electronic legal resources was needed in order to gain a truly detailed understanding of their information behaviour. It was deemed that gap-filling questions were more of a priority than questions checking understanding (particularly as the researcher checked his understanding at several points during the study).	Participants briefed with a 'pithy' focus statement of the project (as recommended by Beyer and Holtzblatt) in order to focus their accounts on *what* they were doing with the current system and *why*. Researcher intervened when participant went quiet or off on a tangent, asking 'what are you doing now?' or re-orientating them towards the task at hand.
Think-aloud (protocol analysis), Ericsson and Simon (1984)	Short probing questions during the think-aloud questions asked by researcher. Coding based on interpretation of *current data* as opposed to from previously observed data or pre-conceived notions about the task.	Researcher needed to reliably understand participants' behaviour (e.g. 'obvious' behaviour or reasons for actions) that would not be verbalised without intervention. A pre-defined controlled coding vocabulary is against the spirit and ethos of Grounded Theory, which promotes 'listening to the data' above all else.	Interventions kept to a minimum length. Researcher only intervened where participant actions required further explanation in order to be reliably interpreted. In-depth questioning was either covered in the preliminary or wrap-up questions.

Table 2: Summary of how our study was informed by and differs from the classic methods of Grounded Theory, Contextual Inquiry and Protocol Analysis, our justification for the differences and the safeguards we employed for avoiding data bias.

In summary, our methodology was influenced by some but not all aspects of these approaches in order to help us fulfil our goal of understanding lawyers' information behaviour and using this understanding to inform design and evaluation. We were not able to adopt any of these approaches in their entirety due to the fact that some aspects were not completely aligned to our goal. We wanted to identify (and define the boundaries of) categories of information behaviour. However, in Grounded Theory, the process of selective coding involves relating the code categories identified to each other in order to identify a 'core category,' hence generating a

Grounded *theory*. Therefore we 'stopped short' of generating a theory by only conducting the open and axial coding elements of Grounded Theory.

We also wanted to understand lawyers' information *behaviour*, not their *cognition* (which is the main focus of Protocol Analysis) or their broader *work* (which is the main focus of Contextual Inquiry). As we are unaware of any specialist methods aimed at examining users' interactive behaviour (as opposed to their work or cognitive processes), we devised our own approach, guided by many of the core principles of Grounded Theory, Contextual Inquiry and Protocol Analysis (i.e. those that *were* aligned to our goal). This proved to be a fruitful approach that yielded a large volume of rich data that has the potential to inform design and evaluation. In the next section, we discuss our data collection and analysis approach in detail.

4.4   Data collection and analysis approach

4.4.1   Overview of data collection and analysis approach

Our data collection and analysis approach focused on observing lawyers' information behaviour in as naturalistic a way as was possible within practical constraints. Our approach involved conducting in-depth interviews and think-aloud observations of both academic and practicing lawyers, with the aim of understanding their information behaviour. In this section, we discuss this approach in detail, including our choice of sample and setting, our recruitment approach, how we carried out the in-depth interview and think-aloud parts of our study and how we analysed and transcribed the resultant data. We also briefly discuss ethical considerations.

4.4.2   Choice of sampling technique and sample

When choosing a sample of participants for our study, we followed a mix of *traditional theoretical sampling* and *evolving* theoretical sampling. Our theoretical sample was *traditional* in the sense that, in order to interview and observe a broad cross section of users of electronic legal resources (and to uncover a broad range of information behaviour), it was necessary to ensure that a vertical slice of academic lawyers was taken, from first year undergraduate level to professor level. A similar vertical slice across practicing lawyers was necessary, this time from Trainee to Associate level (and including support staff). *Evolving* theoretical sampling was employed when it became clear that there were members of (often library) support staff that could provide a complementary theoretical perspective on the information behaviour displayed by the academic or practicing lawyers that they support.

We recruited twenty-seven academic lawyers, who were studying at a large London university and a nearby vocational Law college. This included taught law students at all levels of academia (1st, 2nd and final year Bachelor of Laws (LLB) undergraduates and Master of Laws (LLM) postgraduates) and PhD research students. In addition two vocational students were also interviewed and observed in order to complete the theoretical picture. One was studying a vocational Legal Practice Course (LPC), the other a Bar Vocational Course (BVC).

We also interviewed and observed research staff (some of whom were also involved in teaching).   None of the taught academic students had specialised in a single branch of law (they studied various law modules as part of their degrees). Some of these modules had a jurisdictional focus (e.g. modules focused on International Law) and others had a narrow domain focus (such as modules on Company Law). In contrast, PhD students and research staff specialised in one or more legal areas. These were wide-ranging and included Environmental Law, New Technologies and the law, Human Rights Law, Constitutional Law, EU Law, Labour and Employment Law, Corporate Insolvency Law, Foreign Relations Law, Financial Law and Air Law.

In order to provide a complementary theoretical perspective, we also interviewed five law librarians who were involved in training and supporting academic lawyers from these two educational institutions (two of whom

were also involved in training students from other academic institutions). Location restrictions prevented us from observing them stepping through a recent information-seeking enquiry from a student, although this did not prove to be a problem as all of the academic librarians confirmed that requests for in-depth electronic information-seeking assistance were few and far between.

The practicing lawyers who participated in our study comprised twenty-four lawyers and support staff, working for the London office of a large multinational law firm. These were split across the mainly contentious department of Dispute Resolution (whose staff work on cases where there are multiple parties and a dispute to litigate or resolve) and the mainly non-contentious Tax department (whose staff work on cases involving one or more corporations but no 'dispute' as such). Our sample included lawyers and support staff at all levels of the company hierarchy where the participants deemed that electronic legal information-seeking was 'at least sometimes an important part' of their work (i.e. Trainees, Associates, knowledge support staff). No Partners were included in our sample as time pressures made it difficult for them to commit to taking part (and after some e-mail exchanges with Partners, it became clear that Partners often delegated their information work to Associates or Trainees).

Because the Tax department was smaller than the Dispute Resolution department, the sample of participants from Tax was smaller than from DR (9 as opposed to 14). The sample in Tax only included Trainee and Associate lawyers, not any support staff, as there were far fewer members of support staff working for Tax. Unlike in DR, there were also no embedded LIS staff in Tax that conducted electronic research and only a couple of members of knowledge staff (roughly equivalent to the Practice Development Lawyers in DR). The sample in the Tax department was, however, sufficient as we had already achieved high level of theoretical saturation when interviewing and observing lawyers and support staff in the DR department. The purpose of splitting interviews and observations across the mainly contentious DR department and mainly non-contentious Tax department was not to see whether our findings generalised across different types of practicing lawyers per se, but to test our hypothesis that the information behaviour identified would be similar across different groups of academic and practicing lawyers (even if this behaviour varied in sophistication). This hypothesis was based on the premise that when conducting electronic research, practicing lawyers use many of the same electronic legal resources across departments and firms. Because these resources do not vary in functionality for different types of lawyers (even if they are personalised somewhat to meet a particular departments' needs), we hypothesised that the information behaviour that lawyers across departments display was likely to be similar, simply because it was constrained by the functionality provided by the resources used). This turned out to be the case and, upon observing the first few Tax lawyers, it soon became apparent that it would not be necessary to observe as large a sample of Tax lawyers as Dispute Resolution lawyers and support staff. The number of each type of participant that took part in this study are shown below in table 3.

Type of academic lawyer:	Number		Type of practicing lawyer:	Number
Taught students:			*Trainees:*	
LLB (Bachelor of Law) Undergraduates	9		Trainees working in Dispute Resolution department	4
LLM (Master of Law) Postgraduates	8		Trainees working in Tax department	6
Vocational (Legal Practice /Bar Vocational Course) Postgraduates	2		*Associates:*	
Research students:			Associates working in DR department	3
PhD Students	2		Associates working in Tax department	3
Teaching and research staff:			*Knowledge support staff:*	
Research Fellows	1		Practice Development Lawyers working in DR.	1
Lecturers/Senior Lecturers	4		Practice Development Assistants working in DR.	3
Professors	1		*Library support staff:*	
Library support staff:			Members of the Library and Information Services team embedded in DR.	2
Law Librarians from academic/vocational institution	3		*Legal support staff:*	
Law Librarians from other Institutions	2		Paralegals working in DR department.	2
	Total: 32			Total: 24

Table 3: Numbers of each type of participant that took part in our study.

The academic institution which most of the academic lawyers studied at or worked for was a traditional 'red brick' institution and the law firm for which all of the practicing lawyers worked shared a similar high status, rated as one of the leading law firms in the world. Therefore our sample of lawyers can be considered to be one of high-calibre lawyers. As with studies by Ellis on the information behaviour of social scientists, we did not seek to interview a 'contrasting' group of lawyers studying at non-red-brick universities or practitioners working for a smaller or lesser-known law firm. This was partly for the same reason that Ellis gives in his account of his methodology (see Ellis, 1987), that we were not seeking to make comparisons between the behaviour of different groups of lawyers, but to inform the design of electronic legal resources for these lawyers. It was also due to the fact that, since the behaviours that the lawyers displayed are in part constrained by the behaviours which electronic legal resources support and that only a core set of electronic resources are available to the UK legal market, it would seem sensible that any groups of lawyers with enough inherent variety would display similar behaviours.

In order to balance the need to interview and observe a broad cross-section of electronic legal resource users with the need to gain an in-depth understanding of the lawyers' information behaviour, we felt it necessary to continue slightly beyond what Glaser and Strauss (1967) describe as the point of 'theoretical saturation' of the data (where interviews and observations with further participants does not contribute much to the overall theoretical findings). This was because although it was possible to note when further participants are contributing less and less to the overall picture, it was not possible to predict exactly when theoretical saturation will occur. Therefore we continued slightly beyond saturation point to ensure that as many subgroups of lawyers as possible, across both vertical chains of academia and practice, were represented in the sample. In practice, this involved conducting two or three final interviews/observations to fill in any gaps in the breadth of sample coverage. Also regarding breadth of sample, it is important to note that (in a similar way to David Ellis' studies on information behaviour) we only sought to identify broad differences between groups of lawyers (e.g. academic lawyers vs. practicing lawyers, taught academics vs. research academics, contentious vs. non-contentious lawyers). We did not seek to identify differences between subgroups (such as between Associates working in DR and Associates working in Tax) as we deemed our sample size to be too small to make any such comparison useful or reliable.

4.4.3   Choice of setting

As this was intended to be a naturalistic study, participants were interviewed and observed in their place of work. For academic staff, this was often in their private office and gave access to a desktop computer - which allowed them to participate in the observation part of the study. Law students were interviewed and observed in a private room on campus which was equipped with a computer connected to the university network. Although this room was not located in the department of Laws, all computers on campus when connected to the university network allowed identical accesses to the same electronic resources (and the same access procedures for the resources applied throughout the university). Academic library staff were interviewed in the library itself, usually in a meeting room without access to a computer.

Practicing lawyers working in the Tax department were observed at their desks, using their own computers. Lawyers and support staff in the Dispute Resolution department were interviewed and observed in an office within their department, set up with a computer connected to the firm's network. This was because DR lawyers usually shared their office with others and we wanted to minimise disruption. As with the centralised coverage of the academic institution, this provided participants with access to all of the firm's electronic resources. These could all be accessed in the normal way, with the exception of LexisNexis Butterworths, which on many of the participants' computers was set to remember their username and password. To overcome this problem, participants who encountered password difficulties with this resource and claimed not to have such difficulties on their own machines were offered assistance to log in. We have no reason to believe that any of our chosen settings influenced participants' interview responses or behaviour.

4.4.4   Recruitment approach

Recruitment was primarily achieved by e-mail, although some telephone contact was also made with prospective participants. Academic students were sent a blanket recruitment e-mail and responded if they were willing and able to participate. They were offered payment at the standard UCL participant rate to take part in the study. A similar, but personalised, e-mail was sent to academic staff and practicing lawyers and support staff, explaining the purpose of the study and format of the session in more detail. The only pre-requisite for participation was that the use of electronic resources 'should, at least sometimes, be an important part' of their work. The law firm pre-authorised the contact with each practicing lawyer and member of support staff that was contacted. We do not believe that the academics' self-selection or payment, nor the law firm's pre-authorisation of each participant to be contacted, had any bearing on our results.

4.4.5   Process of in-depth interview part of our study

Interview questions were asked before and after the think-aloud observation. Before the think-aloud observation, academic lawyers were asked about their current place on the academic ladder and practicing lawyers were asked ice-breaker questions on the nature of their job role. These questions were aimed at helping to ensure that participants felt comfortable and at ease. Both academic and practicing lawyers were also asked questions about the amount of electronic information-seeking involved in their work, what electronic resources they use on a regular basis and what types of materials they use these resources to look for. These introductory questions lasted around ten to fifteen minutes and also doubled up as a means of highlighting aspects of legal information work that were not (or were only vaguely) supported by existing electronic resources.

As we became more familiar with the job roles and types of resources used and materials sought by the lawyers, this section of the study became shorter. After we had conducted all of the interviews and observations in the Dispute Resolution department of the law firm, we decided that asking these questions was no longer necessary and truncated the session to focus almost entirely on observing an information-seeking episode. In place of

background questions, the Tax lawyers were asked questions surrounding the context in which the information-seeking episode that was to be observed was taking place. These also acted as ice-breaker questions.

After the observation, the lawyers were asked follow-up questions that had a number of purposes: 1) to clarify any ambiguities surrounding the behaviour displayed, 2) to elicit further details about something the participant had mentioned during the pre-think-aloud questions or during the think-aloud observation itself and 3) to fill gaps in their observed behaviour by probing for information about behaviours that they may not have displayed.

Filling in gaps in observed behaviour also often involved steering the discussion towards electronic resources that had not been used or mentioned by the current participant, but had been by other participants. This was with the aim of identifying behaviours that participants performed when using other electronic resources than the ones used in their think-aloud session. Filling in gaps in observed behaviour often involved asking participants questions about parts of the information-seeking and use process that they had not mentioned or demonstrated during their think-aloud session. This was with the aim of identifying behaviours that were not currently supported (or were rarely supported) by existing electronic legal resources.

Finally, the process of gap-filling involved feeding back insights gained from the ongoing analysis of the data (i.e. insights gained from analysing the transcripts of previous participants). This was with the aim of both identifying behaviours that were supported by current resources but not demonstrated by the participant and behaviours that resources do not currently support. The process of feeding back findings to participants was not straightforward, as much care had to be taken to avoid introducing bias. This was particularly the case when seeking to test hypotheses through feeding back to participants. This always involved asking questions related to participants' comments or interface-level actions. These were usually comments or actions that highlighted similarities or differences with findings observed across other participants. For example, one practicing lawyer mentioned that the process of writing a research note involved taking the information found and '[paring] it down a bit' rather than simply 'regurgitating' it. Previous participants had described a similar process and, in order to consolidate his understanding, the researcher proposed the behaviour 'synthesising' and asked the participant whether it was a 'fair word to describe it?' In this case, the participant suggested the researcher had 'hit the nail on the head.' In other cases, participants corrected the researcher by discussing his question in more detail. Whilst the process of feeding back findings to participants has potential halo effect implications (which can only be minimised, by remaining sensitive to the issue, rather than eliminated altogether), it also serves to test emerging hypotheses and assumptions so as to maximise internal validity as much as possible. As participants seemed happy to correct as well as agree with the researcher, we believe this practice worked well.

Like the introduction questions, these follow-up questions also lasted around ten to fifteen minutes and just as with the introduction questions, as more lawyers were observed, the value of asking these questions reduced. By the last few lawyers, it was not deemed necessary to ask any follow-up questions.

4.4.6   Process of think-aloud part of our study

The think-aloud part of the study lasted around thirty to thirty-five minutes. Lawyers were set the broad task of looking for electronic information that they currently require as part of their work (i.e. to go about satisfying a current information need that involves the use of electronic resources). When participants could not think of a pressing information need, they were invited to step the researcher through a recent information-seeking episode. It was explained to participants that the aim of the study was to observe naturalistic behaviour and therefore they should only try to find information that they currently need or recently needed for their work and they should undertake this information-seeking in the way that they normally would. When stepping through a recent information-seeking task, it was emphasised to participants that it was more important to use the task as a springboard to observe them using electronic resources and therefore they should not attempt to re-construct the information-seeking episode in its entirety – what they actually did when looking for the information originally was not important. Approximately half of participants chose to satisfy a current need. The other half chose to step through a recent information-seeking episode.

Participants were first asked to describe the context of the information-seeking episode in detail, then to think aloud whilst using the electronic resource or resources of their choice (i.e. to 'explain what they were doing as they were doing it' and to 'verbalise any thoughts that were going through their heads'). The broad aim of understanding their use of existing electronic legal resources to inform design and evaluation was explained to participants and they were assured that the think-aloud observation was not a test of their information skills.

Participants were also informed that the researcher would occasionally ask them questions. These took the form of short and seemingly innocuous questions, posed at opportunistic moments during the study to probe participants' information behaviour in detail. This often involved simply probing the participant for more details, or more precise/specific details about their information-seeking and use behaviour. Questions were also posed during the observations when participants mentioned aspects of information-seeking that were not supported by current resources (mostly information use-related aspects, for which there was currently no direct electronic support). These questions were aimed at gaining an appreciation of the wider context of the information work that the participant was conducting. The most common question was 'what would you do next with this information, now that you have found it?' The only other intervention that the researcher made during the observations was to elicit concrete behaviour from the participant when only a verbal description of events was offered (often achieved by asking the question 'how did you go about doing that?') This was necessary as some participants, when stepping through a recent information-seeking episode, reverted to giving abstract descriptions of their actions, rather than concrete demonstrations of those actions. Overall, however, interventions were kept to a minimum and the researcher maintained a passive role for the vast majority of each think-aloud observation.

As an alternative to asking questions during the think-aloud sessions, we did consider asking participants to review their think-aloud sessions and interviewing them about the information behaviour they displayed. This may have avoided the need for researcher intervention during the think-aloud sessions (thereby reducing the possibility of the questions biasing the participants' future actions). However, as highlighted by Van Den Haak, De Jong and Schellens (2003), retrospective think-alouds might lead to participant bias (with participants concealing, inventing or modifying thoughts for reasons of self-presentation or social desirability). Eriscsson and Simon (1993) also emphasise that eliciting retrospective accounts can be time-consuming. Indeed most of the academic staff and practicing lawyers were only able to spare an hour of their time which, in our opinion, was not enough time to conduct an in-depth think-aloud session *and* re-play the session in order to ask questions. Faced with a choice between these two approaches, we decided to ask questions *during* the think-aloud, whilst taking as much care as possible to avoid biasing participants. This involved only asking questions related to the information behaviour displayed or comments made by participants and using careful discretion when deciding if and when to ask questions. As a general rule, questions were only asked if the researcher believed that this might result in the participant articulating details about their behaviour that they had not verbalised. In addition, questions were only asked at points during the session where it was felt they would not guide participants' future comments or actions.

On occasions, participants asked the researcher during the study whether they should proceed in a certain direction. The researcher responded by asking participants to only do so if they normally would as part of their work. Also occasionally, participants asked whether they should illustrate an earlier point that they had made with a practical example. The researcher responded on a case-by-case basis based on his opinion of the likely value of the example. The reason for making this decision on a case-by-case basis was to only allow examples to be presented where the practical element would be likely to illustrate any further information-related behaviour.

As the naturalistic task we set participants involved them '*finding* information,' we were aware that this might bias users towards displaying certain information behaviours but not others, for example behaviours related to information *use* as opposed to information *seeking*. This is one reason why we asked gap-filling questions, aimed at identifying any behaviours not mentioned or demonstrated during the think-aloud session, during our wrap-up interviews (see earlier discussion). We also seized opportunities to ask further questions where, as part

of the warm-up interview, participants mentioned information tasks that substantially differed from those undertaken by previous participants or resource functionality that had not been mentioned or used by previous participants.

We were also aware that by asking users to use *electronic legal resources*, the range of behaviours they could demonstrate were likely to be constrained by the behaviours supported by the resources they used. Therefore we provided participants with a free choice of resources to encourage demonstration of as broad a range of behaviour as possible (i.e. by giving participants the opportunity to use a wide range of resources). We also tried to overcome this limitation by asking relevant gap-filling questions in our wrap-up interviews and by asking opportunistic questions during the think-aloud session itself, centred on the role of electronic resources in lawyers' broader information work. During our questioning, we remained sensitive to the fact that the use of electronic legal resources was only likely to be one part of lawyers' broader information work. We also remained sensitive to the fact that understanding how lawyers' use of electronic resources fits in with their broader work context could lead to the identification of information behaviours that were not currently supported by existing resources. Whilst we considered observing lawyers' broader work tasks in order to identify behaviours not supported by existing resources, this would have required numerous, extended observations. The lawyers in our study often mentioned that their work involved unique time constraints and therefore we do not believe they would have consented to these types of observations.

Overall, we do not believe that restricting our study to observing the use of *electronic* resources to perform a broad task related to *finding* information is likely to have resulted in the identification of a less-than comprehensive list of information behaviours. This is because we remained sensitive to the implications of restricting the study in this way and used interview questions to gain as broad a picture as possible of lawyers' information behaviour. We remain open to the possibility, however, that observing lawyers' broader information work might lead to the identification of additional behaviours – particularly if the nature of this information work differs substantially from the work observed and mentioned by participants in our study. However, we do not expect this to be the case. This is because we have interviewed and observed a broad cross-section of lawyers who performed a range of different electronic information tasks and spoke about a wide variety of broader tasks. Therefore, we have every reason to believe that the behaviours we have identified (and the methods we developed that use these behaviours to frame evaluations of electronic resources in the legal domain) are comprehensive.

4.4.7   Process for analysis and transcription of interviews and observations

The sessions were audio recorded to enable transcription and detailed analysis of the verbal protocols. After each interview and observation was transcribed, the transcript was read sentence-by-sentence and coded in accordance with the 'open coding' and 'axial coding' elements of Grounded Theory in order to identify recurring behaviours and how they might relate to one another. Strauss and Corbin (1998) define open coding as *"the analytic process through which concepts are identified and their properties and dimensions are discovered in data"* (p. 101) and axial coding as *"the process of relating categories to their sub-categories, termed 'axial' because coding occurs around the axis of a category, linking categories at the level of properties and dimensions"* (p. 121). The coding process was achieved by coding parts of the transcripts that appeared to refer to the same behaviour with the same label and refining the analysis through a cyclic process of re-reading the data, re-naming codes (for example when a better or more precise description of the behaviour could be identified), merging codes (when two existing behaviours were deemed to actually be the same), splitting codes (when behaviours that had previously been coded under one code were deemed to actually be different) and by re-coding parts of the data under a different code name or unlinking data from a particular code (when data no longer appeared to fit the code name that it had been assigned to).

The analysis was 'grounded' in the sense that the findings emerged by 'listening to the data' as opposed to seeking to test existing hypotheses. As the study was quite narrowly focused on examining lawyers' information-seeking-behaviour, care was taken to avoid preconceived ideas from biasing the data. This was

particularly important during the coding process, especially as we were previously aware of several models of information-seeking and had already concluded in a previous literature chapter that Ellis's behavioural model was likely to be particularly suitable for informing resource design and evaluation (see chapter 3).

However, our coding process did not involve simply relating our data to different information-seeking models in order to identify which fitted the data best (which might be regarded as 'forcing' as opposed to 'emergence' in a Grounded Theory approach). Instead, the process involved detailed coding of the data using our own terminology and the similarities between our codes and Ellis's model emerged from the analysis. This led us to examine our data in the light of Ellis's model, asking questions of our data such as 'is the behaviour we have identified amongst lawyers similar to that found by Ellis and his colleagues?' 'What behaviours did Ellis and his colleagues find that are missing from our data?' and 'what additional behaviours have we identified in our study that have not been, to the best of our knowledge, identified previously?' To facilitate easy comparison with Ellis's existing model, we chose to use Ellis's existing code labels when we believed our data reflected identical (or highly similar) behaviour, rather than using different terminology. It is important to stress, however, that although the focus of the data *analysis* shifted from a broad investigation of electronic information-seeking to an investigation into the information behaviour*s* displayed by lawyers, the focus of the data *collection* was driven by the desire to understand information behaviour in a way that might inform the design or evaluation of electronic legal resources. It was not driven by the desire to validate or refine Ellis's model (although that was the end result). Our stance, therefore, was partly inductive and partly deductive.

Apart from examining the information behaviour displayed by lawyers, we also conducted several other analyses of our interview and observational data. These included examining the knowledge lawyers hold about the electronic resources they use and the information behaviours they display and an examination of the information expertise lawyers developed (see Makri et al., In Preparation). However, in this thesis, we focus on the findings related to lawyers' *information behaviour* (not their knowledge or expertise) as only the behavioural findings were fundamental to the development of the Information Behaviour methods (which are described in chapter 6).

Use of software to support the qualitative analysis process

To support the coding process, a qualitative research tool called Atlas.ti was used. At the most basic level, this tool allows Grounded Theorists to assign codes to parts of each transcript and then display parts of the transcripts that are related (i.e. share the same code) together. Atlas.ti was also used to support more complicated aspects of the research, such as probing broad differences between groups (e.g. taught vs. research academics, lawyers working within the mainly contentious Dispute Resolution department and mainly non-contentious Tax department etc.). This was achieved by splitting the transcript documents into 'families' of participants and filtering the display of data by family (as well as displaying quotations that share the same code together). Finally, Atlas.ti was used to help support the identification of co-occurring codes, in order to identify relationships between them. It is important to note, however, that whilst software tools aid Grounded Theorists with performing the mechanics of coding (and to some extent diagramming), they do not play any real part in the analysis proper.

Transcription annotation and naming conventions

Both the content of the interview and think-aloud observation were transcribed verbatim, with the letter R preceding researcher interventions and the letter P (and the anonymous participant number) preceding participant comments. Bold italics were used to denote when the participant emphasised a statement, whilst pauses of over five seconds were noted in square brackets at the appropriate point in the transcript, as was when the participant displayed noticeable signs of emotion. This was almost always laughing (usually as a response to the results of their interaction with the system). Also included in square brackets were the participants' interface-level actions, but no interpretation of those actions (for example a search where the participant receives no results due to incorrectly spelling a search term would be coded as: [participant enters the search terms 'undue inflence' into the search box and submits the search, but receives no results]).

In order to probe differences between broad groups of lawyers, each participant transcript was assigned to an appropriate family of documents. For example, the family 'first year students' contained the transcripts of the four first-year undergraduate law students and the family 'DR Associates' contained the transcripts of the three Associates working in the Dispute Resolutions department of the law firm. Transcripts were then filtered and analysed by family. Participants within each family were then numbered at random.

4.4.8   Ethical considerations

This study was undertaken with full ethical approval from the University College London Ethics Committee (application number 0698/001) and informed consent was gained from participants. On the informed consent form, participants were made aware that they were free to withdraw from the study at any time and without penalty and that they could request to review, edit or delete the written transcript of their interview/observation at any time. They were informed that the interview/observation session would be audio recorded and that the data arising from the study would be stored and disseminated in accordance with the Data Protection Act 1998. Practicing lawyer and support staff participants were also informed that the anonymised data would be made available to members of the Knowledge Management team at the firm.

Finally, both academic and practicing participants were informed that any data that could have been used to identify them (including the institution in which they work) would be anonymised from the outset. Rigour in the process of ensuring anonymity when coding data was ensured by following the following rules:

- The name of the academic institution was omitted, along with the name of any teaching staff. For the London office of the large international law firm, the name of the firm was omitted, along with the name of any in-house software, databases, tools etc. that might help to identify the name of the firm. These were substituted with phrases like 'an in house Knowledge-Management database' and 'the London office of a large international law firm.'
- Names of people were omitted, whether first or surnames. These were substituted with participant numbers or general statements such as 'one of my colleagues, my boss.'
- For the law firm, any name or company details that could help identify a client was omitted. These were substituted with general statements such as 'a large steel manufacturer.'
- Precise place names that could help identify the academic institution, a client or the law firm were omitted. These were substituted with more general place names such as 'London' or 'the North-East.'
- For the academic institution, names of departments and course names and for the law firm names of departments, teams or practice areas were not omitted. It was decided that omitting these details might obscure the data. This course of action was deemed to be permissible as the details would be unlikely to *directly* help to identify a participant and still requires some inference to identify him/her.
- The sex of the participant was not omitted. In some rare cases, the participant was referred to as 'he' or 'she.' In most cases the participant was simply referred to as 'the participant.'
- Names of events were not omitted, but any of the above details described in these events were omitted by default.

4.4.9 Summary of data collection and analysis approach

Our study comprised a set of naturalistic interviews and think-aloud observations of a vertical slice of both academic and practicing lawyers, who were chosen by theoretical sampling. At the beginning of the session, participants were asked ice-breaker and background questions and were then observed thinking-aloud whilst they either used electronic resources to find information they currently or recently needed as part of their work. Probing questions were asked during the observation to elicit details about participants' information behaviour that may have not otherwise been verbalised. To conclude the session, participants were asked tie-up questions surrounding the information behaviour that they had displayed. The interviews and think-aloud observations were then transcribed and analysed using the 'open' and 'axial' coding elements of Glaser and Strauss's Grounded Theory in order to identify categories of information behaviour and how they relate to one another.

Chapter 5: Findings and discussion on lawyers' information behaviour

<div style="border:1px solid black">

This chapter at a glance…

In this chapter we:

- Present the behaviour-related findings from our naturalistic study of academic and practicing lawyers' information behaviour.
- Discuss our findings in relation to previous information behaviour work.

</div>

5.1   Overview

In this chapter, we present findings from our study of academic and practicing lawyers' electronic information behaviour. The information behaviour that we discuss in this chapter was displayed through the use of a variety of electronic resources to find legal information currently or recently needed for their work. These resources ranged from general Internet search engines (primarily Google), to specialist legal search engines (such as FindLaw), to digital law libraries and citator services. These included LexisNexis Professional, LexisNexis Butterworths, Westlaw, Current Legal Information, Lawtel, Justis, Kluwer Arbitration, HeinOnline, Practical Law Company, and others. Lawyers also chose to use specialist legal websites such as the Office of Public Sector Information website (which publishes UK Bills of Parliament) and Hansard (which publishes transcripts of UK House of Commons debates). Many practicing lawyers (when looking for internal know-how) also chose to use a Knowledge Management database, which was designed in-house by the law firm at which the study took place.

We identified similar categories of information behaviour to those found by previous researchers. However, our work extends on previous findings both practically and theoretically in six key ways. Firstly we identified four broad over-arching categories that subsume several of the lower-level behaviours found by previous researchers such as Ellis and his colleagues. Secondly we identified two lower-level characteristics (updating and history tracking) which were particularly pertinent to lawyers and have not, to the best of our knowledge, been identified in previous studies of information behaviour. Thirdly, we identified several information *use* instead of information *seeking* behaviours. Fourthly, we identified several lower-level searching behaviours that are conceptually related to some of the cognitive search activities identified by Sutcliffe and Ennis (1998). Fifthly, we identified several subtypes of behavioural characteristics and finally, we identified several levels that each lower-level behavioural characteristic could operate at.

In the remainder of this chapter, we present a refined version of Ellis's model based on the information behaviours identified. We structure our findings and discussion according to the categories of behaviours identified (that together form the model). We begin by highlighting the ways in which our work extends both theoretically and practically upon previous findings in information behaviour research. We then present an overview of the broad high-level information behaviours that we identified in our study and the lower-level behaviours which each broader behaviour subsumes. In this overview, we briefly define the scope of each lower-level behaviour and outline the levels at which we found each behaviour to operate and introduce the subtypes of each behavioural characteristic that we identified (where applicable). We then present our findings, discussing each behavioural characteristic in turn with reference to excerpts from our transcripts that demonstrate the mention or display of each particular behaviour. We also, as part of this discussion, discuss our findings in relation to previous studies of information behaviour (and to Sutcliffe and Ennis's [1998] model of cognitive search activities).

5.2 Refined model of information behavioural characteristics

5.2.1 Introduction to our refined model of information behaviours

Our study identified similar categories of information behaviour to those found by Ellis and colleagues and by Meho and Tibbo (2003) when observing the information behaviour of academics from various scientific disciplines. These categories of behaviour were also similar to those identified by Smith (1988) and described in Ellis (1993) when interviewing a group of English Literature students and can be noted in the right-hand column of table 4. Our study serves to confirm the findings by Ellis and his colleagues in the new domain of law.

Aside from the similarities between the information behaviours displayed by the lawyers in our study and those displayed in other domains, our findings can be regarded as a theoretical (and partly practical) extension of previous findings in the following ways:

1. We identified the **broad overarching categories** of 'identifying and locating,' 'accessing,' 'selecting and processing' and 'distributing.' These categories subsume many of the behavioural characteristics that we have identified in our study (and those that have been identified in previous studies by Ellis and his colleagues and Meho and Tibbo).

2. We identified the additional lower-level behavioural characteristics of '**updating**' (ensuring a *current* understanding of amendments or changes to legal documents and content - i.e. whether a particular case or piece of legislation is *good law*) and '**history tracking**' (ensuring a *historical* understanding of amendments or changes to legal documents and content - i.e. an understanding of the treatment a particular case or piece of legislation has received over time). To the best of our knowledge, these characteristics have not been previously identified in information-seeking studies. We suggest that this is because these characteristics are particularly pertinent to the legal domain (and not necessarily to other domains).

3. We identified the additional lower-level behavioural characteristics of '**collating**' and '**editing**,' which also, to the best of our knowledge, have not been previously identified in information-seeking studies. We suggest that this is because we do not draw as firm a line between information-*seeking* and information *use* behaviour as other studies - Wilson (2000) highlights that information use behaviour *"consists of the physical and mental acts involved in incorporating the information found into the person's existing knowledge base"* (p. 50), whereas information-seeking behaviour *"is the purposeful seeking for information as a consequence of a need to satisfy some goal"* (p.49). We regard any behaviour that could inform the design and evaluation of electronic resources to be of interest.

4. We identified several lower-level searching behaviours of '**search query formulating**,' '**search query refining**,' '**search query reformulating**,' '**search query refocusing**,' '**search query spelling/syntax altering**,' '**search result sorting**' and '**search query/result recording**.' Similar behaviours have been identified in studies of information-searching behaviour, but have not been (to the best of our knowledge) included as components of information behaviour models. We suggest that this is because we do not draw as firm a line between information behaviour and information-searching behaviour - *"the 'micro level' of behaviour employed by the searcher"* (Wilson 2000, p. 49) as other studies.

5. We identified several **subtypes** of behavioural characteristics (presented in brackets next to many of the characteristics in table 4). These subtypes serve to draw distinctions between different (and, for the most part, mutually exclusive) types of each behavioural characteristic. For example, 'updating' behaviour can be performed *directly* by searching or browsing for documents and content and manually checking that the document or content within is up-to-date (i.e. current) or good law. Updating behaviour can also be performed *indirectly* by using an electronic citator service to check whether a particular document or the content within it is up-to-date or good law. Most of these subtypes, to the best of our knowledge, have not been previously identified in information-seeking studies.

6. We identified several **levels** at which the lower-level behavioural characteristics that were identified in our study can operate – at the resource level (i.e. at the level of the electronic resource itself), the source

level (i.e. at the level of an information source or sources within a particular electronic resource), the document level (i.e. at the level of a document or documents within a particular information source), the content level (i.e. at the level of content within a particular document) and the search query/result level. Four of these levels are illustrated in figure 8, which highlights that an electronic resource can contain many sources which, in turn, can contain several documents – each with content within them. For example LexisNexis Butterworths, a widely used electronic legal *resource* contains many *sources* ranging from different series of legal case reports and legal journal articles, to collections of different types of legislation such as Acts of Parliament and Statutory Instruments. Within each source are a number of *documents* (individual case reports, articles, pieces of legislation etc.), each with their own *content*. We do not illustrate the fifth level (the *search query/result* level) in figure 8, but it can be regarded as the means of bridging each of the other levels (i.e. searching for content that is held in a particular electronic resource). Note that, particularly in the Digital Library community, the word 'resource' is often used to describe both digital libraries themselves and electronic sources within a library (e.g. a particular journal series available within). In this article, we refer to 'resources' and 'sources' as separate and distinct entities.
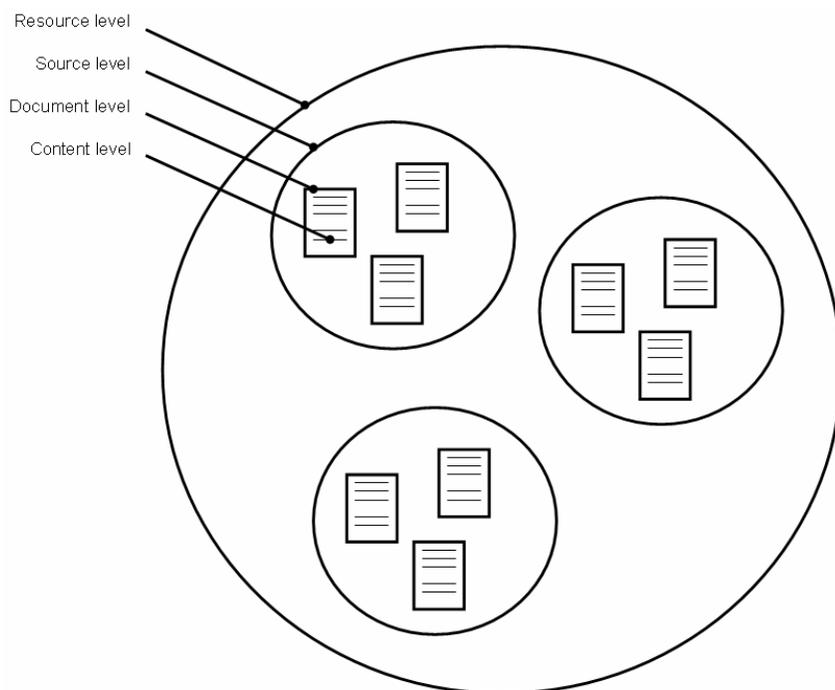
Figure 8: Diagram to illustrate four of the levels at which many of the information behaviours can operate.

We propose that many of the behaviours identified by Ellis and his colleagues can be performed at multiple levels. Some behaviours, such as surveying and monitoring, operate at a combined 'document and content' level. This level is used when it is difficult, impossible or undesirable to separate whether a particular behaviour is performed on the document, or the content within it. For example, surveying and monitoring involve gaining an overview of and maintaining awareness of developments in a particular research area. This involves looking at both documents and the content of those documents. However, the resulting observable behaviour at the document level is likely to be the same as the observable behaviour at the content level and therefore only one combined level is used to describe this behaviour. Whilst these levels may be applicable to paper-based information-seeking (for example, it is possible to regard a paper volume of journal titles as a 'resource,' an individual issue as a 'source,' an article within an issue as a 'document' and the textual content of the article as 'content') this has not been empirically tested as the focus of our study was on *electronic* information behaviour.

5.2.2   Overview of behavioural model

Table 4 provides an overview of the higher and lower-level behavioural characteristics identified in our study, which together form a refined model of information behaviours, partly based on the notion of subsumption (where broader, higher-level behaviours subsume particular lower-level behaviours). As with Ellis's original behavioural model, our model is not intended to be regarded as a *process* model of information-seeking, as the behaviours are not always performed in a linear fashion. Similarly, these behaviours are not entirely discrete (as certain behaviours can be facilitated through other behaviours or performed in parallel). For example, 'surveying' was often facilitated by lawyers searching or browsing for information, probably because this behaviour was not directly supported by many of the electronic resources that they used. In addition, we recognise that although there are numerous relationships between the behaviours identified, these are highly dependent on the information task at hand. Therefore it would be difficult to create an information *theory* based on these behaviours and the relationships between them.

In table 4, several lower-level behavioural characteristics (presented in the right-hand column of the table) are subsumed under each broader higher-level characteristic (presented in the left-hand column of the table). The notion of incorporating some of Ellis's behaviours under broader headings is not new. Indeed, Meho and Tibbo (2003) present a 'summary model' where they place each of the behavioural characteristics they identified under the four inter-related headings of 'accessing,' 'searching,' 'processing' and 'ending.' In table 4, subtypes of behaviours are presented in brackets and those behaviours which might be *expected* to be included under a particular broader heading but were not empirically observed in our study are presented in bold italics. Shaded behaviours are those that have not, to the best of our knowledge, been identified in previous studies of information behaviour.

Table 4 lists the four higher-level behaviours of 'identifying and locating,' 'accessing,' 'selecting and processing' and 'distributing' and the five levels we found them to operate at. In turn, the four higher-level behaviours (presented in the left-hand column of the table) subsume a number of lower-level behaviours (presented in the right-hand column) - many of which have been identified previously by Ellis and his colleagues and by Meho and Tibbo. For example, the higher-level behaviour of 'identifying and locating' might, in turn, involve performing surveying, monitoring, searching, browsing, chaining, extracting or a mixture of these. Similarly, 'selecting and processing' might involve performing 'filtering,' 'selecting,' 'distinguishing,' 'extracting,' 'recording,' 'updating,' 'history tracking,' 'analysing,' 'synthesising,' 'collating,' 'editing' or a mixture of these.

The higher-level behaviours of 'identifying and locating,' 'accessing,' 'selecting and processing' and 'distributing' are empirically grounded in our data. Hence, although the description of the information behaviours in this chapter is presented at a fine level of detail (i.e. by framing our discussion around the behaviours that do not subsume other behaviours in table 4), it would be equally possible to frame a broader discussion around these found higher-level behaviours in order to tell a similar, complementary 'story' about legal information behaviour.

It is important to note that as we observed lawyers' use of *existing* electronic resources, we cannot make any strong claims that these behaviours are exhaustive (particularly for aspects of legal information work that do not involve the use of existing resources). This is because by examining lawyers' use of existing resources, our findings were biased towards *identifying behaviours that were currently supported by* t*hose resources*. That said, we were aware of this limitation and used interview questions and questions at pertinent points during the observations to try to place the information work being demonstrated in a wider context and thereby highlight aspects that were not currently supported, or were only partially or indirectly supported by existing resources.

Higher-level behavioural characteristic and subtypes:	Related lower-level characteristic (and subtypes):
Identifying and locating – R, S, D, C D&C	Surveying (lightly/heavily directed) – *R*, *S*, D&C
	Monitoring (active/passive) – *R*, S, D&C
	Searching – R, S, D, C. *Lower-level searching behaviours not shown in this table.*
	Browsing – R, S, D, C
	Chaining (forwards/backwards), (across resource /within resource), (direct/indirect) – *R*, *S*, D&C
	Extracting – R, S*,* D
Accessing (direct/indirect), (visible/invisible) – R, *S*, *D&C*	
Selecting and processing – R, S, D, C, D&C, Q	Filtering – S, D
	Selecting (direct/indirect) – R, S, D
	Distinguishing – *R*, S, D
	Extracting (direct/indirect) – C
	Recording (manual/automatic) – R, S, D&C, Q
	Updating (direct/indirect) – D&C
	History tracking (direct/indirect) – D&C
	Analysing – C
	Synthesising – C
	Collating – D, C
	Editing – C
Distributing – *R*, *S*, D&C, Q	

Key: R= resource level, S= source level, D= document level, C= content level, D&C= combined document/content level, Q= search query/result level. Theoretical levels (i.e. those which we believe the behaviours can be observed to operate at, but have not been observed in our study are presented in bold italics). Shaded behaviours have not, to the best of our knowledge, been identified in previous studies. Extracting is presented twice as document extracting involves *identifying and locating* documents from sources (and similarly resource/source extracting involves identifying and locating resources/sources). Content extracting, on the other hand, involves *selecting and processing* content from documents.

Table 4: Summary refined model of information behaviour identified in our study along with the levels that each behaviour was observed to operate at.

We now provide an overview and description of each of the behavioural characteristics identified in our study (and presented in table 4), highlighting at which levels our groups of academic and practicing lawyers performed these behaviours and the different subtypes relating to each behaviour that were identified. Our refined model consists of four broad overarching categories - 'identifying and locating,' 'accessing,' 'selecting and processing' and 'distributing.'

The letters next to each overarching category (and the individual behaviours they subsume) indicate the levels at which they were found to operate (R= resource level, S= source level, D= document level, C = content level, D&C = combined document and content level, Q = search query/result level). Letters in bold type indicate levels at which a behaviour might operate *in theory* (but where no evidence was found in our study to support this hypothesis). Although, for example, users demonstrate different behaviour when selecting and processing *resources* ('resource selecting and processing') than when selecting and processing *documents* ('document selecting and processing'), the behaviours are nonetheless conceptually similar and therefore are subsumed under the broader category of 'selecting and processing.'

Several lower-level behaviours are subsumed under the first broad category of 'identifying and locating' resources, sources, documents and content:

- We identified Ellis's behaviour of 'surveying,' which was found to operate at the document and content level. This behaviour is identical in nature to Ellis's original behaviour and, according to Ellis and Haugan (1997), involves the "*initial search for information to obtain an overview of the literature within a new subject field, or to locate key people operating in this field*" (p. 395). This behaviour involves undertaking the physical activities involved with gaining an overview of a research area, such as conducting searches or browsing key sources for important documents in the area (which might also highlight important authors in the field). We found that surveying could be either *lightly directed* or *heavily directed*. Lightly directed surveying involves having few or imprecise details about the area on which an overview is required. Heavily directed surveying involves having many or specific details about the area (such as a lead from a colleague in order to get started in the area).
- Ellis's behaviour of 'monitoring' was also identified and found to operate at the source level and document/content level. We define monitoring as *"maintaining awareness of developments in a field."* This is highly similar to Ellis's original behaviour which involves "maintaining awareness of developments *and technologies* in a field*"* (Ellis and Haugan, 1997, p. 396). We adapt Ellis and Haugan's original definition because the lawyers in our study only maintained awareness of particular legal areas, not *technologies*. Whilst Ellis and Haugan's definition also suggests that monitoring behaviour is achieved by "*regularly following particular sources*" (p. 396), the lawyers in our study displayed monitoring behaviour at both the source level described by Ellis and his colleagues (for example by manually searching or automatically scheduling updates when a new journal issue has been published) and at the document and content level (by manually or automatically searching for recent or updated documents on a particular topic). Although the lawyers in our study only monitored particular legal topics, we define the word 'developments' broadly as it is feasible that a lawyer might also want to maintain awareness of particular case developments, people or firms' activities. We found that monitoring may be *active*, facilitated by pull technologies such as the 'current awareness' sections of LexisNexis Professional and Westlaw (which allow lawyers to browse recent legal developments by topic) or *passive*, facilitated by push technologies such as e-mail alerts or mailing lists.
- The information behaviour of 'searching,' which is not discussed at all by Ellis and his Colleagues was also identified. Although it is included as part of Meho and Tibbo's (2003) summary model of behavioural characteristics, it is not discussed in detail in their paper. Searching involves formulating a query in order to locate information within a particular meta-resource (a resource that catalogues or indexes other resources), resource, source or document. Indeed the lawyers in our study searched for resources (resource searching), sources (source searching), documents (document searching) and content within documents (content searching). Searching behaviour also involves the lower-level behaviours associated with editing the query that is entered into the system. The lower-level search query editing behaviours identified in our study are 'search query refining,' (making minor changes to a search query in an attempt to improve the volume or quality of the results) 'search query reformulating,' (formulating a query again from scratch, often differently) 'search query refocusing' (adjusting the focus of the current query, but not reformulating it altogether) and 'search query spelling/syntax altering' (changing the spelling of query terms/the rules used to instruct the system to connect the query-terms or define the scope of the search). Searching behaviour also subsumes the lower-level behaviour of 'search result sorting,' (arranging the results of a search in a systematic order, such as in date order). The definitions for these behaviours are adapted from the Oxford English Dictionary.
- Ellis's behaviour of 'browsing' was also identified in our study of lawyers. We adopt Ellis's definition of *"semi-directed searching in an area of potential interest"* (Ellis, 1989, p.178). However, whilst studies by Ellis and his colleagues and Meho and Tibbo only describe browsing sources for documents (document browsing), we also identified browsing behaviour at the resource, source, and content levels.
- We also found evidence of Ellis's behaviour of 'chaining,' which was found only to operate at the document/content level (even though it may be theoretically feasible to follow referential connections between resources or sources as well as documents). This behaviour retains Ellis's original definition of *"following chains of citations or other forms of referential connections between material"* (Ellis 1989,

p.178). Both types of chaining that were identified by Ellis (1989) were also identified in our study: *forwards chaining* (which involves following chains of citations or other forms of referential connections between documents which have s*ubsequently cited the current document*) and *backwards chaining* (which involves following references to documents that have been *cited in the current document*). We also identified two other subtypes of chaining, *across resource* vs. *within resource* chaining and *direct* vs. *indirect* chaining. Across resource chaining involves following referential connections between material that leads from one electronic resource to another. Within resource chaining involves following referential connections between material that exists within the current resource. In many electronic resources, chaining can often be facilitated *directly* by following hyperlinks to other referenced materials. Sometimes *indirect* chaining is necessary, where the user has to follow the references manually. This is often the case when chaining across resources.

- Finally, we found evidence of Ellis's behaviour of 'extracting.' As with 'monitoring,' we also slightly adapt Ellis's definition of extracting to "*systematically working though a particular meta-resource, resource, source or document to identify material of interest*" to account for the fact that it may be possible to perform extracting at the resource, source, document and content levels (in this study, however, we only observed extracting at the document and content levels). At the content level (i.e. when lawyers were extracting content from documents), this behaviour can be better regarded as a lower-level example of 'selecting and processing' as opposed to 'identifying and locating' (because unlike identifying resources, sources and documents, identifying *content* usually involves a greater degree of processing than identification). This is why extracting is presented twice in table 4. At this content level, we found that extracting may be *direct* or *indirect*. Direct content extracting involves systematically working through the actual content of sources or documents whilst indirect content extracting involves systematically working through meta-data instead.

The second broad category of accessing is a standalone information behaviour in its own right and was identified as a discrete behavioural information-seeking characteristic by Meho and Tibbo (2003). We define accessing as *"gaining access to resources, sources or documents/content,"* again adopting the Oxford English Dictionary definition. We found that accessing could be *direct* or *indirect*, *visible* or *invisible*. Indirect accessing involves gaining access to a resource, source or document/content by using a third-party site or resource as a gateway (for example logging in using the educational Athens devolved login). Direct accessing involves gaining access without using a third-party gateway (for example logging in directly to a particular resource). Visible accessing involves gaining access to a resource, source or document/content through a procedure that can be seen at the interface level (usually a username/password login screen). Invisible accessing involves using recognition technologies (such as IP recognition) to gain access automatically, without a noticeable access procedure.

Several characteristics were also subsumed under the third broad category of 'selecting and processing resources, sources, documents, content and searches':

- This included Ellis's behaviour of 'distinguishing,' which involves *"ranking sources or documents according to their relative importance based on own perceptions." -* Ellis and Haugan 1997, p. 399).
- We also identified Ellis's behaviour of 'filtering' (the *"use of certain criteria or mechanisms when searching or browsing for information to make the information as relevant and as precise as possible." –* Ellis and Haugan 1997, p. 399).
- A new lower-level behaviour of 'selecting' was also identified. We define selecting as *"carefully choosing resources, sources, or documents as being potentially useful for the information task at hand"* (definition adapted from Oxford English Dictionary). Selecting shares conceptual similarity to distinguishing, filtering and extracting behaviours but is subtly different to these behaviours. It is different to distinguishing because it does not involve ranking resources based on perceived importance. It is also different to filtering because although various criteria are used when selecting which resource to use, these criteria are not used as precise, explicit filters to help decide between candidate resources. Instead, these criteria are often used implicitly to guide lawyers' choice of which resource to use. The boundaries between these behaviours, however, is not clear-cut. For example, the behaviour of 'extracting' documents from sources is different from 'selecting' as it is a behaviour that focuses more on

the process of *locating* the document amongst other documents. However, extracting *content* from documents is highly similar to both selecting and extracting behaviours. We found that selecting may be *direct* or *indirect*. Direct selecting involves examining the actual content of resources, sources or documents when choosing them as being potentially useful for the task at hand whilst indirect selecting involves using meta-information about the content (such as a results snippet or summary) as a substitute for examining the content when choosing them.

- Also subsumed under 'selecting and processing' was the information behaviour of 'recording,' which is similar in scope to 'information managing' behaviour as identified by Meho and Tibbo (2003). According to Meho and Tibbo, information managing involves *"filing, archiving, and organizing information collected or used in facilitating research"* (p. 582). The 'recording' behaviour we identified can also involve filing, archiving and organising information. However, whilst 'information managing' is a rather broad term which can subsume several types of lower-level behaviour, we believe 'recording' is a more precise description of the behaviour we observed. Recording involves making a record of resources or sources used (resource and source recording respectively), of documents and content found (document/content recording) or of the query terms used or results returned in a search (search query/results recording). Recording behaviour can be *manual* (i.e. by hand) or *automatic* (with the help of technology – such as a 'search trail' which automatically keeps a record of search queries entered and results received).

- Also subsumed was the information behaviour of 'document/content updating' which has not, to the best of our knowledge, been identified in previous information-seeking studies. Updating involves ensuring a *current* understanding of amendments or changes to legal documents and content and an understanding of whether a particular case or piece of legislation is *good law*. Updating behaviour can be performed *directly* by searching or browsing for documents and content and manually checking that the document or content within is up-to-date or good law. Updating behaviour can also be performed *indirectly* by using an electronic citator service to check whether a particular document or the content within it is up-to-date or good law.

- We also identified the information behaviour of 'document/content history tracking' which has also not, to the best of our knowledge, been identified in previous information-seeking studies. History tracking involves ensuring a historical understanding of amendments or changes to legal documents and content (i.e. an understanding of the treatment a particular case or piece of legislation has received over time). Although, as with updating, history tracking behaviour can theoretically be performed either directly or indirectly, we only observed direct history tracking behaviour amongst our academic and practicing lawyers, probably due to the fact that the electronic resources that they used did not provide explicit tools to support or automate this behaviour.

- Two other information behaviours subsumed under the broader category of 'selecting and processing' were 'content analysing' and 'content synthesising,' which were both identified but not greatly elaborated on by Meho and Tibbo (2003). Adopting definitions based on the Oxford English Dictionary, analysing involves *"examining in detail the elements or structure of the content found during information-seeking"* and synthesising involves *"combining the elements of the content found during information-seeking into a coherent whole."*

- Finally, we identified the information behaviours of 'document/content collating' and 'document/content editing' which like some of the other behaviours subsumed under the broader behaviour of 'selecting and processing' have not, to the best of our knowledge, been identified in previous information-seeking studies (although both bear some surface similarity to Ellis's 'ending' behaviour and might share some similarity to Meho and Tibbo's 'synthesising' behaviour). Collating involves *"drawing together documents and/or content for later use."* Editing involves *"preparing and arranging documents and/or content for later use by making revisions or adaptations."*

The fourth and final broad category of 'distributing documents, content and search queries/results' is also a standalone information behaviour and involves handing or sharing out entire documents, particular content or search queries/results to others.

Our study has identified behaviours that were identified in previous studies but have either not been discussed in detail, or have only been discussed in relation to paper-based as opposed to electronic information-seeking and use. For example, 'searching' was not discussed explicitly by Ellis and his colleagues as information *searching* behaviour is likely to have been beyond the scope of his studies. Searching was also only briefly mentioned by Meho and Tibbo (2003). This behaviour, however, is the central focus for existing models of the information search process (see Sutcliffe and Ennis, 1998 for an example). 'Accessing' behaviour was identified by Meho and Tibbo (2003), who discuss this behaviour as part of their study of Social Scientists. However, whilst they only discuss accessing in relation to gaining access to *physical* resources such as books, we discuss it in relation to gaining access to *electronic* resources and the documents within them. This has enabled us to discuss how accessing can be performed in new ways using electronic (rather than paper-based) resources.

Our study has also identified several new behaviours, including 'selecting' (which was found to supplement 'differentiating' as a way of choosing whether or not to view or use a particular resource, source or document). Apart from 'selecting,' which can be regarded as a general information-seeking behaviour, the newly-identified behaviours from our study fall into two categories:

1.   Law-specific behaviours that are particularly pertinent to lawyers (and are unlikely to be identifiable in other domains). The new law-specific behaviours identified in our study are 'updating' and 'history tracking.'
2.   Information *use* as opposed to information-seeking behaviours. The new information use behaviours identified in our study are 'recording,' 'collating,' 'editing' and 'distributing.'

We believe that the law-specific behaviours were not identified in previous studies as, to the best of our knowledge, there have been no previous studies of lawyers that have identified specific information behaviours that they perform. We believe the information use behaviours were not identified in previous studies as we adopted a broader scope than many of these existing studies – examining information *search* and information *use* behaviours in addition to information-*seeking* behaviours.

There are also a couple of behaviours that were identified in previous studies of information-seeking in other disciplines, but were not identified in our study. These are 'verifying' and 'networking.' We believe that 'verifying' (*"checking the information and sources found for accuracy and errors"*- Ellis et al., 1993, p.364) was not identified in our study not because accuracy is not as important for lawyers as it is for physical scientists but, quite the opposite, that legal documents must be checked thoroughly for accuracy *before* they are made available in electronic or paper form. We believe that 'networking,' was not identified in our study as although we identified some limited evidence of a social dimension to lawyers' information-seeking and use, our study did not focus on lawyers' broader information practices and therefore this was not identified as a standalone and important behaviour in its own right in the same way as it was in Meho and Tibbo's (2003) study of social scientists.

We have already presented the identified information behaviours under the broad categories of 'identifying and locating,' 'accessing,' 'selecting and processing' and 'distributing.' Taking into account the fact that most of our newly-identified behaviours can either be considered to be law-specific or information use behaviours, it is also possible to split all of the behaviours identified in our study (not just the newly-identified ones) into three sets.

The first set, listed and defined in table 5, can be regarded as 'core' information-seeking behaviours. Most of these behaviours have been identified in a range of disciplines and are primarily focused on information-seeking (rather than information *use*). The second set (see table 6) can be regarded as 'law-specific' behaviours and we suggest that these are particular to information-seeking in the *legal* domain. The final set of behaviours (see table 7) are primarily focused on information *use* (rather than information seeking) and we hypothesise that these or similar behaviours can be observed across disciplines. In tables 5, 6 and 7, we define each of these behaviours. We also summarise the levels at which many electronic resources currently, and could theoretically, support the behaviours. Finally, in the tables, we summarise levels at which the behaviours were actually identified in our study. Some behaviours are discussed together in the tables as they are commonly supported by

the same parts of an interface (for example, selecting, distinguishing and filtering can all be supported through the provision of various items of document metadata).

Core information-seeking behaviours

Behaviour name(s) and definition(s)	Levels at which many electronic resource currently and might theoretically support the behaviour(s)	Levels actually observed in our study
Accessing involves the process of gaining access to an electronic resource or to sources or documents and content within a resource, for example by logging into it.	Many electronic resources put restrictions in place that permit or deny access to the entire *resource*. However, some electronic resources also include restrictions on individual *sources* or *documents* (i.e. when not all users are permitted unlimited access to all parts of the resource and the content within it).	In our study, evidence for accessing was only provided at the resource level (probably because it is far less common for electronic resources to restrict access at the source or document levels).
Surveying involves the initial search for information to obtain an overview of a subject field, or to locate key people operating in a field (e.g. important authors in a certain legal area).	Apart from surveying the research area for *documents and content*, some resources support users in seeking an overview of the *sources* within a particular electronic resource that relate to the research area of interest.	In our study, evidence for surveying was only provided at the combined document and content level (as none of the lawyers in our study had the need to gain an overview of a particular source).
Monitoring involves maintaining awareness of developments in a field.	Like surveying, apart from monitoring the research area for *documents and content*, some resources also support users in maintaining awareness of the available *sources* within a particular electronic resource.	Our study provided evidence of monitoring at both of these levels (although there was far less evidence of monitoring at the source level, probably because it is more important for lawyers to maintain awareness of legal areas as opposed to individual sources).
Searching involves formulating a query in order to locate information. Formulating a query often involves entering search query terms into a search field or a series of fields for submission, however can take other forms (such as stepping through a Query Wizard that guides the creation of the query).	Many electronic resources support searching at the *document* level. However, it is not only possible to search for documents, but also to search for *sources* and to search within documents for particular textual *content*. These levels are supported by some electronic resources.	Evidence for searching at all of these levels was observed in our study.
Browsing and **extracting**. Browsing involves semi-directed searching for sources, documents or content. Whilst browsing, it is common to also perform extracting (which involves systematically working though a particular resource to identify sources of interest, a particular source to identify documents of interest and/or a particular document to identify content of interest).	Like searching, many electronic resources support browsing and extracting. It is possible to browse electronic resources and extract *documents* and also to extract *sources*. It is also possible to browse within documents to extract textual *content*.	Evidence for browsing and extracting at all of these levels was observed in our study.

Chaining involves following chains of citations or other forms of referential connections between sources or documents.	Most electronic resources support chaining between *documents* (although not always to documents that are available in other, often competing, resources). It is also possible to support chaining between *sources* (which is particularly necessary when sources are related to one another - for example some journals are superseded or change name and many overlap in topic areas of coverage) and chaining between *resources* (which is rarely supported by current electronic resources – probably due to the desire to avoid chaining to a competitor's resource).	Only evidence for chaining at the combined document and content level was observed in our study, probably because chaining at the resource level is rarely supported by electronic resources and there is not often any need to chain between sources.
Selecting, **distinguishing** and **filtering**. These are different ways of choosing relevant information (i.e. resources, sources or documents). Selecting involves carefully choosing resources, sources or documents as being potentially useful for the information task at hand (based on own or shared perceptions). Distinguishing is similar to selecting, but involves *ranking* information sources or documents according to their relative importance (again based on own or shared perceptions). This means deciding that one or more source or documents from a group is likely to be more useful than the others. Filtering involves the use of certain criteria or mechanisms when searching or browsing for information to make the information as relevant and as precise as possible (for example restricting a search to return documents by a particular author).	It is possible to select between, distinguish between and filter *sources* (based on the documents within them) and, more commonly, *documents* (based on the document content or various aspects of the document meta-data). It is also possible to select or distinguish between electronic *resources* themselves.	Evidence for selecting, distinguishing and filtering at almost all of these levels was observed in our study (although lawyers in our study tended to select rather than distinguish between resources and therefore distinguishing was not identified at the resource level).

Table 5: Core information-seeking behaviours along with their definitions and the relevant levels that electronic resources might support them.

Law-specific information-seeking behaviours

Behaviour name(s) and definition(s)	Levels at which many electronic resource currently and might theoretically support the behaviour(s)	Levels actually observed in our study
Updating involves gaining a *current* understanding of the importance of a particular legal document (i.e. an understanding of whether a particular case or piece of legislation is currently good law). **History tracking** involves gaining a *historical* understanding of the importance of a particular legal document (i.e. an understanding of the treatment that a particular case or piece of legislation has received over time).	Updating and history tracking only apply to gaining a temporal understanding of the importance of *documents* (and the content within them). Many electronic legal resources support these behaviours, either indirectly or through the provision of dedicated citator tools.	Evidence for updating and history tracking at the combined document/content level was observed in our study.

Table 6: Law-specific information-seeking behaviours along with their definitions and the relevant levels that electronic resources might support them.

Information use behaviours

Behaviour name(s) and definition(s)	Levels at which many electronic resource currently and might theoretically support the behaviour(s)	Levels actually observed in our study
Analysing and **synthesising.** Analysing involves examining in detail the elements or structure of the content found during information-seeking. Synthesising involves combining the elements of the content found during information-seeking into a coherent whole.	Electronic resources do not usually provide much support for analysing or synthesising document *content* that might be or has been extracted from a document.	Evidence for analysing and synthesising at the content level was observed in our study.
Recording involves making a record of resources or sources used, of documents or content found or of the query terms used or results returned in a search.	Most electronic resources provide facilities for recording *documents or content* within them (such as options to download them or print all or sections of the documents). Some resources also provide facilities for recording *sources* (for example by keeping a customisable list of frequently used sources to facilitate quick searching within them) and keeping a record of the *search query terms used or results returned* in a search. It is also possible to keep a record of the electronic resources used when looking for information.	Evidence for recording at all of these levels was observed in our study.
Collating involves the physical act of drawing together documents and/or content for later use.	Although some electronic legal resources support collating *documents* (e.g. by allowing users to download or print them in batches), few support collating of parts of documents (i.e. *content*).	Evidence of collating at both of these levels was observed in our study.
Editing involves preparing and arranging documents and/or content for later use by making revisions or adaptations.	Electronic legal resources do not usually provide much support for editing the *content* within documents.	Evidence for editing at the content level was observed in our study.
Distributing involves handing or sharing out entire documents, particular content or search queries/results to others.	Some electronic legal resources support the distribution of entire *documents*. Few support distributing parts of documents (i.e. *content*) or *search queries/results* to others. It is also theoretically possible to distribute *resources* or *sources* to others.	Evidence for distributing was restricted to the document and search query/result level in our study, probably because there is rarely a need to distribute resource or sources, even though it is theoretically possible.

Table 7: Information use behaviours along with their definitions and the relevant levels that electronic resources might support them.

Although the concept of 'levels' applies to all of the information behaviours identified, not all of the behaviours involve a simple relationship between behaviour and level. A simple relationship does exist for 'distributing' behaviour - we can hand or share out documents, particular content or search query/results to others. However, consider behaviours such as 'surveying,' 'monitoring,' 'updating' and 'history tracking.' These are behaviours

that focus primarily on the *understanding gained*. For example, whilst surveying can be achieved by looking at *documents*, we are primarily trying to gain an overview of a particular topical *area*. Similarly, when updating or history tracking, we are trying to gain an *understanding of the importance* of a particular legal document - either at this particular moment in time or historically. However, the distinction between those behaviours that are aimed at gaining an *understanding* of the information landscape and those that are not is not greatly important when seeking to inform the design and evaluation of electronic resources. This is because designers rarely have the opportunity to design to promote an *understanding* of information. Instead, they must aim to promote understanding as a by-product of supporting particular information behaviours. For example, a designer of an electronic legal resource might hope that providing support for lawyers to find out whether a particular piece of legislation is still in force will provide them with an understanding of the importance of the legislation, or that providing lawyers with an e-mail alerting service when new documents are published related to a particular legal area will help them gain an understanding of developments within that area.

The remaining sections, discussing each of these behaviours and related levels and subtypes, are presented in the order in which they have just been described and are therefore structured under the broad headings of 'identifying and locating resources, sources, documents and content' (section 5.3 of this chapter), 'accessing resources, sources, documents and content' (section 5.4), 'selecting and processing resources, sources, searches, documents and content' (section 5.5) and 'distributing documents, content and search queries/results' (section 5.6). Each behaviour and subtype is discussed in turn and reference is made to excerpts from participants' transcripts in order to provide evidence for each behaviour and subtype. Excerpts can be identified by the preceding 'A' to denote 'Academic lawyer' or 'P' to denote practicing lawyer. The academic role/level of academic lawyers and job role or practicing lawyers is then presented in brackets (e.g. A21 (Professor) denotes an academic Professor of Laws). Lawyers' interface-level actions are presented in square brackets within each excerpt. Square brackets are also used where the lawyer is referring to something earlier in the excerpt or transcript (which would otherwise be unclear) and where we have omitted details (for example the name of the firm's in-house Knowledge Management database) for confidentiality reasons. Three dots within square brackets […] denotes places within the excerpt where text has been omitted for reasons of clairty. An illustrative transcript of an Associate Tax lawyer looking for information on a rather complex Capital Gains Tax issue is presented in its entirety in appendix 1.

5.3   Identifying and locating resources, sources, documents and content

Identifying and locating resources, sources, documents and content involve several lower-level behaviours, all of which have been previously identified (albeit not at different levels or separated into subtypes) by either Ellis and his colleagues, by Meho and Tibbo or by both. These behaviours include *surveying*, *monitoring*, *searching*, *browsing*, *distinguishing, filtering, selecting*, *extracting* and *chaining* and are discussed and in detail below.

5.3.1   Surveying (D&C)

Surveying behaviour, *"characteristic of the initial search for information to obtain an overview of the literature within a new subject field, or to locate key people operating in this field"* (Ellis and Haugan 1997, p. 395), was found to be a common information behaviour amongst all four groups of lawyers (i.e. taught students, research staff and students, Dispute Resolution lawyers and Tax lawyers) who surveyed the research area for documents and content. Electronic surveying behaviour was displayed or mentioned in four ways:

- By using secondary electronic commentary sources to gain an overview of the area.
- By using Internet search engines, almost always Google or Google Scholar, to gain an overview of the area.
- By using personal collections of resources (for example stored bookmarks of Internet sites).
- By using shared collections of resources (for example a departmental Intranet site with links to resources).

Surveying behaviour in the legal domain involved much more 'gaining an overview' of particular legal areas as opposed to 'locating key people' that have written legal documents in that area, probably due to the relative lack of importance of who wrote a particular document (provided it is published in a reputable source) and the greater importance of the content and legal principles arising from the document.

We identified one pair of subtypes of surveying behaviour, *lightly directed surveying* and *heavily directed surveying*. Strictly speaking, these are *not truly mutually exclusive* subtypes as 'directedness' can be regarded as a continuum. However, for practical purposes, lawyers often displayed surveying behaviour at the extremes of this continuum and therefore they are discussed as mutually exclusive in this section. Lightly directed surveying was undertaken when the lawyer had few or imprecise details about the area on which an overview is required. Heavily directed surveying was undertaken when lawyers already had many or specific details about the research area (either due to prior knowledge of the area or due to a 'research lead' from a colleague or lecturer).

Lightly directed surveying was common across all groups of lawyers, as it was often the case that the information-seeking problem was vague and not as well defined at the start of the information-seeking episode as it would become by the end:

> *A4 (PhD Student): The first thing I do is if I have absolutely no idea about the area then I will turn to secondary sources.*

Heavily directed surveying was common amongst taught students, but not at all common amongst research students and staff and practicing lawyers. These groups of lawyers tended to scope out their own direction for surveying a research area rather than rely on others to provide them with direction. Taught students often obtained 'research leads' from colleagues or, more frequently, academic staff:

> *A30 (2ⁿᵈ year LLB): We were referred to a number of publications and articles, so I guess the first thing was to basically get as many of the articles that we could, mainly from online resources. Also just to build a bibliography of relevant sources, so that at the stage where we were actually building the dissertation we would know where the relevant material was to be found.*

Surveying was almost always facilitated by searching or browsing for documents and was only identified at the combined *document and content level*. This may be because most information tasks are focused on finding information (i.e. documents and content) as opposed to finding useful sources or resources and this makes the widespread observation of 'source surveying – gaining an overview of potential sources to use' or 'resource surveying – gaining an overview of potential resources' unlikely.

Using secondary electronic commentary sources

All groups of lawyers used secondary electronic commentary sources to gain an overview of a research area, and often as a springboard to facilitate chaining, as one PhD student explained:

> *A4 (PhD Research student): Always the first step in legal research unless you are very familiar with the subject is to go to the secondary source first and try to see roughly 'what's going on' and from those secondary sources you can find a citation or references to other sources.*

The lawyers used an assortment of electronic legal resources to facilitate document and content surveying behaviour. One of the most frequently used electronic secondary sources used by DR lawyers (contained within the LexisNexis Butterworths electronic legal resource) was Halsbury's Laws, as explained by a Dispute Resolution Trainee:

> *P6 (DR Trainee): Usually the sources that I'd look at here, there's one called 'Halsbury's Laws of England' which is a commentary source, so I'd click on that [selects appropriate source in list of*

databases]. It's got basic commentary for basic areas of law. So it might not have the answer to the specific question you want, but it might give you a good overview.

Particularly for Tax lawyers, other commentary sources such as Simon's Direct Tax Service (again contained within LexisNexis Butterworths) were also used to gain an overview of legal areas. This Tax Trainee, for example, conducted a search of all the commentary sources in LexisNexis Butterworths in order to gain an overview of the legal topic of 'Capital Allowances:'

P13 (Tax Trainee): Normally I would go to LexisNexis because you get more of a textbook overview to begin with than you would get on the Internet, which is likely to be some more specific advice that was given but won't reflect the specific fact pattern that I'm looking at. So as a Trainee, I didn't really know much about Capital Allowances, full stop. So I would load up LexisNexis and I would go through to the search menu [clicks on 'search' tab]. I would do a search through the commentary sources first off. [Clicks on the 'commentary' tab, which restricts the search to commentary sources].

Both DR and Tax practicing lawyers also mentioned that certain electronic legal resources, other than LexisNexis Butterworths, were particularly suitable for gaining an overview of various legal areas. These were resources that contained secondary materials such as Practical Law Company (PLC) and Current Legal Information (CLI). Both of these resources were mentioned, although not frequently, by both Dispute Resolution and Tax lawyers. In the excerpt below, a DR Trainee illustrates the use of PLC to find overview information on the topic of 'Competition Disqualification Orders:'

P14 (DR Trainee): The other thing I'd use is PLC, Practical Law Company, because that's got loads of articles. So if you just go to the PLC tab on the Intranet, you can just do a general broad search. So I could just put in 'competition disqualification orders' and that would throw up articles that had been written on it or general overviews, a definition of it, so that's quite useful if you don't know anything about the area. And that might lead me to narrow my search. For example, it might tell me that there's Office of Fair Trading guidance on the issue and I'd then go to the OFT website and look at them.

Using Internet search engines

The use of Internet search engines (almost always Google) to gain an overview of a research area was widespread amongst both academic and practicing lawyers and, like electronic secondary sources, were often used to find starter documents that would provide hyperlinks to chain from:

A22 (LLM Student): I guess Google is best for providing a start off really, if you don't really know what you're looking for and need to get as much material as possible. Start off with Google and when you have an idea of where you're heading, come back to Westlaw with a more specific search.

Academic lawyers also, along with Google, used the Google Scholar search engine to find scholarly legal articles. In addition, practicing lawyers occasionally used their own in-house knowledge management software (which performs a similar function to an Internet search engine) to facilitate surveying. However, this resource was used far less than either Google or Halsbury's Laws.

The first two ways of facilitating surveying (using secondary electronic resources and Internet search engines) differ somewhat from the ways in which other disciplines, such as social scientists (Ellis, 1989) carried out this behaviour (by seeking out people that knew about the area, reading reviews of materials, consulting bibliographies, abstracts, indexes and library catalogues). This may be partly due to differences between disciplines, however it is just as likely to be due to the availability of new technology (such as Internet search engines) since Ellis's study was conducted.

Using personal collections of resources

Using personal collections of resources was a widespread behaviour reported by Meho and Tibbo's (2003) academic social scientists. However, this behaviour was rare amongst the lawyers in our study (and was only displayed by two of the eight members of research and teaching staff). This indicates that surveying through the use of personal collections might be particularly important for academics, but not necessarily for other groups. One member of academic teaching and research staff did, however, use her personal collection of Internet bookmarks as her primary method of surveying the documents and content within a research area:

> *A7 (Lecturer): My starting point is almost always [pauses] unless I remember that I've already got a hardcopy of something, because I prefer reading hardcopies [pauses] then my starting point would almost always be the Internet and it would almost always be one of the bookmarked pages. So that's where I would start and use that as a jumping off point.*

Using shared collections of resources

The final way of surveying, using shared collections, was not mentioned in previous studies and was rather uncommon in our study of lawyers, although one practicing lawyer mentioned using these rather than personal bookmarks as starting points:

> *P6 (DR Trainee): I work in Dispute Resolution, and we also have a specific Dispute Resolution website that has been put together by our Knowledge Management team, I think, which is usually a very good starting point. If you've got a question and you think 'I've got absolutely no idea,' then it'll have the basics on the core things, so for here it will have something on appeals or part 36. This probably doesn't mean anything to you! But very basic stuff, if you need to get a general overview before you get into more complicated areas, that's probably where you'd go first to get an idea.*

Summary of surveying behaviour

Surveying behaviour was found to be a common information behaviour amongst all four groups of lawyers in our study and was displayed in four main ways: by using secondary electronic commentary sources to gain an overview of the area, by using Internet search engines, almost always Google or Google Scholar, to gain an overview of the area, by using personal collections of resources (for example stored bookmarks of Internet sites) and by using shared collections of resources (for example departmental Intranet sites with links to resources). We found that surveying may be *lightly directed* or *heavily directed* and observed the behaviour to operate at the combined document and content level.

5.3.2   Monitoring (S, D&C)

Monitoring behaviour involves *"maintaining awareness of developments in a field"* (definition adapted from Ellis and Haugan, 1997, p. 396). In effect, this involves monitoring the research area for sources, documents and content. Whilst this behaviour was rarely mentioned or displayed by taught students, there was slightly more mention and display by research students, research and teaching staff, and practicing lawyers. This difference might be due to the fact that taught students primarily conducted prescribed electronic research tasks, which resulted in little need to perform monitoring behaviour.

We identified one main set of subtypes of monitoring behaviour. We found that monitoring may be *active*, facilitated by pull technologies such as the 'current awareness' sections of the LexisNexis Butterworths and Westlaw electronic legal resources (which allow lawyers to browse recent legal developments by topic) or

passive, facilitated by push technologies such as e-mail alerts or mailing lists. We found that active monitoring was achieved in three main ways:

- By manually conducting regular searches on a particular legal topic in electronic legal resources.
- By regularly browsing particular sources in electronic legal resources or on law-related Internet sites.
- By regularly following previously bookmarked Internet pages.

Passive monitoring was displayed in one main way, by subscribing to e-mail alert lists.

None of these ways of achieving monitoring were identified by Ellis and his colleagues or by Meho and Tibbo. However, these forms of monitoring are conceptually similar to some of the ways that monitoring was found to be achieved in previous studies. For example, just as the social scientists in Ellis's (1989) study regularly consulted sets of journals which were deemed to publish material of interest, some of our lawyers periodically followed previously bookmarked pages to Internet sources deemed to publish material of interest. Just as the social scientists read secondary sources such as book publishers' lists and reviews and the English Literature academics interviewed by Smith (1988) consulted secondary sources such as annual journal bibliographies, book reviews and publisher catalogues, our lawyers browsed commentary sources in electronic legal resources and on law-related Internet sites. In addition, just as some of Smith's (1988) English Literature academics also reviewed primary sources, some of our lawyers manually searched primary materials within electronic legal resources (on a particular topic of interest).

Although active and passive monitoring were identified as a main set of subtypes, we also identified another two sets of subtypes which were less important for describing our data and therefore are not elaborated on greatly in this section. These were *manual* or *automated* monitoring (a subtype alerted by the fact that automated monitoring tools such as e-mail alerts were rarely used by any of the groups of lawyers) and *formal* vs. *informal* monitoring (a subtype alerted by the fact that much non-electronic monitoring behaviour appeared to occur through inter-personal contact between colleagues). The latter set serves to highlight the importance of personal contacts during information-seeking. However, as our study was focused primarily on the use of electronic resources (as opposed to 'people' resources), we do not discuss personal contact any further (and nor does it feature as a separate information behaviour in the same way as Meho and Tibbo's [2003] 'networking' behaviour).

The lawyers in our study displayed limited monitoring behaviour at the *source level* described by Ellis and his colleagues (for example by manually searching or automatically scheduling alerts when new information from a particular source is published) and some monitoring behaviour at the *document and content level* (by manually or automatically searching or browsing for recent or updated documents on a particular topic).

Conducting regular searches on a particular topic using electronic resources

Conducting regular topic searches using electronic legal resources is an example of active document and content monitoring (it is active in the sense that it must be performed manually, and operates at the document and content level because the searches are likely to span multiple sources within a particular resource). Whilst few taught students mentioned or displayed active monitoring behaviour, the exception was a Master of Laws (LLM) student who spoke about the need to regularly search for documents in electronic legal resources that might add to an academic debate of interest:

> *A17 (LLM student): One of the good things about online resources is that they are very up-to-date. It's good to check online to see if there's any recent additions to this academic debate or to see if anyone's written an article recently that isn't available in the library. So it's really about updating yourself, filling in the gaps [pauses], trying to find anything that you've missed.*

In addition, very few practicing lawyers mentioned or demonstrated conducting regular manual searches on electronic legal resources in order to facilitate monitoring behaviour. However, one DR lawyer did conduct weekly topic searches on behalf of a client in order to stay abreast of recent developments:

> **R: You mentioned that part of your work involves weekly checks of laws, I'm guessing to see whether they'd apply to your client's situation?** *P3 (DR Paralegal): Yes, that is right, because obviously it is a specific industry and they will want to know if there have been any developments. It would be quite a basic search, for example opening the 'Bills Before Parliament' and you would do just a simple word search to see if there have been any developments or not.*

Regularly browsing particular sources

An example of active source (as opposed to document and content) monitoring was regularly browsing particular sources, for example to determine whether any pieces of legislation had recently come into force or been amended. The academic and practicing lawyers below both mentioned regularly browsing the websites of different legislative bodies:

> *A11 (2nd year LLB student): The other thing I do research on is the Law Commission website, because as a law student you're supposed to know what's happening and the Law Commission is the body responsible for new laws and statutes coming out. For example, in Criminal Law the law of murder is in reform and it's going to come out in a few months and we're supposed to read the Law Commission proposals and they're only available on their website.*

> *P3 (DR Paralegal): If it's weekly checks or reports, then I wouldn't be using those databases, but go straight to Governmental websites to check whether any recent Bills or legislation has been introduced. So this would be either the UK Parliament website, Welsh Assembly, Scottish Parliament, Northern Irish. And recently I discovered one that's pretty good, the Cabinet Office Public Service website which gives you the details of Statutory Instruments, Acts and Bills in Progress.*

In addition, one Bar Vocational Course (BVC) student mentioned browsing the 'current awareness' section of the electronic legal resource Lawtel for information on a particular legal area:

> *A33 (BVC student): It has this section called 'current awareness' [pauses] things that have happened recently in one place, which is quite useful if you want to have an overview. You can scan down the page and see if there's anything even vaguely related to what you're doing.* **R: On a particular topic?** *A33: Yeah. I think you can select current awareness for whatever area [pauses] immigration [pauses] human rights.*

Regularly following bookmarked Internet pages

Another example of active source monitoring was demonstrated by a law lecturer, who periodically re-visited her bookmarks of Internet sources that published documents related to her research area:

> *A7 (Lecturer): I would have links to all the various documents in the areas that I researched; in particular that was European Union Employment Policy. So here [shows folder list] are my various links and it's an annual process of producing policy documents and I would know that around about March every year and in the summer there will be something new.*

Subscribing to e-mail alerts

Passive source and document/content monitoring was only carried out in one way by the lawyers in our study - by subscribing to e-mail alerts. These alerts focus on informing the user when new or updated documents are published. One level that e-mail alerts can operate at is the source level, where the alerts are focused on informing the user of new or updated documents from a *particular subscribed source*. For example, the same law lecturer performed passive source monitoring by subscribing to e-mail mailing lists written by various European Union Law organisations:

> *A7 (Lecturer): I subscribe to various mailing lists, legal and other mailing lists from government departments such as the Treasury or publishers telling me what they have recently published. You can also subscribe to things from think-tanks or one of the main ones I subscribe to is one of the units responsible for European Employment Law and Policy. They send weekly mails about the work they're doing and it just links you to this week's news on Employment Law and Policy in the EU.*

Similarly this Tax Trainee subscribed to e-mails from the PLC (Practical Law Company) electronic resource that sends him weekly digest alerts detailing when new Corporate or Finance-related articles have been published:

> *P24 (Tax Trainee): I've got PLC Corporate and PLC Finance updates. They will e-mail me and tell me what articles they've produced or what new topics are under consideration. So I've kept these e-mails dating back to when I joined and then you can see the development. The idea being, again, to make sure that you've always got the latest law to hand, and the best way of doing that is making sure you have the latest article or commentary to hand. And this is the best way to do that.*

Another level that e-mail alerts can operate at is the document and content level, where alerts are focused on informing the user of new or updated documents *related to a certain research area* (but published in multiple sources). As one DR Trainee explained, it is also possible to request e-mail updates on individual documents, such as a specific legal bill - another example of passive document and content monitoring:

> *P6 (DR Trainee): We get updates sent through to our inbox as well of any important developments, so if a really key case comes through that's going to affect an area of the law which Dispute Resolution lawyers use all the time, then you would get an update into your inbox flagging up that there had been some changes [pauses] there is a system that you can set up, which I've never had to use. If you're tracking a certain area - for example you know it's relevant for a particular case you're working on [pauses] you can get updates on a specific area sent to your inbox, you can set that up. Or you can get updates on a specific case, for example, or on a specific Bill if you're tracking a Bill that is going through Parliament. And, in fact, I know that my supervisor has got that set up for something that he's doing, but I've never had to use that before.*

Summary of monitoring behaviour

Monitoring behaviour was not observed particularly frequently amongst the lawyers in our study (and was particularly rarely observed amongst taught students). This may, however, be due to the nature of the naturalistic task set to participants – which may have discouraged them from demonstrating monitoring behaviour due to the need to find information 'currently needed as part of their work.' We found that monitoring may be *active*, facilitated by pull technologies or *passive*, facilitated by push technologies and found that whilst passive monitoring was by subscribing to e-mail alert lists, active monitoring was achieved in three different ways by manually conducting regular searches on a particular legal topic in electronic legal resources, by regularly browsing particular sources in electronic legal resources or on law-related Internet sites and by regularly following previously bookmarked Internet pages. We observed monitoring behaviour to operate at the source level and the combined document and content level.

5.3.3 Searching (R, S, D, C)

Searching involves formulating a query in order to locate information within a particular meta-resource (a resource that catalogues or indexes other resources), resource, source or document. The lawyers in our study searched for resources, sources, documents and content and we now discuss each of these levels of searching in detail. Although 'searching' behaviour was found to be highly common across all groups of lawyers, it is not discussed at all by Ellis and his Colleagues (other than in the context of 'filtering' behaviour, which involves *"use of certain criteria or mechanisms when searching or browsing for information to make the information as relevant and as precise as possible."* - Ellis and Haugan 1997, p. 399). In addition, although it is included as part of Meho and Tibbo's (2003) summary model of behavioural characteristics, it is not discussed in detail. Similar searching behaviours and lower-level search tactics were, however, identified by Sutcliffe and Ennis (1998) and we discuss these below.

Resource searching (searching meta-resources for resources)

Although we observed searching at the resource level (i.e. lawyers searching meta-resources for electronic resources), searching at this level was somewhat rare. Resource searching was observed to be most common amongst taught students, who often used meta-resources such as the university library's database search tool to locate and access particular electronic resources. For example one 2nd year Bachelor of Laws (LLB) student would always access the electronic legal resource Westlaw through the university homepage:

> *A24 (2nd year LLB student): I normally go through the [university] Website. I just type in 'west law' here [types search terms in the search box integrated on the university homepage, which invokes a Google search restricted to the university site].*

In addition, commonplace behaviour across all groups was to use Google to locate other electronic resources, particularly when the lawyer was unsure of the name or Internet address of the required resource:

> *A12 (PhD student): I was looking for the Ministry of Justice [pauses] they have a special unit that vets all bills before they are sent to parliament. So I tried to find that unit and couldn't find it, so I Googled it because I didn't know the exact address, found it and looked through their website […].*

Source searching (searching resources for sources)

Along with the previous behaviour of resource searching, searching at the source level (i.e. searching resources for sources) was also quite rare. As with resource searching, this level of searching was mentioned and displayed more by taught students than the other groups of lawyers (although some source searching behaviour was observed across groups). One LLM student demonstrated searching at the source level by searching the Westlaw directory of journals covered, so as to ensure that she was using the correct abbreviation for the 'Common Market Law Review' journal in one of her citation field-restricted searches:

> *[Conducts search and does not retrieve any results]. A9 (LLM student): That might mean that I might not be able to get this on this particular website. [Navigates to Westlaw directory of journals covered and searches for 'common,' selecting the 'starts with' radio button]. Common Market Law Review [pauses] there it is. I got the wrong citation I think. It's C.M.L. Rev. not C.M.L.R.*

Other taught students mentioned and demonstrated searching for journal titles on their university library website. Similarly one DR Practice Development Lawyer spoke of performing a similar task by searching the firm's library catalogue:

P4 (DR PDL): If you need a really old article, you'd find it in CLI and find that it belongs to this particular journal and the date and the issue number and then you'd go to the library catalogue online and look 'do we have that journal? Does it go back that far?'

Source searching amongst practicing lawyers usually involved searching within particular news or commentary sources. For example the Tax Associate below searches within the 'Company Tax Manual' source on LexisNexis Butterworths. The Tax Trainee in the second excerpt below performs a similar search within Lexis, this time restricted to the 'Taxation Magazine' source:

P21 (Tax Associate): All the references are pointing to one particular manual out of 25 or 30, called Company Tax Manual. [Selects 'Company Tax Manual' source and begins to search within it]. I'll try to search it for "loan relationship" AND "participator" AND "paragraph 2." [Conducts search].

P7 (Tax Trainee): I'm going into a specific journal magazine just to see if they have anything more. I don't think they will because I think all this stuff should've come up when we did the initial searches, but I thought that if I just go into each individual one then it might bring more back. [Searches for 'European law AND avoidance' restricted to only the 'Taxation Magazine' source]. This result came up before in the general search. So I probably won't look at that. So based on the headings, I don't really see anything useful.

Document searching (searching meta-resources, resources or sources for documents)

By far the most common level of searching was document searching (i.e. searching for documents). However, several lower-level searching behaviours were also identified in our study. These were '*search query formulating,*' '*search query editing,*' and '*search result sorting*' (arranging the results of a search in a systematic order, such as in date order). The behaviour of 'search query editing' comprises further lower-level behaviours of '*search query refining,*' (making minor changes to a search query in an attempt to improve the volume or quality of the results) '*search query reformulating,*' (formulating a query again from scratch, often differently) '*search query refocusing*' (adjusting the focus of the current query, but not reformulating it altogether) and '*search query spelling/syntax altering*' (changing the spelling of query terms/the rules used to instruct the system to connect the query-terms or define the scope of the search). All of these behaviours are presented in table 8 below as an expansion of 'document searching' behaviour. It is important to note, however, that these lower-level behaviours are not unique to searching at the document level. However, since searches at the resource, source and content levels were relatively straightforward (i.e. required little added search syntax or query editing), they were not identified as lower-level behaviours in their own right.

Document searching	Search query formulating	
	Search query editing	Search query refining
		Search query reformulating
		Search query refocusing
		Search query spelling/syntax altering
	Search result sorting	

Table 8: Expansion of the 'document searching' behaviour showing the lower-level searching behavioural characteristics as identified in our study.

Similar searching behaviours to the ones found in our study, along with related lower-level search tactics were identified by Sutcliffe and Ennis (1998) (although the authors present these as cognitive processes rather than the physical correlates to these processes that result in a related information behaviour). We propose that their cognitive activity of 'query formulation,' which involves *"identifying search terms and transforming them into the query language supported by the search system"* (p. 327), is highly conceptually similar to our information behaviour of 'search formulating.' Similarly we propose that 'evaluating results' (which involves *"scanning the*

results set or examining the contents in detail in order to decide whether to accept the retrieved results or continue searching" p. 327), is conceptually similar to Ellis's 'differentiating/distinguishing' and our 'content extracting' behaviours.

Search query formulating

Formulating a search query involves choosing appropriate query terms and search syntax. For example the Tax Associate below, who was searching for information about whether a company can issue loan notes to directors of a company before making a takeover offer, used Boolean syntax to connect the search terms used and decided to search for 'management' rather than 'company director' because she was aware that she was searching the firm's own Knowledge Management database and that, in this database, it would be likely that directors would be referred to as 'management':

> *P12 (Tax Associate): We're looking at section 135 of the legislation and want it particularly in the context of loan notes and probably, as I said we're considering this specifically for the directors at the moment, I know from past experience from the way that usually things are written up in our Tax news that finds its way onto the Info bank, people probably refer to that as 'management.' [Searches for "I35" AND "loan note" AND management].* **R: Why did you search for 'loan note' rather than 'loan notes?'** *P12: Because then I get 'note' singular and plural.*

Search query formulating also involves defining the scope of the search. For example, the DR Practice Development Assistant below restricted his search to a particular type of legal document (case judgements) and to a particular court. In addition formulating a search also involves deciding which search field within a particular electronic resource to use:

> *P19 (DR PDA): I think I went into Case Track next and went to search for a 'judgement.' Obviously you have to select the type of court you want to search, as you can't go across-the-board [selects relevant court type using radio button]. I think I actually put my search in the free-text because I couldn't think of where else to put it down here. So I think I'd just put 'article 6.'*

The behaviour of defining the scope of a search (which was observed in our study as a potential means of search query formulating, search query refining and search query refocusing) is similar in nature to Ellis and Haugan's (1997) behavioural characteristic of 'filtering.' Ellis and Hogan define filtering behaviour as the *"use of certain criteria or mechanisms when searching or browsing for information to make the information as relevant and as precise as possible"* (p. 399) and give the example of restricting searches (for example through the use of keywords or date restrictions), which is akin to the scope definition and editing behaviour identified in our study.

It is not always possible, however, to define the scope of the search as intended. This is illustrated in the transcript excerpt below, where a Tax Associate attempts to search a number of sections of Simon's Direct Tax manual at the same time in LexisNexis Butterworths. The Associate ticked the checkboxes next to the sections she would like to search within, but struggled to find a way of searching only within the selected documents as opposed to within the entire Simon's Direct Tax source:

> *P12 (Tax Associate): At the moment it's only letting me search within the Inland Revenue manuals, which is way too big. I don't want to search all the manuals for my search term. That would be absolutely useless. I just want to search one very small bit of the manuals. I'll have to say that I actually don't know how to do that. So at this point I'd probably give up on that.*

The search queries that were formulated varied in sophistication. Whilst all lawyers demonstrated the use of basic Boolean syntax, only a small number demonstrated the use of advanced syntax. An example is given by the Tax Trainee below, who used the within 'w/' operator in LexisNexis Butterworths to search for the search

term 'remittance' within ten words of the search phrase 'capital gains tax,' which, as he explained 'means that hopefully they'll be in the same sentence':

> *P23 (Tax Trainee): 'W/10' means 'within 10 words of,' so that means 'remittance within 10 words of capital gains tax' in the article, which means that hopefully they'll be in the same sentence, which means that hopefully they'll be talking about the same thing rather than being two incidental references.*

Search query editing

Search query editing comprises further lower-level behaviours of '*search query refining*,' '*search query reformulating*,' '*search query refocusing*' and '*search query spelling/syntax altering*.' Search query refining involves making minor changes to a search query in an attempt to improve the volume or quality of the results. This behaviour was quite common across all groups of lawyers and could be achieved in several ways:

- By broadening the query (e.g. by removing search terms or expanding the scope of the search).
- By narrowing the query (e.g. by adding search terms or restricting the scope of the search).
- By altering the spelling or syntax of query terms.

The first two ways that lawyers demonstrated refining a query was to broaden it by removing search terms or expanding the scope of the search, or to narrow it by adding search terms (as with the PhD student below) or restricting the scope of the search (as with the DR Practice Development Lawyer below):

> *A4 (PhD student): That's brought back too much [pauses] 315 results. So that's too much for any reference. So I try to add another key word, which is 'insurance'. So "good faith" and 'insurance.' [Edits search].*

> *P2 (DR PDL): An immediate scroll down of the first page, which has got 10 results, shows that the thing that I'm looking for hasn't even come up. So what I'm gonna do now is look under 'internal author.'*

It is also possible to refine a search for documents by altering the spelling or syntax of query terms:

> *P14 (DR Trainee): So I'm going to go back to the other results and I think I'm going to amend my search to read 'director disqualification order AND subsidiary.' I'll put in an AND in capitals so that it brings back all those words in the results, not just some of them. [Conducts search, which retrieves a number of results].*

Altering the spelling or syntax of query terms is discussed later in relation to the separate low-level behaviour of 'spelling/syntax altering.' This behaviour is considered to be separate as not all instances of altering the spelling or syntax of search queries result in search query refinement.

Search query reformulating, which involves formulating a query again from scratch, was displayed far less than search query refining or search query refocusing (in fact only by a couple of taught students). The LLM student below, for example, struggled to find a case in Westlaw that was referred to by her lecturer simply as 'the UHT case' and, after searching for 'uht' to no avail, decided to abandon the search and, instead, conducted a broad Google search on the area of EU law that the case was reported to be about:

> *A9 (LLM student): It might just be that there's no article by that name [with 'uht' in the title]. Let me try [pauses] [begins typing]. I'm just trying to find keywords that might be in the title of a relevant article [pauses] hmm [pauses]. I think I'd go back to Google and maybe do a general search and*

maybe it'll give me more clues as to where I can go from there. [Goes to Google and searches for 'risk policy eu'].

Search query refocusing involves adjusting the focus of the current query as opposed to reformulating it altogether. This search behaviour was not as rare as search query reformulating, but not as common as search query refinement. As with search query refining, lawyers adjusted the focus of their queries in several ways:

- By using synonymous search query terms (for example PhD student A4 changed his search terms to 'uberimma fides' - the Latin for the principle of 'utmost good faith' which he was originally searching for but found no relevant documents on).
- By removing search terms or expanding the scope of the search, as illustrated by the following 1st year LLB student, who explained that it is possible to search for a case by searching for just one of the two party names in the 'party names' segmented field in Westlaw:

  A5 (1st year LLB): [Types in R and Soerl into 'Party Names' boxes]. **R: So whenever it's a case involving the Crown you can put 'R' in the party names box?** *A5: Yeah, you can also just put the other party in. It didn't have it [bring back any results], so this time I'll just put 'Soerl' in and see if that works.*

- By adding search terms or restricting the scope of the search.
- By changing search terms, as illustrated by the following Legal Practice Course student:

  A28 (LPC student): This search that I'd done, my first search when looking for the Collective Enfranchisement stuff, was just for 'Collective Enfranchisement' and I think it came up with quite a few cases. And obviously, I clicked on some of them, read one or two and realised that 'hold on, this search isn't going anywhere because these cases are talking about different areas.' At that point, I realised that I had to narrow this down a bit. So what's a narrower word or term related to what I'm looking for? Resident landlord. Yep. And when I typed that in, it brought up exactly the sort of case that I was looking for.

The final lower-level search editing behaviour was search query/syntax altering, which was far more common than the other lower-level search editing behaviours. Search query/syntax altering involves changing the spelling of the query terms or changing the rules used to instruct the system to connect the query-terms or define the scope of the search:

> **R: Why do you think it didn't find any documents?** *P19 (LLM student): Oh, I typed wrongly. [Corrects spelling mistake related to the name 'Stephen' and re-searches, this time finding documents].*

> *A28 (LPC student): I'm gonna assume that I've made a mistake with my original name of the legislation. If I remember correctly in Lexis you have to put AND everywhere [pauses] between each word. [Adds AND connectors between all words in the legislation title and conducts search].*

Search result sorting

Search result sorting, which involves arranging the results of a search in a systematic order, was rather uncommon, although it was displayed by a couple of DR and Tax practicing lawyers. One example of this behaviour was displayed by a DR Practice Development Assistant, who asserted that he preferred to display the search results from their internal Knowledge Management database in order of date rather than in the default order of machine-determined relevance:

P11 (DR PDA): Normally, as a rule of thumb, I will swap it for the content date straight off so that I have the most recent documents at the top [sorts results by reverse date order]. Not necessarily the way to do it, but it's the way I do it because I work backwards down.

The same procedure of sorting results by date was also performed by the Tax Associate below, who was also searching the firm's Knowledge Management database. She explained that it is more useful to find recent as opposed to older documents:

P12 (Tax Associate): I usually have mine prioritised by 'effective date,' so you get the most recent results first. Just because, I figure, if there's something relevant then I'd rather know what we were taking about recently because that is most likely to point to the current issues between our interpretation and maybe points that the Revenue that have come up with or tricky points that people have faced on recent deals.

Content searching (searching within documents for particular content)

Aside from searching at the resource, source and document levels we also observed a small number of lawyers searching within the content of a particular document. Lawyers usually achieved this behaviour by conducting a simple search within their Internet Browser or Word Processing package (e.g. by going to 'Edit' and 'Find' on the menu):

[Searches for 'R v.' within Internet Explorer to find mention of other cases within the current case report]. A19 (LLM student): This only works for criminal cases because I know for sure that the first party must be R!

P21 (Tax Associate): Now I'll just do a search for the word 'participator' [searches for word within the word document of the Budget document]. It comes up in the document, but nothing is relevant.

Another way in which content searching was achieved was by using the electronic resource itself, as opposed to the Internet browser search command, to search within the current document. The Tax Associate below explained how a previous version of LexisNexis Butterworths provided a search box along with the full-text of documents to users to search for particular words or phrases within the currently open document text, but that it was still possible to search documents for particular content by returning to the results list, ticking the checkbox next to the documents of interest (in this case sections of the Simon's Direct Tax commentary source) and typing in search terms:

P20 (Tax Associate): You see this system I find horrendous to use. In the good old days, you would click on that [points to 'Control of Foreign Companies' hyperlink] like the whole of this section, 'Control of Foreign Companies' and you would have a search box up here so that you could search within the highlighted results, but all you seem to have up here in this version is an option to 'search source,' which will search the whole of the Simon's direct source. But I think you can do that if you go back to here [presses back button and returns to previous page which lists the sections headings that make up Simon's]. Here you can achieve the same thing. [Ticks checkbox next to sections of interest and searches within the selected documents for 'accounting period!'].

In addition, one law lecturer had found a rather novel way of performing content searching, by installing a 'Google Toolbar' (which has a feature to allow the fast searching of document text displayed on-screen in the user's browser window):

*A27 (Lecturer): I have a Google toolbar, so I have to use the Google toolbar to search through the AustLi generated text to find what I'm looking for [pauses]. **R: What advantage does the Google toolbar have?** A27: It allows you to search the text, so that if the text is on the screen somewhere down*

there, you can use the toolbar facility for searching for particular words on the screen that has now loaded [pauses] and that's often invaluable. So even with Lexis, its search facility is so poor, in fact it doesn't have a search facility for words in the reports that it's displaying. It highlights those words, but let's say it's 120 pages of complicated text [pauses] trying to go down to find the highlighted word is just horrendous, so having the Google toolbar is quite useful.

Summary of searching behaviour

Although searching at the resource, source and content levels was quite rare amongst the lawyers in our study, searching at the document level was very common. When editing searches, the lawyers displayed more search spelling/syntax altering behaviour than search refinement, reformulating or refocusing (although these were all identified as searching behaviours in their own right). Search result sorting was also an observed, but relatively rare, search behaviour. As only a small proportion of search queries were subsequently edited by the lawyers (and therefore there is relatively little evidence of many of the searching behaviours), it is not possible to make any strong assertions about how common each lower-level behaviour might be across groups of lawyers.

5.3.4   Browsing (R, S, D, C)

Browsing involves *"semi-directed searching in an area of potential interest"* (Ellis, 1989, p.179) and was a common information behaviour amongst all groups of lawyers in our study. Previous studies by Ellis and his colleagues identified the browsing of physical shelves as a common behaviour across several domains. However, electronic browsing was restricted to electronic indexes and abstracts for the social scientists in Ellis's (1989) study. The social scientists in the study by Meho and Tibbo (2003) displayed a greater range of electronic browsing, presumably due to the increased support for browsing on the Internet since Ellis's study was conducted. However, all of these studies describe browsing at the document and content levels (i.e. browsing sources for documents and documents for content). In our study, we observed browsing at four levels – the resource level (i.e. browsing meta-resources to locate resources), the source level (browsing resources to locate sources), the document level (browsing sources to locate documents) and the content level (browsing documents to locate content).

Resource browsing (browsing meta-resources for resources)

Resource browsing was a fairly common behaviour amongst lawyers and was displayed almost entirely by academic lawyers (mostly by the sub-group of taught students). This level of browsing was mentioned and displayed in two ways:

- By browsing university library or law library pages that listed electronic journals and the electronic legal resources that carry each journal.
- By browsing the university list of electronic legal resources directly.

Similarly to when 'resource searching,' resource browsing behaviour usually involved browsing university (or law library) pages that listed electronic journals and the electronic legal resources that carry each journal (i.e. browsing to find out which electronic resource carries a particular journal title):

> *A2 (3rd year LLB student): I don't think anyone has 'Hastings Centre Report' for example. You can find that also by if you go to the [university library] main page and go to 'electronic journals,' you can see what's listed and it'll tell you where you can access them.*

The other common means of resource browsing amongst taught students was to browse the university list of electronic legal resources directly (i.e. browsing to find and access a particular electronic resource):

A3 (1st year LLB student): So then I go to 'databases' and 'law databases' and down to Westlaw [participant logs in to Westlaw].

Source browsing (browsing resources for sources)

Source browsing was also a fairly common behaviour amongst the lawyers in our study and was predominantly displayed by taught students (although some evidence for source browsing was also observed amongst research students and staff). This level of browsing was mentioned and displayed in three ways:

1. By browsing within electronic legal resources to locate a particular journal title source (which would then lead to browsing the source at the document level).
2. By browsing within electronic legal resources to locate a particular database within the library (which would then lead to searching the database at the document level).
3. By browsing university library or law library pages that listed electronic journals.

The first form of source browsing was achieved in electronic legal resources such as HeinOnline, where lawyers browsed to locate a particular journal title source within the electronic resource:

> *A4 (PhD student): Usually I use HeinOnline when I know exactly what law journal I want. **R: And in this case you don't?** A4: I am trying to find 'Modern Law Review.' But I know in this case that I'd found a textbook which mentioned an article which comes from the Modern Law Review. **R: So you're scrolling down the list of sources to find that particular journal?** A4: Yeah. [Scrolls down list and selects 'Modern Law Review Index' rather than 'Modern Law Review, which is also on the list].*

Source browsing within HeinOnline was usually followed by document browsing (i.e. locating a particular volume and issue number of the journal title series and then looking through the list of journal articles in that particular volume and issue).

Other source browsing within electronic legal resources was achieved by browsing to locate a useful database within the electronic resource. This was usually in order to then conduct a search restricted only to that particular source. This means of source browsing is demonstrated by an LLM student, who browses the Westlaw directory of sources:

> *A13 (LLM student): Let's go to 'UK multiple databases' [pauses] it's one of these side buttons. Directory [pauses] ummm [pauses] 'all Westlaw databases' and then I go to [notices choice of U.S. State or U.S. Federal materials listed in the top level of the directory]. U.S. State or U.S. Federal? U.S. Federal material [pauses] 'federal cases and judicial materials' and then 'all federal cases,' choose just the broadest database possible because I'm not exactly sure where this case will be.*

Overall, most source browsing was displayed by taught students and (as with resource browsing) involved looking down the university library or other law library list of journal titles. This list supported both resource and source browsing in the sense that it was possible to browse down the list to locate a particular journal title (source browsing) but also to find out which electronic resource carries particular volumes and issues of the journal title (resource browsing).

Document browsing (browsing sources for documents)

The most common level of browsing to be mentioned and observed was document browsing, which involves browsing sources within a particular electronic resource in order to find documents within those sources. This level of browsing was demonstrated and mentioned by all groups of lawyers, although was particularly

prominent amongst practicing lawyers. As one DR Trainee explained, browsing for documents in electronic legal resources can often be a useful alternative to searching:

> *P8 (DR Trainee): By its very nature, you're only asked to do legal research if the point isn't clear and somebody doesn't know it off the top of their head. So it can be that you think you're interpreting the search terms correctly, but the computer doesn't like it. So I find it's much more useful to see lots of information and say 'that seems like it might be vaguely useful.'*

Document browsing was mentioned and demonstrated in three ways:

1. By browsing a particular source for documents within it (e.g. browsing a particular journal title for articles within it or database for documents within it).
2. By browsing a hierarchy of related documents.
3. By browsing within a search result list to identify documents within a particular category.

The most common form of document browsing was browsing a particular source for documents within it. One example is demonstrated by a 3rd year LLB student, who talks about browsing Westlaw for articles within a particular journal title:

> *A2 (3rd year LLB student): There was a part where it said 'contents' and then it had 'crime' and it was similar to this [points to Westlaw directory] but then you could expand that and then it would have the date of the journals and then you could go to the date that you wanted and expand it. But then you've got to read every single title in that particular volume or issue of the journal.*

The above example illustrates both document browsing and source browsing (since the student describes the process of locating a particular journal title from a list as well as locating individual articles within a particular journal title). Another example of browsing a particular source for documents within it is illustrated by a DR Associate, who browsed sources within LexisNexis Butterworths (in this case when he was researching an unfamiliar area and wanted to see what particular sources 'have to offer'):

> *P18 (DR Associate): Within LexisNexis, it's Halsbury's that I use most [participant switches back to the LexisNexis browser window]. For legislation, I use the UK Parliaments Acts section and with that I find it quite easy just to browse, so if I'm just looking for a particular Act, this is a relatively quick way of getting there [pauses] when I'm new to an area, I'll start by seeing what they've got to offer me and take it from there, so go into sources and see what they've got in them, so a bit more exploring involved.*

Another means of facilitating document browsing was by browsing a hierarchy of related documents. This was somewhat rare amongst lawyers, although some evidence of document browsing (particularly within commentary sources such as Halsbury's laws) was observed. In this example, a DR Trainee browsed the table of contents of Halsbury's and located a neighbouring piece of commentary on the same topic as the previous piece that he had read:

> *P14 (DR Trainee): You can look on the left-hand side in the table of contents. It's quite useful to look there sometimes to find out if there's anything else that might be relevant. So you could look at other areas where directors may be disqualified [scrolls through Table of Contents] and because I'm particularly interested in Competition Law, then I might select this one 'competition undertakings' on the left-hand-side [clicks on neighbouring commentary piece].*

Similarly in this excerpt, a Tax Associate browsed the 'Company' section and 'Control of Foreign Companies' sub-section of Simon's Direct Tax manual in order to gain an overview of the tax rules surrounding the control of foreign companies:

R: What did you just click on there? P20 (Tax Assoxiate): That's just browsing into the 'Company' section in there, browsing into the different bits, the different sections. This resource is basically a resource for all direct corporate tax issues, so I'm probably thinking that this is going to be under the 'Corporation Tax' heading [expands collapsible tree and then decides to look in another section of Simon's]. There we go, it's under the big picture area of 'Control of Foreign Companies.'

The final way in which document browsing was mentioned or demonstrated was by browsing within a search result list to identify documents within a particular category. This form of document browsing also serves as a means of restricting search results:

P18 (DR Associate): I use these things on the left a lot. I don't know what you call them, they're trees that you can open out and look at a particular type of information, like law bulletins, legal journals, materials, commentary.

Content browsing (browsing documents for content)

Browsing at the content level (i.e. browsing documents for content) was another common level at which browsing behaviour was observed. Content browsing was observed in three main ways:

1. By browsing (i.e. scanning) through the textual content of a document.
2. By browsing through the headings within a document (or jumping to a particular heading).
3. By jumping between instances of the search terms or particular words or phrases within a document.

At its simplest, content browsing involves scanning through the text of a document (often to facilitate extracting behaviour). Because of the low-level of granularity of this behaviour, it was usually observed as skimming through textual content:

[Participant scrolls through body of case]. A9 (LLM student): Yeah, sickness insurance, power authorisation [pauses] it's more along the lines of what I'm looking for. Again, I'm scanning [pauses] [highlights relevant sections with mouse during scanning].

However, sometimes lawyers would skim through the headings within the document as opposed to scanning through the entire textual content:

R: When you went through the actual source, you just skim-read through it like you're showing me now? P6 (DR Trainee): Yeah, I tend to just skim-read through. Sometimes it's easier to go through and actually look at some of the headings [points to headings within the PDF document] to see whether they're actually useful.

As a Professor of Law explained, skimming through headings within cases becomes easier with practice, as it is possible to build a mental picture of the 'anatomy' of a legal case:

A21 (Professor): I know the anatomy of a decided case now and I know where to go. I can scan down knowing pretty well, given the logic of the judge's mind and the conventional order in which the material is presented, where it's going to come in the case.

In some electronic resources, lawyers also demonstrated that it was possible to *jump to* as well as skim through headings within documents. This Tax Trainee, for example, read the table of contents of an article in the PLC electronic resource and clicked on the 'transfer pricing' hyperlink in the table, jumping automatically to the 'transfer pricing' section of the article:

P24 (Tax Trainee): So I've loaded this practice document and we've got the transfer pricing rules there [jumps to 'transfer pricing' section by clicking on hyperlink in the document table of contents and scrolls through the section].

Finally, content browsing was facilitated by jumping between instances of the search terms or particular words or phrases within the document, as illustrated by this Tax Trainee who searched within the current document for the phrase 'capital allowances' and jumped between instances of the search terms in the document text, reading the surrounding sentences out loud:

P13 (Tax Trainee): I'd look to see if there's a discussion about Capital Allowances in this context and at this stage I'll do a word search [searches within document for 'capital allowances' and jumps between instances of the search terms]. Because it's instructions to Counsel, I know due to the nature of the document that it will go into more detail than this quick overview, so I would flick through to the next 'Capital Allowances' results [jumps to next instance of search terms in document text].

Similarly this Tax Trainee attempted to jump to the highlighted instance of her search term in a document within LexisNexis Butterworths:

P29 (Tax Trainee): I also see there's only 1 hit in the entire decision of my search term [points to '1 hit' navigation buttons at bottom of screen]. So I'm just going to go to that [clicks on navigation arrows next to '1 hit' text].

Incidentally, she was unsuccessful as the navigation buttons in LexisNexis Butterworths only allows users to jump *between* search term instances (i.e. the buttons do not allow users to jump *to* a single instance of the user's search terms).

Summary of browsing behaviour

Browsing behaviour was found to be common amongst lawyers as a whole, although browsing at the document level was far more common than browsing at the resource, source or content levels. We found that each level of browsing could be facilitated in a number of ways: Resource browsing was facilitated either by browsing university library or law library pages that listed electronic journals and the electronic legal resources that carry each journal or by browsing the university list of electronic legal resources directly. Source browsing was also facilitated by browsing university library or law library pages that listed electronic journals and by browsing within electronic legal resources to locate a particular journal title source or to locate a particular database within the library. Document browsing was facilitated by browsing particular sources for documents within them, by browsing hierarchies of related documents and by browsing within search result lists to identify documents within a particular category. Finally content browsing was facilitated by scanning through the textual content or the headings of documents, and by jumping between instances of the search terms or particular words or phrases within a document.

5.3.5   Chaining (D&C)

Chaining behaviour - *"following chains of citations or other forms of referential connections between material"* (Ellis, 1989) was found to be another common behaviour amongst the lawyers in our study. Chaining is an inherent and important part of scholarly research as explained by one law lecturer who compared it to a spider spinning a web:

A7 (Lecturer): Research always changes as you progress. But it is always the case that you read one document and that leads you to another document [pauses] so one document will refer to an earlier document or a journal document will obviously refer in its footnotes to other articles, so you do it sort of

like a spider spinning a web [pauses] your starting point leads you in unexpected directions and it makes you aware of things that you're missing and then you start a search.

Chaining behaviour was found to operate at the combined document/content level (even though it may be theoretically feasible to follow referential connections between resources or sources as well as documents). Like Ellis (1989) and subsequent studies, we identified both forwards chaining (which involves following chains of citations or other forms of referential connections between documents which have subsequently cited the current document) and backwards chaining (which involves following references to documents that have been cited in the current document). We consider this to be a pair of subtypes of document/content chaining behaviour. We also identified two other subtypes of document/content chaining, *across resource* vs. *within resource* chaining and *direct* vs. *indirect* chaining. Across resource chaining involves following referential connections between material that leads from one electronic resource to another. Within resource chaining involves following referential connections between material that exists within the current resource. In many electronic resources, chaining can often be facilitated *directly* by following hyperlinks to other referenced materials. Sometimes *indirect* chaining is necessary, where the user has to follow the references manually. This is often the case when chaining across resources.

Forwards and backwards chaining

Backwards chaining was by far the most common type of chaining observed amongst the lawyers that took part in our study. Legal documents routinely listed hyperlinked citations to previous documents and the lawyers in our study often followed these hyperlinks. For example, this Bar Vocational Course student searched for a known case involving a company called Ramsey Walkers Snack Foods Ltd. and, in the text of the case, found reference to another case that discussed the same principles and followed the hyperlink to this case:

> *A33 (BVC student): If it linked to other cases such as Linford and Carey [reads related case from screen] which is an earlier case that discussed the same principles then I'd just click on the highlighted link on the page and it takes you straight to it.*

Not only was backwards chaining achieved by following hyperlinked citations, but also by following referential connections within the text of a particular legal document. The DR Trainee below, for example, mentioned reading that the Office of Fair Trading (OFT) had published guidance on the Enterprise Act 2002 during a search for 'competition disqualification orders' and visiting the OFT website to locate this particular guidance document:

> *P14 (DR Trainee): So I could just put in 'competition disqualification orders' and that would throw up articles that had been written on it or general overviews, a definition of it, so that's quite useful if you don't know anything about the area. And that might lead me to narrow my search. For example, it might tell me that there's OFT guidance on the issue and I'd then go to the OFT website and look at them.*

Just as the study by Ellis (1989) found forwards chaining to be rare and the study by Smith (1988) found no evidence at all of this subtype of chaining behaviour, we also found forwards chaining to be a behaviour that was rarely mentioned or displayed (despite the fact that, in recent years, tools to facilitate forwards chaining have been incorporated into many electronic legal resources). However, some evidence of forwards document chaining was displayed by academic and practicing lawyers alike:

> *A17 (LLM student): One of the useful things about databases like Westlaw is that they will quote not just articles which have Franck as the author [points at example] but also articles which have quoted Franck.*

However, none of the lawyers in our study who used Westlaw (which were predominantly academic lawyers) mentioned or demonstrated use of tools within the library to support forwards chaining, for example the 'related info' tab which lists referential connections between the current legal document and other types of legal material (such as cases which cite the section of an Act which is currently being viewed).

As one DR Associate explained, forwards document and content chaining can also be achieved by using citator tools to find details of cases that had subsequently followed a previous case:

> *P18 (DR Associate): With fraud I needed to look at the whole body of case law from Derry and Peek, which was one of the very early cases, through to the present. And so I would then be looking at all the other cases that followed, which were apparent from the footnotes. We used to have a facility before our subscription changed and the layout changed on our Butterworth's service with something called 'CaseSearch.' Basically [it] enabled you to put the name of a case in and the citator would tell me every time that case has been applied.*

Overall, mention or usage of citator tools was rare. However some limited mention was made of using case citator tools to support forwards chaining behaviour. Limited mention was also made by practicing lawyers of using statute citator tools, however this was only mentioned in the context of supporting 'updating' and 'history tracking' behaviours (discussed later in this chapter), and not forwards chaining.

Direct and indirect chaining

Direct chaining is often prevented in some electronic legal resources such as Lexis Professional and Westlaw by barriers, such as a hyperlink being unavailable or only a summary instead of the full-text of a case being presented. Although this may be due to the fact that the electronic legal resource simply does not carry the full-text of the material, often there is an expectation that the full-text *will* be available and confusion can arise when it is not. One second year undergraduate, for example, was presented with the summary of a case rather than the full-text. She noticed a citation for the case mentioned in the summary but, upon performing a citation search for the case, was brought back to the Westlaw summary that she had originally indirectly chained from:

> *A24 (2nd year LLB student): I can get to the cases easily, but I can't get to see the whole thing [searches by citation using the citation presented in the summary document and is taken to the same summary document again]. It's a summary again [pauses]. I think it's taken me back to where I was actually. **R: Why do you think that it's done that? A24: I don't know. Maybe because these ones are not available to look up. I don't really understand. I don't know, sorry!***

Another Postgraduate student, however, was aware that the full-text of the article was unlikely to be found on Westlaw (which indexes a large collection of legal journal articles but does not, however, hold the full-text of all of these articles). The LLM student, instead, performed indirect document/content chaining using a different electronic resource:

> *A17 (LLM student): One of the problems with Westlaw is that they quote articles for which they don't actually have the article for. So, for example, they're quoting that Thomas Franck wrote this, but you can't actually print it out because they don't have the article online. **R: What would you do if you needed to find that article?** A17: Ummm [pauses] I'd have to find out where it is online and normally that involves going into the [university] site which lists the database or Googling it to see which online database subscribes to this particular journal. [Finds the citation in hand is likely to be on HeinOnline by scrolling down the University Library webpage]. So I'd just go to HeinOnline and pull out the reference for it.*

We also identified evidence of indirect chaining by performing segmented field searches, sometimes with the citation (as with the LLM student above) and sometimes with other details, such as the party names of a case or

the title of a journal article. The Tax Associate below performed indirect document chaining by copying a document number listed on the summary page of a document in the firm's Knowledge Management database and pasting it into the search form of another document database provided by the firm:

> *P12 (Tax Associate): I hate printing things from [the KM database] because I think it's done really badly. So I would, instead, take the number off the bottom of this [copies document number listed at the bottom of the summary into the computer's clipboard] and I would instead search for it on the firm's document database.*

Indirect document/content chaining was, however, less commonplace than direct document chaining across all groups of lawyers.

Across and within resource chaining

Both across and within resource chaining were common information behaviours amongst all groups of lawyers. As explained by the DR Associate below, across resource chaining is often necessary when looking for legal documents as different legal documents are available on different electronic legal resources:

> *P18 (DR Associate): I might be looking at one footnote to a particular point that's made in Butterworth's. That footnote might contain, say, five different cases. The point I'm making is just to find those, you could find yourself going between four sources. One of them might link straight through to something in Butterworth's. You might find one in Justis, one in Lawtel and one you'll have to go and find a hardcopy out of the library. So it's an extraordinarily inefficient system, but it is what it is.*

Another DR Associate explained that different electronic legal resources have different 'strengths' and described a potential pattern of across resource chaining where a commentary article might provide leads for potentially relevant cases, which could then be located in another electronic legal resource and a citator tool used to see whether the case is still good law:

> *P5 (DR Associate): Well I just see different strengths in the different databases, so if I want to understand something a bit better and get commentary on it then it'll be in Halsbury's. And that'll give me ideas for cases that I can search, so I'll open Justis to get the full case and print it out. Then I'll have the Citator open to just double-check if this is good law and what other cases have mentioned it.*

Within resource chaining was also commonly observed. In this example, a DR Associate examined the footnote within the 'damages' section of the 'misrepresentation and fraud' chapter of Halsbury's laws and followed a hyperlink to another commentary article (this time on the slightly different topic of fraudulent misrepresentation). The Associate also pointed out the possibility of following referential connections within the body of the commentary article to the cases that are mentioned within the article:

> *P18 (DR Associate): So my first stop, as you can see, is going to be the footnote because the footnote might point me somewhere else within Halsbury's. So let's take the 'fraudulent misrepresentation' footnote [clicks and is taken to the relevant fraud section] and these footnotes give me the background to what fraud is, what I need to prove for fraud and this tells me the meaning of fraud. If I want to check any of these points and look at them in more detail, particular if the question that I'm looking at isn't exactly answered within that paragraph, then I'm going to go into the cases [points to hyperlinks to cases in footnote] to see if in discussing the points that are set out here, any of these cases have discussed any of the issues that I'm concerned with or whether there's anything else in there that helps me with what I'm trying to get at.*

This Tax Trainee, referring to changes in 'transfer pricing rules,' highlights the fact that although many electronic legal resources provide hyperlinks to other documents within the same resource, they do not

necessarily provide links to *all types* of documents. For example the PLC article on transfer pricing that the Trainee refers to below only provides links to *articles* about particular cases that have dealt with the legal issue of transfer pricing, and not the case reports themselves:

> *P24 (Tax Trainee): So you're now looking at the changes and what they talk about here is why it's changed and how it's changed. For why it's changed, again they will reference you to certain cases that have come up and these links will link you through to various articles, only articles, on the cases and what they said.*

Summary of chaining behaviour

Document and content chaining behaviour was found to be common amongst the lawyers who took part in our study, although forwards chaining between documents was far less common than backwards chaining. Other than the forwards vs. backwards subtypes of chaining, which had also been identified in previous studies by Ellis and his colleagues, we also identified two other subtypes of document/content chaining, *across resource* vs. *within resource* chaining and *direct* vs. *indirect* chaining. These can all be considered as orthogonal sets of subtypes. Across resource chaining involves following referential connections between material that leads from one electronic resource to another (whether those connections be to previously or subsequently written material). Within resource chaining involves following referential connections between material that exists within the current resource. Document/content chaining was often achieved *directly* by following hyperlinks to other referenced materials. It was also just as frequently achieved *indirectly*, by following the references manually. This was often the case when chaining across resources.

5.3.6   Extracting (R, S, D)

Extracting involves "*systematically working though a particular resource, source or document to identify material of interest*" (definition adapted from Ellis 1989). At the content level (i.e. when lawyers were extracting documents from content), this behaviour can be better regarded as a lower-level example of 'selecting and processing' as opposed to 'identifying and locating' (because unlike identifying resources, sources and documents, identifying content usually involves a greater degree of processing than identification). This is why extracting is presented twice in table 4 and is discussed at the content level in section 5.5.4. Some limited evidence was also found of extracting at the *resource* level (i.e. systematically working through meta-resources to identify resources of interest) and at the *source* level (i.e. systematically working through resources to identify sources of interest). We do not discuss extracting at the resource and source levels in this section, although evidence for extracting at these levels can be noted through our discussion of 'resource browsing' and 'source browsing' in section 5.3.4.

There was also relatively limited evidence of *document* extracting, perhaps due to the fact that most modern electronic resources provide facilities for searching and browsing within the resource and therefore it is unnecessary to 'systematically work though' a source to identify a document, for example. Some evidence was, however, found to suggest that lawyers occasionally perform this behaviour. This undergraduate student, for example, believed that the only successful way of finding a journal article in Westlaw where the volume and issue number are known was to browse through and locate it (i.e. to systematically work through the source, selecting the volume and issue number of the journal in order to locate the article):

> *A2 (3rd year LLB student): You could expand that [points to collapsible tree] and then it would have the date of the journals and then you could go to the date that you wanted and expand it but then you've got to read every single title in that particular volume or issue of the journal. Literally the only way that you can find journals on Westlaw is through the titles and not through an author search.*

Similarly one digital law library, HeinOnline, has a prominent browsing interface that requires lawyers to systematically select the journal title, volume and issue number in order to locate a known article and therefore this is the method that lawyers used most when using this resource (incidentally, some lawyers also mentioned being unsure whether searching was possible within HeinOnline):

> *A17 (LLM student): So I was looking for a particular journal [article] in the Harvard Law Review. So I knew it was in HeinOnline and I'll now click on 'Law Journal Library' and then H [pauses] Harvard Law Review [pauses] 77.* **R: How did you know that you were looking for volume 77?** *A17: From the top of my head, I remember the article being in there. [Scrolls through and reads list, not finding the article that was thought to be there]. Ok, I take it back, but the general principle is the same. I think I've got the wrong volume.*

Document extracting was also performed, albeit to a limited extent, in LexisNexis Butterworths. This Tax Trainee, for example, browsed documents within the Simon's Direct Tax source by year and then by section in order to extract section 580 of the manual (which he was referred to in a legal textbook):

> *P22 (Tax Trainee): Here we go, Simon's Tax Cases. And the reference I had was 1994, 580. So let's go to 1994 [browses to year and scrolls down to 580].*

5.4   Accessing

Although Meho and Tibbo (2003) identified accessing as an important information-seeking activity amongst social scientists, all of their transcript excerpts relating to accessing demonstrate physical problems accessing paper-based information, such as long travel distances and difficulty getting hold of published materials within certain countries. No mention was made of access issues surrounding electronic resources, sources or documents. However, we found resource accessing, which we define as *"gaining access to resources, sources or documents/content,"* to be an important information behaviour amongst the lawyers in our study that was observed widely across all groups of lawyers. This involved accessing resources (along with the sources, documents and content within them).

We found that accessing could be *direct* or *indirect*, *visible* or *invisible*. Indirect accessing involves gaining access to a resource, source or document/content by using a third-party site or resource as a gateway (for example logging in using the educational Athens devolved login). Direct accessing involves gaining access without using a third-party gateway (for example logging in directly to a particular resource). Visible accessing involves gaining access to a resource, source or document/content through a procedure that can be seen at the interface level (usually a username/password login screen). Invisible accessing involves using recognition technologies (such as IP recognition) to gain access automatically, without a noticeable access procedure. Theoretically, it may be possible to observe accessing behaviour not only at the resource level, but also at the source and combined document/content levels (particularly if it is possible to subscribe to only certain sources within an electronic resource, or if it is necessary to pay for accessing content on a document-by document basis). However, we did not find any evidence of accessing behaviour at these levels in our study.

5.4.1   Direct and indirect resource accessing

Academic lawyers rarely displayed direct resource accessing behaviour as UK academic institutions all access the major electronic legal resource platforms either 'indirectly' and 'visibly' through a third-party gateway (known as Athens) or 'invisibly' through IP recognition technology. The exception to this was one or two academic lawyers who also worked in practice and had individual usernames and passwords to certain electronic legal resources. Due to the need for academics to log in through the Athens gateway, direct resource accessing behaviour was almost always observed amongst practicing lawyers, and involved entering username and

password details (sometimes an individual set of login details, sometimes one set of generic login details that were used firm-wide):

> *P4 (DR PDL): So I've loaded this from our Intranet and logged in using the username and password box.*

As using a third-party gateway to access electronic legal resources was only necessary for academic lawyers, this was the only group of lawyers who accessed resources indirectly through Athens. As this LLM student explained, this is the way in which academic lawyers must access electronic legal resources such as Westlaw:

> *A22 (LLM student): The reason that you log in through the university website as opposed to going through Westlaw.com straight away is that you need an Athens gateway.*

Recently, the introduction of a 'devolved Athens login' process meant that academic lawyers could log in to Athens using their university username and password details and no longer had to remember a separate set of login details. This change pleased one law lecturer, who was 'frustrated' by the 'problem of having too many passwords':

> *A6 (Lecturer): This is really a big improvement, the new Athens login, because I was perpetually forgetting my Athens password and there's nothing more frustrating than if you're sitting here looking for something and then you realise your Athens password is missing. It really was incredibly frustrating. And it was the problem of having too many passwords. Too many different things to remember.*

5.4.2   Visible and invisible resource accessing

Both visible and invisible subtypes of resource accessing were observed in our study, across all groups of lawyers who took part. Visible resource accessing often took the form of logging in to an electronic resource with a username and password (which serves to illustrate that the visible and direct resource accessing subtypes are closely related and are certainly not orthogonal). One DR Practice Development Assistant, for example, was issued with a username and password for a specialist Insurance Law resource called Insurance Day by the firm's library staff. Logging in was 'visible' in the sense that it was necessary to log in with the username and password in order to access documents within Insurance Day – it did not happen automatically:

> *P10 (DR PDA): You have to login first [logs in to Insurance Day electronic resource with personal username and password]. **R: Is this a separate username and password that you set up for yourself?** P10: This is a separate username and password that Insurance Day people have set up for me. So this is one that we've subscribed to. We've got the hardcopy and the electronic version as well, but there's only one user, so it's only me that's got the electronic user rights to this.*

As explained by the Tax Trainee below, access to the LexisNexis Butterworths electronic legal resource usually occurs automatically and invisibly, through IP recognition technology:

> *P23 (Tax Trainee): The firm has a system where you don't have to go through the normal login procedure, but it provides an automatic login system. That means that you don't have to type in your username and password every time you go into it because you might spend all day on Lexis, so it would be a bit of a fag having to log in all the time.*

However, as another Tax Trainee below explained, IP recognition is only performed when the user logs in through the firm's link to LexisNexis Butterworths, as opposed to from the main LNB website:

This DR Associate presents a second example that illustrates that access restrictions often vary depending on how the resource is accessed and that this can often cause confusion. The Associate was surprised to be faced with a login screen which, as the Associate correctly asserted, was presented because he was not using his own computer (which stores a cookie that remembers his username and password). Incidentally, this username and password was not essential for accessing the resource (as IP recognition technology already controlled access) but was primarily used to keep a log of the Associate's search trail and saved current awareness searches/alerts):

> *P18 (DR Associate): I remember starting off by [pauses] [Logs into LexisNexis and is faced with login screen]. What is this? On mine it goes straight through.* **R: So it normally doesn't present a login screen?** *P18: No. I think the first time you use it, it does. It's probably because I'm not using my computer. It must store cookies or something. So it came up with another registration screen and gave me an option of either 'register now' or 'register later' so I took 'register later' because as far as I know I'm already registered on my system and it appears just to let you go through, I presume because it recognises our IP address, so notes that we're paying our fees or whatever to use it.*

5.4.3 Summary of accessing behaviour

Behaviour surrounding the access of electronic resources has not, to the best of our knowledge, been identified in any studies of information behaviours. However, we found resource accessing to be common amongst lawyers in our study. We found that accessing could be *direct* or *indirect*, *visible* or *invisible*. Direct accessing, on the whole, was displayed only by practicing lawyers whilst indirect accessing (through the Athens educational gateway) was displayed only by academic lawyers. The 'visible' and 'invisible' subtypes of resource accessing were displayed by all groups of lawyers.

5.5 Selecting and processing resources, sources, searches, documents and content

Along with 'identifying and locating,' another broad category that we identified in our study was that of 'selecting and processing.' This category subsumes several behaviours which involve the selection of resources, sources, document and content (i.e. distinguishing, filtering, selecting and extracting). This category also subsumes the behaviours of recording (making a record of resources, sources, documents or search queries), updating (ensuring a current understanding of amendments or changes to legal documents and content and an understanding of whether a particular case or piece of legislation is good law), history tracking (ensuring a historical understanding of amendments or changes to legal documents and content and of the treatment a particular case or piece of legislation has received over time), analysing (examining in detail the elements or structure of the content found during information-seeking), synthesising (combining the elements of the content found during information-seeking into a coherent whole), collating (the physical act of drawing together documents and/or content for later use) and editing (preparing and arranging documents and/or content for later use by making revisions or adaptations).

Selecting and processing behaviour shares some conceptual similarity to Meho and Tibbo's (2003) 'information managing' behaviour, which involves *"filing, archiving, and organizing information collected or used in facilitating research"* (p. 582). However, we were able to identify several lower-level behaviours that go beyond 'filing, archiving and organising' and therefore present 'selecting and processing' as a high-level category that subsumes several other, more precise, behaviours as opposed to a low-level category in its own

right. This broad category of selecting and processing is also loosely related to Ellis et al.'s behaviour of 'ending,' which involves *"the assembly and dissemination of information or the drawing together of material for publication"* (Ellis et al, 1993, p. 365). However, our 'selecting and processing' behaviour is by no means the same as Ellis et al.'s 'ending' behaviour. Ellis and his colleagues reported that typical 'ending' activities involved finding further information at the end of the information-seeking process (Ellis et al., 1993) and carrying out small-scale searches targeted toward specific unsolved questions (Ellis and Haugan, 1997). In our study, however, these types of activities are covered under 'searching' behaviour and, for our lawyers, the 'drawing together' of material involved the conceptually different lower-level selecting and processing behaviours of 'analysing,' 'synthesising,' 'collating,' and 'editing' (which we discuss later in this section).

The behaviours of distinguishing, filtering, selecting and extracting are highly related as they all involve the selection of resources, sources, document and content. However, they are subtly different. Distinguishing involves *"ranking sources or documents according to their relative importance based on perceptions"* (definition adapted from Ellis and Haugan 1997, p. 399). These perceptions may be individual or shared and may be based on firm criteria or more fluid and subjective criteria. This distinction is part of the subtle difference between filtering and selecting behaviours. Filtering involves the *"use of certain criteria or mechanisms when searching or browsing for information to make the information as relevant and as precise as possible"* – Ellis and Haugan 1997, p. 399). 'Selecting,' on the other hand, does not involve applying concrete criteria or mechanisms as a filter on information, but *"carefully choosing resources, sources or documents as being potentially useful for the information task at hand"* (definition adapted from Oxford English Dictionary) and is heavily based on subjective perception. There is, however, overlap between these behaviours. For example, 'carefully choosing' sources might sometimes involve distinguishing between them (i.e. ranking them based on own or shared perceptions), perhaps even subconsciously. Similarly, distinguishing might involve employing criteria or mechanisms in order to perform such ranking.

There is also overlap between individual levels that these behaviours can operate at. For example, selecting content from documents is likely to involve extracting behaviour (*"systematically working though a particular resource, source or document to identify material of interest."* - definition adapted from Ellis 1989). Therefore although all four of these behaviours have attributes which make them distinct, the boundaries are less clear-cut than with other behaviours discussed in this chapter.

We now turn to discuss each of the above behaviours at the levels they were most commonly observed. This is followed by a discussion of the other behaviours which involve the selection and processing of information: recording, updating, analysing, synthesising, collating and editing.

5.5.1   Distinguishing (S, D)

Distinguishing involves *"ranking sources or documents according to their relative importance based on perceptions."* – definition adapted from Ellis and Haugan 1997, p. 399). This behaviour was not very common and was observed far less than document selecting behaviour (which is discussed in section 5.5.3). This may be due to the fact lawyers did not tend to 'rank' resources based on their perceptions of how useful they might be, but instead selected them based on a number of 'hard' and 'soft' criteria. It might also be argued that once a lawyer has selected an electronic resource to search or browse within, it is often unnecessary to distinguish between individual sources within that particular resource. This is especially the case if the lawyer is conducting a resource-wide search for UK cases or legislation, for example. Similarly lawyers usually decided whether a particular document was likely to be useful for their information search *without* ranking the various possible documents that might be useful. For example, it was rare to observe a lawyer examining several search results on a results list and deciding to prioritise certain results over others, as explained by this Tax Associate, who stated that, for him, it was only necessary to distinguish between documents (i.e. rank them according to their relative importance) when there are lots of search results:

R: Was it the case that you picked the most likely one to be relevant by weighing them up? P21 (Tax Associate): No, it was a more gradual step-by-step thing. I looked at the first one and decided it was not relevant, then I looked at the second one and decided it was not relevant. But that was only because there was 4 of them. If there was like 30 I would have probably gaged them all against the other.

Despite the relative lack of evidence of distinguishing behaviour, there were some occasions in which source distinguishing was mentioned or demonstrated to be necessary. Distinguishing was achieved in three main ways:

1. Based on the perceived authority of the content within the source (i.e. the reputation of the source).
2. For legal cases, based on the level of court in which the case was heard.
3. Based on the date which the document or source was produced.

Unlike with Ellis's social scientists, who used the generality or technicality of documents as important distinguishing criteria, lawyers always distinguished between sources based on the perceived authority of the content within the source (i.e. the reputation of the source). For example, one PhD student scanned his journal article search results in order to find which ones came from 'credible' law journals:

A4 (PhD student): Searching journals is for me a bit tricky because there are not a lot of credible law journals. There are a few journals which you find are just basically law news. They are ok to read, but there is no academic value in citing them. Like New Law Journal is good if you want to learn new things, but we do not find it very credible to cite. Nor do we find it useful for discussion. So basically I tend to read through the title and what journal it is in.

When it came to looking for other types of legal documents, no lawyers mentioned or demonstrated a need to distinguish between sources containing legislation, probably because competing sources all display identical wording of statutes and statutory instruments as passed by the relevant government. However, some source distinguishing could be noted when deciding when to use one particular report series over another. One DR Paralegal explained that sometimes it is possible to distinguish between sources based on the interpretation presented by writers of each case report series and suggests the need to read more than one report of the same case (i.e. from different report series):

P3 (DR Paralegal): Umm. I suppose if I wanted to have a particular case and wanted to check how I understand, how I interpret what they said, then I might just look at the second website - a different database. And obviously the wording would be different and it might help me. R: Is that because it would be from a different Report Series? P3: Written by different people. So you might have an interpretation and see that both of them have picked this up. It is just easier to make sure. This is how we were taught, we were told to use both in conjunction to be sure sure.

Similarly this Bar Vocational Course student talks of a 'hierarchy' of law report series based on their authority:

A33 (BVC student): Seven cases have come up [points to results list] and two of them appear to be the same, from the AIT [pauses]. So I'll take the one that's been published in the All England Reports. R: Why did you choose to look at that version? A33: I chose it over the other alternative which was the AIT case itself because of the hierarchy of journals that you should use when citing cases. All England Reports are higher than the other options that were there.

The concept of valuing certain sources above others when looking for cases is also discussed by the Tax Trainee below, who explained that it is preferable to choose cases from the 'Simon's Tax Cases' and 'All England Reporter' sources above other sources, particularly digest sources such as 'Specialist Case Digest,' which only presents summary versions of case transcripts:

R: How did you know to pick that version of the case and not another version? P24 (Tax Trainee): There's not a lot of difference. In this department we tend to go for either of these two, to be honest [returns to results and points to the cases from 'Simon's Tax Cases' and 'All England Reporter']. Simon's Tax Cases is extremely well respected and very strong authority and so is All England. Equally I could have clicked through to that one. What I'm looking for is the transcript of the case, we're now looking at the specific detail. R: So they shouldn't really vary? P24: They shouldn't really vary at all. What we're trying to do now is look at exactly what was said and then we put our own interpretation on how it came across. This 'Specialist Case Digest' version of the case you wouldn't use in practice. These two are the strongest - Simon's Tax Cases and All England Reporter.

Other lawyers were less concerned about which report series they choose. This DR Associate suggested that the decision is usually arbitrary, although for practical reasons it is sometimes useful to follow a citation to a particular case report in order to make chaining easier:

P18 (DR Associate): It doesn't always matter if you go to exactly the same citation. You can look at the same case in a different report series if it happens to be reported it two places, but it is useful if you're working from Halsbury's or from a textbook because a lot of the time they'll say in the in the citation is '92 at 103' and 92 is the page in the law report book where the report starts and 103 tells you the page on which the point being referred to is made. So that's why sometimes it is worth going to exactly the citation you're given, because that then tells you exactly where you want to be looking.

Similarly this 1st year undergraduate student suggested that differences between cases from different report series are likely to be minor:

A3 (1st year LLB student): If it's gone through the courts, you'd look at the one from the highest court to look at the decision - so the one with the latest date. But otherwise, if it's just been reported in two different places then it doesn't make much difference because they're both written as the judgements came. So there shouldn't be any differences and if there are they will probably only be minor semantic differences.

Aside from distinguishing legal cases by source (i.e. by report series), lawyers also distinguished between cases based on the level of court in which the case was heard. As the Tax Trainee below explained, it is possible when searching LexisNexis Butterworths for cases to retrieve results that list the same case, but heard at different levels of court. Often lawyers were interested in the final decision that was made with regard to a particular case (after it has escalated through various levels of court) as opposed to previous decisions that may have since been overruled by a higher level of court:

P29 (Tax Trainee): You can also see that numbers 2 to 6 in the results list are all the same case, just at different various levels. So I don't need to read the lower-level decisions, so I'm going to go into the highest level that it's been decided on, which is the Court of Appeal. [Clicks on case title].

A final example of distinguishing documents was by date. In the excerpt below, a Tax Trainee looks for a legal definition for the term 'transfer pricing' in the PLC electronic resource. As several versions of the definition are presented in the search result list, each with a different date next to it, the Trainee looks at the most recent one and not the others:

P24 (Tax Trainee): This is another feature that Lexis and other places don't have. This is where anything that hasn't got anything after it [points to lack of descriptive metadata in the search results list, other than 'transfer pricing' and the date] then we know it's a definition, which senior people don't need. Just to check we're in the right area. Let's go to the most up-to-date one. [Clicks on definition of 'transfer prices' with the most recent date].

5.5.2 Filtering (D, C)

Ellis's behaviour of filtering involves the *"use of certain criteria or mechanisms when searching or browsing for information to make the information as relevant and as precise as possible"* – Ellis and Haugan 1997, p. 399). Filtering documents, often by restricting searches by various aspects of meta-data, was common amongst all groups of lawyers. Filtering *content* was less common amongst lawyers and was often achieved by highlighting instances of the search terms entered in the document text or searching for particular words or phrases within the document text.

The most common way of filtering documents was by date (e.g. the date a particular legal case was heard, piece of legislation came into force or journal article was published), as illustrated by the PhD student in the excerpt below:

> *A4 (PhD student): What we have here is 258 cases, which is by no means workable for any lawyer or researcher. So [pauses]. Maybe I will try to limit it to the previous ten years because I only need to know the recent developments [restricts search by date].*

Another common way of filtering documents was by document type, as illustrated by this Tax Trainee, who explained that it is possible to filter search results by 'resource type':

> **R: You clicked on 'resource type.' What will that do?** *P22 (Tax Trainee): Resource type would separate them into memoranda, legal opinions, counsels' opinions, things like that. At which point you could filter out most of the transactional documents again and get to ones that were more legal opinion.*

A similar example is illustrated by another Tax Trainee, who conducts a search in LexisNexis Butterworths restricted to legislation:

> *P22 (Tax Trainee): So I was going to look on LexisNexis. [Searches for '81(3) AND "finance act 1996"]. I'll just tick the legislation box.* **R: You ticked that box because the Finance Act is a piece of legislation?** *P22: Exactly. There's not much point in getting all the commentary up on it.*

Another way of filtering document was by legal area. This Tax Trainee used pre-set search preferences to restrict searches within the firm's Knowledge Management database to documents provided by his own department (thereby ensuring that all documents will be Tax-related):

> *P22 (Tax Trainee): So if I go with '81(3) AND "issue."' [Submits search]. Oh, hang on, the Tax filters are off. [Turns 'preferences' on].* **R: What did you just do?** *P22: We have preferences within [the KM database] which we set up ourselves when we join the department which means that it's just drawn from within the Tax part of [the KM database].*

Similarly this Tax Associate set up preferences that filter results not only by the department within the firm that uploaded the documents in the first place (such as 'Tax' or 'Corporate'), but also by document type (the Associate also has a preference set up that excludes newsletters and memos):

> *P12 (Tax Associate): In [the firm's KM database], you can create 'preferences' which are sort of refined searches. So for Tax, that's just all of the Tax documents, but I have lots of other preferences [loads preference list]. I've got a 'Corporate,' 'Corporate intra-group,' 'Tax with no news' and if I know that I'm looking for a point that I know is most likely to come up in the Corporate department, because it's their type of work, then I'll take my 'Tax' preference off and just search within 'Corporate.'*

The final way of filtering documents was by jurisdiction, as explained by the same Tax Associate. She explained that it is possible to filter results by the countries which the law in the document affects. However, in this example she did not find the filtering process useful as she believed that most of the results had been

incorrectly classified as 'Worldwide Law' rather than 'UK Law,' even though, upon closer examination, they appeared to deal with UK issues:

> *P12 (Tax Associate): Once you've got these results you can go to 'relevant law,' for example, which is not going to be particularly relevant in this case and actually, we always sort this fairly badly [points to text that indicates that most of the current results have been classified as 'Worldwide Law' on the system rather than 'UK law']. Most of these documents should be under 'UK law' because this section of legislation that we're looking at is UK statute, so I'm not sure why it says 'UK law' 1 result and 'Worldwide law' 10 results out of 11, because as far as I'm concerned it should probably be 'UK law' 11 results out of 11. But if I was doing something were I thought it was more general, it wasn't a specific piece of UK legislation, then I might say 'there's no point looking at the 'French law' results,' for example, 'I'll get rid of those.'*

Filtering *content* as opposed to documents was often achieved by highlighting instances of the search terms entered in the document text or searching for particular words or phrases within the document text (actions which also facilitated content browsing and content extracting behaviours). Excerpts to illustrate these actions are presented elsewhere in the context of content browsing and content extracting.

5.5.3 Selecting (R, S, D)

The behaviour of 'selecting' was identified at the resource, source and document levels and was found to be common across all groups of lawyers that took part in our study. We define selecting as *"carefully choosing resources/sources/documents as being potentially useful for the information task at hand"* (definition adapted from the Oxford English Dictionary). At the document level only, lawyers were found to select documents *directly* (by examining the full-text of the document) or *indirectly* by examining various aspects of meta-data about the document. The direct-indirect pair of subtypes was not identified at the source level due to the fact that although some limited evidence was found of lawyers selecting sources indirectly by examining meta-data about the source (such as the temporal coverage of the source), no evidence was found of lawyers searching or browsing within a particular source in order to decide whether the source contained documents of interest and should be searched or browsed further. The direct-indirect pair of subtypes was also not identified at the resource level due to the fact that the criteria used to select resources did not fall neatly into the categories of 'direct' and 'indirect.'

Resource selecting

The criteria that lawyers used to select resources were different in nature to those that they used to select or distinguish between material at the source and document levels. This is because selecting which electronic resource to use to find information is often based on opinion and perceptions about the resource. Lawyers tended to select rather than distinguish between resources, perhaps due to the fact that no two resources look or function the same, nor do they contain exactly the same materials. The criteria users used to select which resource to use have not been discussed by Ellis and his colleagues, nor by Meho and Tibbo. We found that lawyers selected resources based on several *hard criteria* and several *soft criteria*.

The hard criteria that were identified were:

- The subject and nature of the content of the resource.
- The structure of the content.
- The perceived authority of the content.
- The perceived comprehensiveness of coverage of the content.
- The perceived cost of accessing the content.

The soft criteria that were identified were:

- The perceived ease of use/simplicity of the resource.
- The perceived speed/time savings offered by the resource.
- Prior positive experiences that the user has had with the resource.
- The user's familiarity of the resource.
- Whether the resource had been recommended by others or not.

'Hard' (content related) criteria

The first set of criteria lawyers used to select electronic resources can be described as 'hard' criteria. Although by no means objective in nature, these criteria are based on reasonably concrete criteria to do with the content of the resource. We now discuss each of these criteria.

Subject and nature of content

The first hard criterion used when selecting resources was the subject and nature of the resource content. One PhD student frequently visited a specialist government website to find specific types of materials, such as transcriptions of parliamentary debates and other government documents that may not be available from other sources:

> *A12 (PhD student): I use lots of parliamentary documents, so that would be in the UK is called Hansard, which is the transcription of parliamentary debates online. I look at that a lot. I look at government documents such as white papers and explanatory memorandums and bills - so really the bill that is sent.*

Along with specific types of materials, some resources are also particularly suitable for finding information on specific *subjects*. For example this 3rd year LLB student explained that LexisNexis Professional was likely to be unsuitable for finding information on her research topic of legal and economic issues surrounding the British Broadcasting Corporation and its license fee:

> *A2 (3rd year LLB student): This is why I wouldn't really use Lexis unless I was stuck, because there's no media-related journals I don't think. So I'd tend to use Lexis mainly for cases because they've got the majority that you'd need.*

Instead of LexisNexis Professional, the LLB student decided to use Google Scholar to find information on the BBC license fee.

Structure of content

Another hard criterion which lawyers took into account when selecting resources was the structure of content within the resource (i.e. how the content is displayed on-screen and on-paper when printed). Sometimes particular electronic resources were deemed to display content in more useful ways than others. For example, several lawyers in both academia and practice spoke of being frustrated by the fact that Westlaw displays and facilitates acts of parliament section by section and does not have the option to view or print the entire act at once. Incidentally, this is a misconception as Westlaw *does* allow the full-text of an Act to be viewed and printed in PDF format, provided the user is aware of this function. However, just like many of the other lawyers in our study, this Senior Research Fellow is unaware of this facility and instead turns to the Office of Public Sector Information Website to access entire acts of parliament:

A25 (Researcher): With legislation, the way it's set out actually on Westlaw is quite often in sections and so if it's English then I'll go back and have a look at the Stationary Office Website because it's easier just to pull Acts off there.

Sometimes the aesthetic aspects of how content is displayed on-screen could also play a role in the decision of which resource to use for finding certain types of material, as this LPC student pointed out. Once again, this led to the student using the OPSI website to look for legislation rather than one of the electronic legal resources:

A28 (LPC student): I personally find it a bit difficult to read legislation when it comes up in the Lexis or Westlaw site. The print might be a bit too small or in a different format. So once I've found that piece of legislation, I'll often go to the HMSO [Her Majesty's Stationary Office, now OPSI] website and read it from there because I often find it easier to read, easier to follow I suppose. Lexis and Westlaw have links in everywhere [hyperlinks]. Sometimes you just want to read the law, because it's complicated as it is. The HMSO website lets you do that.

Similarly, when this DR Associate was a trainee, she was directed to use Lawtel rather than Justis due to the fact that the printed output from Lawtel fits more of an A4 page than the output from Justis:

P18 (DR Associate): When I was a trainee and I was working for particular people they would probably prefer me to use Lawtel for quite a strange reason which is that in Justis, you'll notice that the pages don't fill the screen and when you print that off, that will only be half a page.

Perceived authority of content

Another hard criterion used to select resources was the lawyer's perception of the authority of the content within the resource. For example, one final year undergraduate student explained that Google Scholar was likely to provide more authoritative content than the regular version of the search engine as this content will be academic in nature:

P2 (3ʳᵈ year LLB student): I find that if I type into just Google plain and simple then I get a lot of things that aren't really specific or of much use to me, but if it's academic then I know that it's of a certain reliability.

Authority is also a means of distinguishing between sources, and was discussed in section 5.5.1.

Cost of accessing content

The penultimate 'hard' criterion used to select resources was the cost of accessing the content within them, which was only discussed by practicing lawyers (probably because academic lawyers incurred no direct costs for accessing any of the resources that they used). This DR Practice Development Assistant explained that he often starts off an information search by consulting pre-paid resources (i.e. those where the firm has already paid a subscription for unlimited use) before going on to use resources that charge on a per-search basis:

P17 (DR PDA): You start off with what I call the 'free' sources, which are the once we've already paid for and then end up, if you want to make absolutely sure then for cases you'll go into Lexis and that's sort of the most reliable source.

Comprehensiveness of coverage

The final 'hard' criterion that lawyers mentioned or demonstrated for selecting resources was the comprehensiveness of coverage that the resource provides. Such comprehensiveness may be based on:

- The type of legal materials contained within the resource (e.g. case reports, statutes, statutory instruments, commentary, articles).
- The jurisdictions covered by the resource (i.e. which countries' laws can be found by using the resource).
- The time periods covered by the resource (i.e. how far back the resource catalogues a particular type of legal material).

Two examples of resource selecting based on the *type of materials contained* in a particular resource are presented by a law lecturer (who preferred to use the Office of Public Sector Information website when looking for statutes and statutory instruments) and a first year LLB student (who preferred to use LexisNexis Professional and Westlaw to look for UK cases). These preferences were both due to the fact that the academic lawyers' perceived these resources to be comprehensiveness in terms of coverage for these types of materials:

> *A7 (Lecturer): I don't use Westlaw for legislation or statutes. I either go direct to the source, so if it's the government I go to the Houses of Parliament [OPSI] site and get the full-text because they have a very comprehensive site for statutes and statutory instruments.*

> **R: Why did you turn to LexisNexis Professional and Westlaw?** *A1 (1^{st} year LLB student): Because they are the only two resources offered by the university that contain the most [pauses] vast amounts of cases, so it's more likely that I would find the information there.*

An example of resource selecting based on the *jurisdictions covered* by a particular resource is presented by a DR Practice Development Lawyer, who uses different electronic legal resources depending on the jurisdiction of law she is looking for information on:

> *P4 (DR PDL): For English cases you'd be looking at Butterworths and CLI and Lawtel and if it's international cases, depending on which jurisdiction, you might look at Lexis or WorldLi.*

Finally an example of resource selecting based on the *time periods covered* by a particular resource is presented by a 1^{st} year undergraduate student, who uses a particular electronic legal resource (HeinOnline) when looking for older cases:

> *A1 (1^{st} year LLB student): Sometimes for very old cases, like in the 1600s, I think there's also another source that's on the [university] website. I think it's called HeinOnline, but I'm not sure about it. But in some rare situations, I'll use it as well.*

'Soft' (subjective and social) criteria

The second set of criteria used to choose between resources can be described as 'soft' criteria. These criteria were far more subjective than the 'hard' criteria discussed previously and were often influenced by personal preference and perception and potentially by the wider social context of information-seeking within the workplace. We now discuss each of these soft criteria in turn.

Ease of use/simplicity

When selecting resources, an important soft criterion is the perceived ease of use and simplicity of the electronic resource. This second year undergraduate explained that Google 'made her like computers' due to its simplicity, approachability and the perception it gives of user control:

> *A11 (2nd year LLB student): You get summaries on Google as well, you get explanations - certain terms that you can't understand, you get so many definitions just like that. It's just simple. **R: In what way is it simple?** A11: It's gonna sound funny, but even on its own website it's just so simple. I don't like computers. I used to hate computers. So Google is something simple and looks approachable to me. I don't like heavy websites like the [university] website. It used to scare me off because it's got so much colour and so much information. With Google, you just define everything [pauses] you're in control with Google, well that's what you think anyway, and I like that. **R: In what way do you think you're in control?** A11: You can define everything, you can choose everything. **R: Are we talking about choosing results or choosing what to type in?** A11: Choosing what to type in. Even specialising like with Scholar [pauses] that I discovered like two months ago and I'm like 'woah' a specialised Google. It's just easy and I like it! Google made me like computers!*

As one law librarian explained, ease of use when searching Google makes a large contribution to the phenomenal popularity of the search engine when looking for information on the Internet:

> *A16 (Librarian): I think law students are the same as all other students, are the same with all other people who aren't involved in the information profession. They just think that Google is a gift from heaven and it's fabulous. **R: Where do you think that view stems from? What exactly is it about Google?** A16: Ease. Ease of use. Solely and specifically ease of use. One box, search terms in, voomph! Twenty seconds later, results back.*

Speed/time savings

Another soft criterion which lawyers use to select resources is the speed in which material is likely to be found and/or perceived time savings if one resource is used over another. Once again Google was mentioned as a positive example of a resource which helps users to locate material faster than the alternatives, especially when the lawyer is pushed for time:

> *A6 (Lecturer): If I've lectured between 11 and 12 and I have another one between 1 and 2, I come here for an hour and I answer my e-mails for fifteen minutes. Then if I'm looking to pull up research using the time available to me, I'm really only playing with half an hour. And as a consequence, three or four minutes of clicking on things is rather a lot of time. It's often quicker for me to go to Google for example. Google still remains a very effective way of finding material. I'll often just use Google because it will often find something quicker than almost anything else [pauses]sometimes. If I've enough information.*

Prior positive experience

Lawyers also regarded prior positive experience with a particular electronic resource to be a reason for choosing to use one resource over another. Lawyers spoke of certain resources providing 'better,' 'more reliable' or 'more useful' search results. For example, this DR Associate explained that she now prefers to use Lawtel rather than Justis when searching for legal cases. This is based on the prior experience of not finding any relevant cases when using Justis, and finding relevant cases when typing exactly the same search query terms into Lawtel:

P5 (DR Associate): I've just found recently that if I've just got search terms, rather than knowing what case I'm looking for, it's always brought some things up, whereas Justis doesn't always do that. Justis isn't, I think, as effective when you're just using terms. It might give you hundreds of cases that are irrelevant or just not recognise the terms. But Lawtel seems to. And it might be fluke, fluky, but it seems to always give me something. There's an area I was researching recently on indemnities and I couldn't get any cases anywhere, but then I pumped the same search terms into Lawtel and got two cases that were helpful. So it seems to be that the search mechanism is useful. So I think to myself in future when I'm in a similar situation using search terms, I'll go to Lawtel first.

Familiarity

The penultimate soft criterion used by lawyers to select resources is familiarity. Several lawyers cited familiarity with a particular electronic legal resource as a reason for choosing to use it to find certain material (when they felt as though the content would be available on more than one electronic resource):

A26 (Lecturer): I've become so accustomed to using Westlaw that I haven't gone back to Lexis.

R: How do you decide between when you will use Lexis and when you will use Westlaw? *A28 (LPC student): It used to be completely at random, but now I'll now start with LexisNexis, purely because I've had a bit more practical use of it.*

As one law lecturer explained, familiarity can indeed play an important role in the perceived usefulness of a resource. The lecturer recalled incidents where colleagues, less experienced than him at using Google, would think that he had 'miraculous powers' when actually he attributed his search success to familiarity with how to search Google that had been gained through experience:

A6 (Lecturer): Colleagues will come up to me and say 'I can't find this South Africa case anywhere' and they know it's vaguely, they've got some sort of approximate name, perhaps about bus fares in Cape Town. And they say they can't find the case anywhere on Westlaw etc. And I'll shove the terms into Google and on the first page or so it'll come up or I'll find a reference that will enable me to locate the case [pauses] and they'll think I have miraculous powers. It's because they've never spent time on Google learning to refine search terms and playing around with search terms. And maybe trying seven different combinations and going to several pages on each before finding them.

As explained by a law librarian, firms who develop electronic legal resources rely on their users gaining familiarity in order to secure their future use of the system and hence future revenue for the firm. Electronic legal resource firms also make student representatives available to help their peers find legal information using a specific digital law library, a practice which this particular law librarian is rather sceptical of:

A15 (Librarian): You do get people coming through who have already got a specific preference for one or the other, usually because their university law department could only afford one or the other. But they have student representatives now for Lexis and Westlaw kind of milling around law departments. It's quite frightening! It's like the mafia or something! I do think that's a definite policy of trying to hook students in at an early stage.

Recommendation

The final soft criterion identified in our study to influence resource selecting is that of recommendation, which may come from a member of academic staff (for academic lawyers), library staff (for both academic and practicing lawyers) or a colleague (for both academic and practicing lawyers). For taught students in the early

years of their academic careers, recommendations for which electronic resources to use are 'prescribed' by teaching staff:

> *A8 (3ʳᵈ year LLB student): Sometimes I surf the [pauses] teachers give me some specific website links to find articles, but I don't actually look at other big databases of legal information unless a tutor tells me to do so.*

In another example, one law librarian recommends Westlaw to students for finding certain types of legal information because he prefers the resource himself over its competitors:

> *A20 (Librarian): I personally prefer Westlaw. I prefer the look and feel of Westlaw, so I tend to use it more when I'm demonstrating and where I'm recommending [pauses] where I'm showing people where to go for material. But that's just a personal preference.*

For practicing lawyers, recommendations of which electronic resources to use in order to find particular types of legal information are primarily provided by Professional Support Lawyers (i.e. legal information specialists) and members of library staff:

> *P6 (DR Trainee): For the question of knowing whether you're using the right resource, sometimes you'll get pointed in the right direction because the person giving you the research is higher up than you, knows more than you, and will say 'oh, I think you might find something on this, this and this.' But I think in terms of selecting appropriate electronic resources, you're expected to do that yourself. I wouldn't really bother senior people within the firm asking about that. I would probably go to the Support Lawyers and the library staff if I was concerned if I wasn't using the right resources and would ask them if they had any ideas. But selecting appropriate resources is something we're expected to know how to do.*

Source selecting

Little evidence was provided of selecting distinctly at the source level. This might be due to the fact that sometimes decisions about choosing which sources to use to find information in are made at the resource level. For example, some lawyers chose to use particular resources due to the comprehensiveness of coverage they provide (i.e. due to the fact that they carry a wide range of sources). Sometimes lawyers do not have a choice of which source to choose, particularly when the information required can only be obtained from one particular source. This might explain why there was more evidence of source *distinguishing* than source selecting (i.e. there was only a true choice about which source to select when a number of alternative sources, each with overlapping or identical coverage were available within a particular electronic resources and, in that case, it was the lawyers' natural instinct to rank the value of the sources and choose *between* them).

Where source selecting was demonstrated, this often involved using criteria such as the perceived authority of the source and the coverage of the source (discussed in section 5.5.1). Other times lawyers selected sources simply because they needed to find material that they knew would be found within them. This is illustrated by the Tax Associate below who found that, when looking for information about the Finance Act 1996, many of the references he found pointed him towards the 'Company Tax Manual' source, which led him to search within it:

> *P21 (Tax Associate): All the references are pointing to one particular manual out of 25 or 30, called Company Tax Manual. [Selects 'Company Tax Manual' source and begins to search within it]. I'll try to search it for "loan relationship" AND "participator" AND "paragraph 2." [Conducts search].*

The Tax Trainee below illustrated that, sometimes, there are no concrete reasons for selecting a particular source to search or browse within, other than in hope that it might contain useful information:

[Participant searches for 'European law AND avoidance' restricted to only the 'Taxation Magazine' source]. P7 (Tax Trainee): This result came up before in the general search. So I probably won't look at that. So based on the headings, I don't really see anything useful. **R: Why did you pick 'Taxation Magazine' as something to search in?** *P7: Just because it was the first on the list! [Laughs]. That was simply the reason.*

Document selecting

Just as Ellis's (1989) social scientists distinguished between information *sources* by substantive topic, our lawyers selected d*ocuments and content* by legal topic or issue. Lawyers also, like Ellis's social scientists, selected documents by quality, level and type of treatment (but in slightly different ways). For example, whilst the authority or reputation of journal titles was important in both the social science and law domains, lawyers tended to select cases based on the level of court in which a case was decided rather than the report series it was written in. In addition, like the Chemists in Ellis et al.'s (1991) study, our lawyers also selected by author and like the English literature academics in Smith's (1988) study, our lawyers also selected based on currency (more specifically, based on whether the laws mentioned in the document were likely to be still in force). In addition, we identified several further criteria for document selecting (mostly based on meta-data about the document), including the date in which case reports were published or legislation was introduced or updated and the title of the legal document (or the party names for cases) to name just two. The full list of criteria are listed later in this section. We found that document selecting could either be *direct* or *indirect*. Direct selecting involves looking at the actual content of a document when carefully choosing it as being potentially useful for the information task at hand. Indirect selecting involves using meta information about the content (such as a summary or a results list or snippet) when choosing it as being potentially useful.

Direct document selecting

Reading the actual content of documents to decide whether they might be useful was rather rare and indirect document selecting was observed far more than direct selecting. However, when lawyers *did* decide to consult the text of a particular legal document in order to decide whether it might be useful, they tended to skim read it in order to determine whether the legal issues covered were relevant to the information-seeking problem at hand. One example was from a DR Paralegal who was asked by a client to find out whether there were any restrictions of use of any materials for the packaging of the product, specifically, whether there were any restrictions of use of certain 'heavy metals.' The paralegal read through the text of Article 6 of the Packaging Essential Requirements Act 2003 in order to try to find the answer to this question:

> *P3 (DR Paralegal): I think it was Article 6 that dealt with the issue that we're interested in. Yeah, I remember.* **R: And you remember that Article 6 was useful from when you did this search previously?** *P3: Yeah. So that section dealt with concentration levels of regulated metals present in packaging. So this is totally relevant to the question that the client asked us about. And here, we see the exact provisions and also the legislation gives us dates as to when and how the companies, producers, manufacturers, whoever need to reduce the presence of metals in their packaging. And here [points to one of the provisions in the Article], it says 'regulated metals.' And if you go to the definition of 'regulated metals' [scrolls to definition section of the Act], I think it includes 'heavy metals,' which we were asked about.*

Just as the DR Paralegal above chose to only read the Article of the Packaging Essential Requirements Act 2003 most relevant to the information-seeking problem at hand, very few of the lawyers who mentioned or displayed direct document selecting behaviour read through the entirety of a legal document. This LPC student, for example, skipped to the dictum of the case report he was reading which, as he explained, is the part of the report where the commentator summarises what the judge has ruled in the case and other judges state whether they agree with his decision:

A28 (LPC student): The most important thing in law is the dictum - what the law justices have said about the case [pauses] because you use that dictum to form your argument and ideas. You take that dictum as being 'the law' and it lists out that Lord Justice Riggs has given about fifty paragraphs of commentary about the case [pauses] sorry 34 paragraphs [pauses] and the last two Lord Justice, Lord Justice Longmore and Gibson agree with what Lord Justice Riggs has said.

Indirect document selecting

Indirect document selecting was accomplished in a variety of ways, all of which involved using meta-data about the legal document rather than the full content to help decide whether the document may be useful. These ways included selecting based on:

- The level of court that a particular legal case was reported at.
- The date that a particular case was heard, piece of legislation was introduced/amended or journal article was published.
- The title of a particular legal journal article or piece of commentary or piece of legislation and the party names involved in a particular case.
- The author of a particular legal journal article.
- The source in which a particular legal journal article was published.
- The headnote or summary of a case or abstract or contents page of a journal article.
- The keywords or index terms used to describe a legal case, piece of legislation, journal article or piece of commentary.
- The context in which query search terms are mentioned within the document.
- How many times search query terms are mentioned within the document.
- How relevant the electronic resource deems the document to be (e.g. the percentage relevance rank assigned to the document).

We now turn to discuss each of these ways of achieving indirect document selecting.

Indirect document selecting by level of court

The first way in which lawyers selected documents (in this example, legal cases) was by the level of court in which the case was heard. As this 3rd year undergraduate student explained, the UK court system is hierarchical and therefore cases that were heard in different courts are likely to vary in importance:

A2 (3rd year LLB student): There's different stages in a case; there's the trial stage which doesn't tend to be reported, but then if there's an appeal it will go to the Court of Appeal and you can also appeal the Court of Appeal's decision in the House of Lords, so they'll be reported in different areas.

Indirect document selecting by date

Another way in which lawyers performed document selecting is by date. This DR Trainee, for example, asserted that when looking at cases, she prioritised those cases that are most recent (and were heard in a higher level of court):

P6 (DR Trainee): The other thing I would do as a lawyer when I'm sifting through cases is see how recent they are, because quite often you would get tons and tons, so you would prioritise the most recent ones and the ones that are higher authorities, those that have been in the House of Lords or Court of Appeal rather than in the lower down courts, because they have more weight, they're more important.

Also related to selecting by date, this 3rd year LLB student was looking for a case that her lecturer had referred to as the 'MG baby' case. She expected the would have only been recently reported in the last couple of days and therefore, when she found a case from 2005 with 'MG' as one of the party names, she decided not to look at it:

> *[Finds a case with 'MG' as one of the party names]. A2 (3rd year LLB student): I don't think this is it.* **R: Why don't you think this is it?** *A2: Because this case was reported in 2005 and the decision was yesterday I think.*

Similarly this DR Practice Development Assistant was aware that the Company Law Reform bill had changed recently and therefore also declined to look at a document from 2005:

> *P10 (DR PDA): And this one is from 2005, but the Company Law Reform bill has been changed recently and there have been some new changes brought in this year, so really looking at something from 2005 is probably not going to be very reliable.*

Indirect document selecting by title (or for cases party names)

Lawyers also selected based on the title for legislation and journal articles and by the party names for cases. As asserted by DR Practice Development Assistant P19, "*you shouldn't judge a book by its cover, but I do judge an article by its title, especially on this when you're doing a search!*" This assertion was echoed by many lawyers. The PhD student below used the titles of Acts returned in his search results to perform selecting behaviour:

> *A4 (PhD student): So we have 18 results and I will look at them one by one. I know that the Proceeds of Crime Act 2002 has nothing to do with me, so I won't look at that. But I'll look at the Pensions Act.*

Similarly the 1st year LLB student below used information about the name of parties in the title of a case to select potentially useful cases. In this case, because the case included the abbreviation 'DDP' (which stands for Director of Public Prosecutions) and the party name 'Soerl,' the law student was reasonably sure that he had found the correct criminal case:

> *A5 (1st year LLB student): Yeah, I've found it [the Soerl case].* **R: How did you know that you'd found it? Because it's the right year?** *A5: Yes, and because it's the Crown against Soerl as opposed to Soerl and Soerl. I just did [searched for] all the ones with Soerl in the title. All the other ones are civil cases, but this one is a criminal case. It's actually a bit more complicated than that. You know it's a criminal case when it's Crown vs. Soerl, but in old cases in a Magistrate's Court you'd have the name of the chief constable as the defendant or nowadays the equivalent of that [the Crown] is DDP - the Director of Public Prosecutions. [Selects case].*

This Tax Trainee also selected documents to read in more detail by reading the titles of the commentary articles in the search result list:

> *P13 (Tax Trainee): This one has come up with the result 'Annuities provided by partnerships' [reads out title of first hit], which suggests to me that it's not dealing with Capital Allowances. So I would look mainly at the headings, I guess, to establish whether or not it's what I'm looking for. So this heading here, 'husbandry,' is clearly irrelevant. 'Capital allowances on a change of ownership,' however, **does** look like it will be slightly relevant, so I'll click to open that in a new window.*

Indirect document selecting by author

Another way in which lawyers performed indirect selecting behaviour was by using details about the author of a particular legal journal article. One LLM student, for example, was looking for an article written by Higgins and therefore scanned the results list looking for that surname:

> *A13 (LLM student): It's the article by Higgins, so I usually just search by last name when I'm looking for the article.*

Indirect document selecting by headnote or summary (for cases) and abstract or contents page (for articles)

Lawyers also indirectly performed document selecting by reading the headnote or summary (for cases) and abstract or contents page (for articles):

> *A18 (LLM student): Sometimes students and lawyers, if they're pushed for time will just look quickly through the headnote, so it is quite important. But it is also quite difficult to get the decisions out of a document. People have different views about why a judge ruled in this way. So it's a very important task.*

> *A22 (LLM student): Then I'd read the abstract. This abstract talks about the Criminal Justice Act, parts 9 and parts 10, which talk about 'the modification of rule against double jeopardy.' So that's very relevant.*

Similarly, in the excerpt below, a Tax Trainee read the summary of a document on the firm's Knowledge Management database and decided that as it included text 'directly related to the facts' he was looking at, he would read through the full-text of the document:

> *P13 (Tax Trainee): You can see on the screen [points to summary text in results list] that this doesn't actually tell me what the document contains, it just says 'partnership based leasing structure combining instructions to…' [Continues to read summary text]. So that's not very specific, so I needed to click in to see what it's about. [Loads full-page document summary and reads through summary]. I can see here, something directly related to the facts I'm looking at in this little description [highlights part of the un-highlighted text with mouse cursor], which is good. At this stage I would probably click into the document itself and start having a look through.*

Indirect document selecting by keywords

Keywords were also read as a means of indirectly selecting documents. In this example, a DR Professional Development Lawyer noted that a particular document was unlikely to be relevant because it contains the keyword 'mediation':

> *P4 (DR PDL): I then pretty much looked at, let's say 80% of them by looking at the keywords here. I would know that this one here isn't relevant by looking at the keywords [points at keyword 'mediation'] because it's looking at stay of proceedings in order to have a mediation instead and I know that's not going to be discussing massive multi-party court proceedings.*

Keywords representing chapter headings were also read as an indictor of the potential usefulness of particular pieces of commentary within Halsbury's laws:

P18 (DR Associate): I look at this quickly and I'm looking down this column. [Points to 'location' column on results list]. These are chapter headings from Halsbury's and they tell me the are - the chapter of Halsbury's that it's from. So with only 12 search results, I can look through those very quickly and see what's relevant to me. I'm not interested in 'bailment,' I'm not interested in 'copyright,' 'damages' to an extent. But 'misrepresentation and fraud,' that is what I'm interested in, so I start concentrating on what I've got here.

Similarly keywords within the body of a case report were also used by some lawyers to decide whether to read the case in greater detail:

P29 (Tax Trainee): I'll start by just looking at the key words and I can see that it's probably not going to be what I want.

Indirect document selecting by source (e.g. journal title)

The source in which a legal journal article was published was also used as a means of indirectly selecting documents. In this example, a DR Practice Development Assistant described using the journal title as a means of knowing whether the article might be useful or not:

P19 (PDA): The 'source' [points to 'source' field in results list and pauses briefly] in this case because I'm not too au fait with the whole Human Rights Act, I don't know what are the main journals for it and also I don't know what some of these abbreviations stand for. But in other cases you get to know what the main journals and abbreviations are. You know that CJ is Construction Journal or whatever. So in that way you can pick them out by source.

Indirect document selecting by the context that query search terms are mentioned

The penultimate way in which lawyers indirectly selected documents was by looking at the context in which their query search terms were mentioned in the text of the document (and often reading the surrounding paragraphs):

A2 (3rd year LLB student): I tend to just go down the page and read in what context subscription is mentioned. I did this search already actually and there wasn't much. [Continues reading the first few pages of results].

P12 (Tax Associate): So I would basically just start with the most recent result [clicks on document in results list and loads full-page summary] and then kind of find where, quite usefully, it highlights where your search terms are. I'd just have a look at those and see if what I was looking at was relevant to my question.

Similarly, some lawyers read the document result snippets rather than the text of the document itself, again identifying the context in which their query terms had been mentioned. In the following example, a DR Trainee used the 'expanded list' function within LexisNexis Butterworths to display a result snippet of the query terms in context (i.e. the 'chunk of text within which your search term is'):

*P6 (DR Trainee): It's come up with these results in Halsbury's and what I would tend to do is click on a setting here called 'expanded list,' which is very helpful. **R: Why do you use the expanded list?** P6: Because say here I've got 31 hits and I'm trying to work out which ones might be helpful. If you just see the titles to the paragraphs within Halsbury's which it's deemed relevant [shows non-expanded list], you can't always tell from them whether something's going to be relevant or not. Whereas if you use 'expanded list,' it's very handy. It basically gives you the chunk of text within which your search term*

is. So you can just very quickly scan through and see 'no, I don't like the look of that one!' Say this one here, for example, that looks relevant.

Indirect document selecting by how many times search query terms are mentioned

Another way in which lawyers indirectly performed document selecting was by looking at how many times their search query terms were mentioned in the document text, as mentioned by the LLM student and Tax Trainee below:

> *A19 (LLM student): In this case I would focus on how many repetitive terms of 'Wigmore' will appear in the paper. For example, here there is only one mention and I would skip it because it may not be relevant at all.*

> *P29 (Tax Trainee): If I've clicked on one and it says '25 hits' or 40 then obviously it will be a good indication that, at least in very broad terms, that the case is on-point.*

Similarly, the undergraduate student below referred to a feature within LexisNexis Professional similar to the 'expanded list' feature within LexisNexis Butterworths. She suggested that it is possible to use this feature to decide whether a document might be useful by looking at how much of the screen is devoted to the result snippet (i.e. how many times the search query terms are mentioned):

> *A2 (3rd year LLB student): So it lists where those terms, 'access legal advice,' have been mentioned in the case. I think it's good because you can go through it and tell quite quickly if a case is going to be relevant based on how much was said on it. If there's a big chunk of the screen devoted to it - you see there it's been mentioned three times, then I'll probably have a look at this case.*

Indirect document selecting by how close search terms are to one another

Indirect document selecting was also achieved by examining how close the search query terms are to one another within the document. This was a rare form of content extracting, illustrated only by one Tax Trainee, who dismissed a commentary article as unlikely to be useful because the search terms were not close together in the document text:

> *P22 (Tax Trainee): That's a tiny bit useful, but not wildly. **R: And how did you determine that?** P22: Right, because what I was focusing on there was '81(3)' and 'issue' within close proximity to each other and some level of discussion around the definition of 'issue,' but it seems here that there's no real substantive discussion of the terms. **R: So you wanted the two highlighted search terms to be close together in the commentary?** P22: That's right. Yeah.*

Indirect document selecting by examining how relevant the electronic resource deems the document to be

Finally, indirect document selecting was also achieved by examining how relevant the electronic resource itself deems particular documents to be. Lawyers, such as the Tax Trainee below, sometimes examined search results to see how close to the top of the results list they were placed:

> *P13 (Tax Trainee): The best results are when you see a heading quite close to the top of the screen that tells you 'this is exactly what I'm looking for' and you **know** that you want to look at that document rather than trying to figure out whether the document is going to tell you what you want it to.*

Similarly, another Tax Trainee selected a document from the firm's Knowledge Management database because the resource assigned a high relevance ranking percentage to the document based on the search query terms entered:

> *P23 (Tax Trainee): That's potentially relevant [points to top result in the list].* **R: How did you know that the top result was potentially relevant?** *P23: Partly because it gave me a figure of 61 [points to relevance ranking percentage] and that's the sort of percentage relevance that it deems, and it ranks them in order. [Loads summary of document and skim-reads text].*

Summary of selecting behaviour

Selecting behaviour was observed at the resource, source and document levels and was found to be common across all groups of lawyers that took part in our study. Lawyers selected resources based on several *hard criteria* (such as the subject and nature of the content of the resource, the structure of the content, the perceived authority of the content, the perceived comprehensiveness of coverage of the content and the perceived cost of accessing the content). They also selected resources based on several *soft criteria* (such as the perceived ease of use/simplicity of the resource, the perceived speed/time savings offered by the resource, prior positive experiences that the user has had with the resource, the user's familiarity of the resource and whether the resource had been recommended by others or not). Selecting at the source level was not as common as at the other levels, perhaps due to the fact that sometimes decisions about choosing which sources to use to find information in are made at the resource level and due to the fact that sometimes lawyers do not have a choice of which source to choose, particularly when the information required can only be obtained from one particular source. Lawyers selected documents indirectly by level of court, date, title, author and source name. They also selected documents indirectly by reading the headnote or summary of a case or abstract or contents page of a journal article and the keywords or index terms used to describe a document. Finally, they selected document indirectly by examining the context in which query search terms are mentioned within the document, how many times they were mentioned in the document and how relevant the electronic resource deemed the document to be.

5.5.4 Extracting (C)

In section 5.3.6 we discussed document extracting (*"systematically working through sources to identify documents of interest"*). We now discuss extracting at the content level. This behaviour is highly similar to 'content browsing' behaviour. Indeed, browsing content can facilitate the extraction of content deemed useful. Document selecting behaviour can also work hand-in-hand with content extracting behaviour, as lawyers often used the substantive topic of the document to help them decide whether or not the document might be useful to them.

Content extracting behaviour was found to be either direct or indirect. *Direct* content extracting involves systematically working through the actual content of sources or documents whilst *indirect* content extracting involves systematically working through meta-information. Both direct and indirect subtypes of *content extracting* behaviour were equally common and unlike extracting at the document level, content extracting was a fairly widespread behaviour. We now discuss each of these subtypes of content extracting behaviour.

Direct content extracting

Although direct document extracting behaviour was not as commonly observed as indirect document extracting, both direct and indirect subtypes of *content extracting* behaviour were equally common. Unlike extracting at the document level, content extracting was fairly widespread (even though extracting overall was not a particularly widespread behaviour).

Lawyers achieved direct content extracting behaviour in the same way that they achieved document selecting – by reading the textual content of legal documents in order to determine whether the content might be useful for the information-seeking problem at hand. In this example, a DR Trainee had a rather complicated question to answer. She wanted to:

> *P14 (DR Trainee): Find out if the Office of Fair Trading can apply a competition disqualification order, which basically disqualifies a director from being a director for 15 years, whether they can they can apply one of those to the director of a parent company where the subsidiary company has been in breach of competition law and the subsidiary company is based overseas.*

In order to answer this question, she skimmed through the full-text of a commentary piece from the electronic edition of 'Buckley on the Companies Act' (a commentary source within LexisNexis Butterworths) and pinpointed the part of the commentary that began to answer her question:

> *P14 (DR Trainee): This bit's relevant here [begins to read through text of a paragraph that did not include one of the search terms]. It says 'a parent company may be accountable for a breach of one or more of its subsidiaries on the basis that the subsidiary has no real independence from the parent and the parent and subsidiary constitute a single economic entity.' So that's talking about the consequences for the parent if a company is in breach. [Continues to read text].*

Similarly this Tax Associate skimmed through the entire text of an Act in order to see whether it mentioned a legal definition of the word 'participator:'

> **R: As you skimmed through that document for a second time, what bits were you reading?** *P29 (Tax Associate): I actually skimmed through it all. There was one bit that mentioned 'participator' but it wasn't exactly what I was looking for because it told me when being a particular was relevant, but not what goes into the definition of 'participator.'*

This Tax Trainee also skimmed through the text of an entire document, this time a case report. He explained that reading the entire text of cases is encouraged in the Tax department in order to avoid misunderstanding the case and in order to make sound inferences about the main decision of the case:

> *P22 (Tax Trainee): So it's brought this up [points to case text] and I would have read the whole case because when we read cases in the Tax department we are expected to read them properly, rather than do the sort of flick through approach that I was taking to some of the resources before. So word searching or whatever is probably not sufficient.* **R: Is that for any particular reason?** *P22: Well, because if you just pick parts out of the case then you run the risk of misunderstanding what they're saying or failing to fully explain it in context and also because you've got, in this case, a number of judges [pauses], you need to decide what the main decision is because they could have all said different things. You need to decide what they have in common.*

Rather than read the entire text, this Tax Trainee 'skims through the start of the paragraphs,' and noted that the commentary article he was currently viewing does not deal with the precise issue he wants to find information on:

> *P35 (Tax Trainee): I guess what I've done there is that I haven't read the text fully, I've sort of skimmed through and looked at the start of the paragraphs to see what it's about and I can see that this passage of text that I'm looking at here is mainly about when a new Partner comes into the Partnership rather than dealing with the situation that I was looking at. So I would discount that. [Continues to look at results].*

Just as with direct document selecting, not all direct content extracting behaviour was achieved by reading or skimming the entire textual content of a document. Sometimes lawyers searched within the document in order to jump to specific parts of the content:

> *P6 (DR Trainee): If you're searching for a specific part of a case, then you can obviously tap in a keyword and it'll take you immediately to that part of the judgement which may be 60 pages long, whereas if you're looking at the hardcopy of a case, you never have time to read through the whole thing.*

Other times lawyers browsed the document and only read content within headings of interest. For example, the DR Practice Development Assistant below wanted to find out if any Statutory Instruments had affected the Olympics 2012 and Paralympics games Act 2006. The PDA went about finding out the answer to this question by scrolling to the 'notes' section of the Act, which is where he expected to find this information:

> *R: So you mentioned that you wanted to find out if any S.I.s had affected this Act since it was enacted? P17 (DR PDA): Yes. This is probably an unorthodox way of doing it, but this document is only 62 pages so I was able to go through it and just look if anything is in the 'notes' section, because that's where it would be.*

Similarly this Legal Practice Course student scrolled through the text of a case report and only read the 'judgement' of the case in order to avoid reading the full-text of the case in its entirety:

> *A28 (LPC student): It's brought up all the facts of the case and I think with a bit of luck it's brought up what the judge's said. [Scrolls down document and finds 'Judgment' heading]. Yeah, it's brought up the opinions of the judges.*

Finally this Tax Trainee scrolled through the full-text of a journal article and read the section entitled 'tax avoidance' because the heading seemed relevant:

> *P7 (Tax Trainee): This section with 'tax avoidance' in the title deals with the tax avoidance issue, which is an issue that I'm interested in. I just tend to see if there's things that look quite useful and then if there is, I will just print the document out and then read it. [Prints document].*

Indirect content extracting

Indirect content extracting was achieved in some similar ways to those used to achieve indirect document selecting (indeed many of the ways in which we observed lawyers to achieve content extracting have already been discussed in the context of document selecting). The ways in which they achieved indirect content extracting behaviour were by examining:

- The headnote or summary of a case, summary or contents page of a piece of legislation and abstract or contents page of a journal article.
- The context in which query search terms are mentioned within the document.
- How many times search query terms are mentioned within the document.

In our first example, a 3rd year LLB student read the headnote of a case in order to 'draw the points out of it' quickly:

> *A2 (3rd year LLB student): If you're pushed for time, like property law for example, you've got so many cases to read, you just want to draw the point out of it quickly and sometimes it's not necessary that you need to know all the legal reasoning in a case, so if you're just looking for the main principle then you just go for the headnote.*

Similarly a DR Paralegal examined the contents page of a PDF document containing a Statutory Instrument:

> *P1 (DR Paralegal): Sometimes you might get a document that is like 300 pages long though, so you'd have to have a look at the contents page and see which section you want to look at.*

Three examples of examining the context in which the search terms have been mentioned in the text of the document are presented by some Trainees. The first trainee read the paragraphs of the content of a legal document that surround her search terms in order to decide whether the document was likely to be relevant. This behaviour is identical to document selecting behaviour and illustrates that document selecting can often be facilitated by content extracting (i.e. the two behaviours can work hand-in-hand):

> *P14 (DR Trainee): The search terms come up in red, which is quite good for skimming through the text to see where the terms you want are located. To be honest, I would have to read this in quite a lot of depth in order to decide whether it was relevant or not. [Continues to skim read through full-text of document, scrolling between the search terms highlighted in red within the text]. And this is a pretty long case, unfortunately. [Reads parts of surrounding paragraphs where search terms are highlighted aloud]. Ok, this isn't relevant. I think it just mentions the terms that I searched for in it, but that's not what the case decided. I need a case that says 'this is the law on this area.'*

The second Trainee skim-read the full-text of a case report and noted that, even though the case mentioned each of her keywords ('tax,' 'abuse' and 'European Law'), it was still not useful because it 'deals with a different point, in a different area':

> *P7 (Tax Trainee): Having read one of the sections in this one, I don't think it's that useful, even though it has everything I wanted in it, which was 'tax,' 'abuse,' 'European Law,' but it deals with a different point, in a different area, so I don't think it's that helpful [continues to look through results and reaches the end of the list].*

Finally the third Trainee searched a case report for the mention of the legal term 'reciprocity' (illustrating that content searching behaviour can sometimes facilitate content extracting). She then mentioned that she would read the surrounding text in detail:

> *P8 (DR Trainee): Let's say you have a feeling that Lord Browne Wilkinson has said something very specific in this case, then you might even do a CTRL and Find [presses CTRL-F shortcut to find words within the text] and type in something specific like 'reciprocity' [finds highlighted term] and then I'd read that in quite a bit of detail.*

Indirect content extracting was also facilitated by looking at how many times search query terms are mentioned. In this example, similar to the previous examples of using the LexisNexis Butterworths 'expanded list' feature to facilitate document selecting, an LLM student examined how many times his search results were found in the text of the document and, after deeming the journal article to be relevant to his topic of 'predatory pricing,' began to skim the full-text of the article:

> *A17 (LLM student): You can have an expanded list as well which will actually tell you why the computer has given this result back. This is in the substance of the article itself. You can find articles which are very much concerned with 'Predatory Pricing' if you see that they have lots and lots of quotes that mention 'Predatory Pricing.' And you get articles that quote all the way down the page [pauses] and you know that article is very much concerned with Predatory Pricing.*

Summary of content extracting behaviour

Content extracting behaviour was found to be common amongst the lawyers who took part in our study and was often used to facilitate document selecting behaviour. We found that content extracting may be *direct* or *indirect*. Direct content extracting involves systematically working through the actual content of resources, sources, documents or content whilst indirect extracting involves systematically working through meta-information instead. We found that direct content extracting was usually performed by reading or skimming the textual content of legal documents, or parts of the content. Indirect content extracting was performed in three main ways (which were also ways of performing document selecting): by examining the headnote or summary of a case, summary or contents page of a piece of legislation and abstract or contents page of a journal article, by examining the context in which query search terms are mentioned within the document and by examining how many times search query terms are mentioned within the document.

5.5.5 Recording (R, S, D&C, Q)

Recording involves making a record of resources or sources used (resource and source recording respectively), of documents and content found (document/content recording) or of the query terms used or results returned in a search (search query/results recording) –the 'Q' in the title above. Recording behaviour can be *manual* (i.e. by hand) or *automatic* (with the help of technology – such as a 'search trail' which automatically keeps a record of search queries entered and results received).

Resource recording (keeping a record of resources found)

Evidence for recording at the resource level (keeping a record of resources found) was observed. However, recording at this level was not particularly common. This may be due to the fact that these sorts of behaviours are less likely to be uncovered by asking lawyers to find information that they currently need or have recently needed for their work (the broad task that the observation portion of our study was based around). However, when resource recording behaviour was mentioned or displayed, it almost always involved saving and re-visiting Internet bookmarks, as explained by this law lecturer, which is a form of manual resource recording:

> *A7 (Lecturer): My starting point would almost always be the Internet and it would almost always be one of the bookmarked pages. So that's where I would start and use that as a jumping off point.*

As explained by the same law lecturer, this method of recording resources by using Internet bookmarks does not, however, help lawyers to remember *why* they originally bookmarked a particular resource or what types of legal information it is useful for finding:

> *[Loads LexisNexis Academic QuickSearch and finds a recent case given on a student handout]. **R: What is this 'Academic QuickSearch' version of Lexis?** A7 (Lecturer): I don't know actually, I must have had a reason to bookmark this version rather than another version when I was doing my doctoral research. So I just tend to use this version now, because I would have decided that it was the best version for my research. But because I never record the reason why I bookmark sites, I'm not sure why I originally bookmarked it.*

Some practicing lawyers also mentioned use of pages on the firm's intranet, which also provided a list of links to electronic legal resources (although, of course, this is a pre-defined and shared record of resources rather than a user-created record of useful resources):

R: In terms of accessing the various electronic resources that you use, do you have them saved in your favourites? Or if not, how do you get to them? P14 (DR Trainee): I just go through the Intranet because there's a 'research' tab tool that's just got it all there.

Source recording (keeping a record of sources found)

Similarly to recording at the resource level, we found evidence for source recording (keeping a record of sources found). However, recording behaviour at this level was also not very common. A couple of lawyers did, however, mention the fact that LexisNexis Butterworths allows lawyers to keep a partly manual and partly automatic record of the individual sources that they search within the library. As explained by this DR Associate, LexisNexis Butterworths provides a pull-down list of sources which automatically displays sources within the library that the lawyer has recently searched within at the top of the list:

> *P18 (DR Associate): The way that they've done this new pull-down list [clicks on source combo box], they've got a list of hundreds of sources and I have yet to find a quick way of [pauses, then scrolls through source list]. I think actually on my PC, because I've used it before, Halsbury's would now show up at the top of the list, but in the meantime that's pretty useless to me.*

This Bar Vocational Course student, referring to the same recently used sources list, explained that it is also possible to manually customise the list by adding frequently used sources within LexisNexis Butterworths to the list:

> *A33 (BVC student): Because I've used LexisNexis Butterworths before, the last thing that I've used, Harvey's Employment Law, is already there in my sources [points to the list of recently used sources]. It wasn't under the original list, so you have to add it to your list.*

Document and content recording (keeping a record of documents/content found)

Document and content recording was far more commonly mentioned and observed than recording at the resource and source levels and often involved simply printing or saving documents and occasionally e-mailing the document to oneself. As explained by one DR Trainee (and echoed by several academic and practicing lawyers), lawyers often expressed a preference to print long documents rather than read them on-screen:

> *P6 (DR Trainee): I tend to not like looking at things on the screen much, I tend to miss things. So I'd think 'that looks like a useful resource' and I'd print it out into a hardcopy and then I'd have another look at it, go through and highlight anything that I thought was relevant.*

As with the DR Trainee above, lawyers across all groups mentioned either 'highlighting' or 'tagging' content within documents that they had printed by hand. This was also discussed by the Tax Associate below:

> *P20 (Tax Associate): It's not worth printing out the odd pages here and there. It's much easier to print the whole section on Control of Foreign Companies and then if you need to tag or highlight particular pages within that printout, then you can. [Retrieves printed document and loads the document again from the search list, then searches once again for 34(4), notes the page number and flicks to the corresponding page in the printed version of the document and adds a paper sticky tag to the page]. That will save me a couple of minutes or so of drafting by letting me cut and paste from the underling document into the memo of advice that I draft.*

Despite the use of highlighting and tagging on printed documents, almost no evidence was obtained of lawyers making a note of *when* a particular document was obtained, *where* it was obtained from and what search terms or browsing categories were used to locate it.

As this DR Associate explained, the output of many research notes produced by Trainees in the firm also includes hardcopies of materials in an attached binder:

P18 (DR Associate): People here are generally quite keen on producing memorandums or notes of their research, so they'd normally come back with a full note dealing with the question that's been put to them and setting out the sources in footnotes and probably attaching a binder of materials printed off or photocopied in hardcopy.

These files or binders are often stored by practicing lawyers in a filing cabinet or on shelves so that the material can be easily returned to if the need arises, as explained by the Tax Associate below:

P20 (Tax Associate): Usually I just keep these printed materials, especially this sort of textbook secondary material, in a file so that they can be used on a communal basis and they're kept in personal files such as these here [points to binders on shelves]. It's good to have it all in one place though as sometimes another point comes up or you'll be reading the section you're interested in and it'll say 'see section D102a' and you've got that in this printout, instead of having to go back and find it again.

Saving legal documents onto a hard drive or removable disk was another common means of keeping a record of documents found. As this law lecturer explained, saving documents can also be used in conjunction with printing documents in order to have 'two sources' stored. In this case, the law lecturer saves documents under various folders and sub-folders, arranged by legal subject:

A7 (Lecturer): Over the years since I did my doctorate, I used to start off by photocopying and printing out documents that I needed so these [documents filed on the wall] are all the documents that I needed for my doctorate. Increasingly I got to the point that it was just too much information. My primary sources would be quite detailed documents which were published by the European Commission which would be anything from 20 to 200 pages long, so I stopped printing them out and I would download them onto my laptop, so I would have two sources of information. Documents that I had saved onto my laptop under files of primary sources or articles and I would also keep links to the web pages filed under various folders and subfolders.

The same law lecturer explained, however, that maintaining such a hierarchy of folders in order to save documents is difficult and it is often necessary to re-arrange them:

A7 (Lecturer): I'm constantly having to re-arrange my folders because I find that [pauses] I have to separate them out not just chronologically [pauses] but when you start out your research, everything is just one amorphous mass and you realise that you have to distinguish between different categories. So, for example, in terms of Employment Law I might originally have a folder which is 'labour law articles' but then soon I realise that I have to separate them out into articles on equal pay, articles on unfair dismissal, articles on contracts and employment. And sometimes even in that I might have 'equality' and I would have to separate them out between gender and race equality, sexual orientation, religion, age and theories and concepts [pauses] because I find that if there are more than about 20 or 30 articles in one folder, I totally lose track so I just have to think logically [pauses] what's the best way of dividing up. So I divide race and gender into sexual orientation, religion and age and that's the logical division for me and I'll know that's what I've done. So the filing system, as it were, just increasingly grows and just develops as you realise that there are subtle nuances and differentiations that you can make.

Finally, e-mailing legal documents to oneself was also found to be a means of document recording. This LLM student e-mailed documents to herself in order to avoid finding the same document again sometime in the future:

A9 (LLM student): I'm glad I know in Westlaw how to [pauses] when you've got a case, you can either quickly print it, but I'm glad you've got the option to e-mail because if you don't have time and you can

e-mail yourself a document it's a lot easier than sort of going back to it, going through the whole hassle of finding it all over again.

Search query/result recording (keeping a record of search queries used in or results returned from a search)

Like resource and source recording, recording at the search query/result level was also not particularly common, although it was mentioned and observed to a limited extent across all groups of lawyers in our study. The value of keeping a record of search queries or results was not appreciated by all. This DR Trainee emphasised the importance of keeping a research trail:

> *P6 (DR Trainee): You have to be very careful about having a research trail so you can see exactly how you got the information that you've got. If you're online, it's easy to just click on a link and jump about and suddenly you've reached an answer or found something out that is useful, but you're not entirely sure how you got there.*

However this trainee in the same department of Dispute Resolution asserted that there was 'not enough time' to keep a record of the entire research process and therefore does not keep record of her search queries or results:

> *R: Is there ever a need to keep a log of what you're searching, or not really? P14 (DR Trainee): It's something they told us we should do at law school, but I've never done it here just because at law school we had to show exactly how we've found the information, whereas here they're not so interested in that, they just want the information. [Laughs]. You haven't got time to go through your whole researching process.*

Whilst the value of keeping a record of search queries and results may not always be recognised by Trainees, this DR Practice Development Assistant (who undertakes information-seeking work on behalf of other people within the firm) explained that there is often a need to keep a research trail in order to send this information to the enquirer. As he explained, the information recorded includes the search query terms used and the order in which searches were submitted:

> *R: When you send material off to someone that's made an enquiry, is there ever a need to send an audit or research trail of what searches you've typed in? P19 (DR PDA): Yeah. Quite often they're interested in the search terms you used, what order you did the searches in, how you narrowed those down. I mean sometimes people just want one thing, but quite often if you don't tell them they come back and ask for you to tell them what you've searched for.*

When probed for more details about keeping a record of search queries and results, the Practice Development Assistant explained that he keeps this record manually and sometimes types the details into an e-mail and other times takes screenshots:

> *R: If you were sending them a document that structures all of that, how would you structure that document? P19 (DR PDA): It's usually just in the e-mail attaching it. I'll usually just start 'blah blah blah, here are the ones you asked for' and then just go through the databases and the order, then the search terms I've used and then put 'see attached the results at the bottom.' R: So you type them out and then give them screenshots of the results? P19: Yeah, but it depends on exactly what they want and where they are. Sometimes it's screenshots, sometimes I'll just save the end result and save it as a PDF. Quite often they want to know the number of hits that are going to come up too. Is this going to create 300 or 3,000 or 3 hits? Just so that they know how to go about narrowing it down or whether they need to re-think the search.*

Overall, it was more common for lawyers to keep manual records of the search queries used (and sometimes brief information about the number of results returned) than to use automatic means of search result/query recording. Indeed very little mention was made of other tools within electronic legal resources that might support search query/result recording, such as features that keep a record of the user's 'research trail,' saving a time-stamped list of the searches that they have performed and allowing users to re-run these searches. An exception was by this Tax Trainee, who explained that he found LexisNexis Butterworth's research trail facility, which automatically keeps a record of searches carried out, to be useful:

> *P23 (Tax Trainee): Lexis records all of your recent searches as you may or may not be aware. The saved searches are quite good just so you don't end up repeating the same search terms over and over again or to make sure you use enough combinations of different search terms, as you don't want to waste time doing the same thing again.*

Summary of recording behaviour

Evidence of recording behaviour was found amongst the lawyers who took part in our study, although more evidence was found of recording at the document level (i.e. keeping a record of documents found) than at the resource, source or search query/result levels. We found that recording behaviour could be achieved *manually* (i.e. by hand) or *automatically* (with the help of technology – such as a 'search trail' which automatically keeps a record of search queries entered and results received). Although recording behaviour is conceptually similar to Meho and Tibbo's 'information managing' behaviour, it has a more precise definition and a narrower scope.

5.5.6 Updating (D&C)

Document and content updating (ensuring a current understanding of amendments or changes to legal documents and content and an understanding of whether a particular case or piece of legislation is good law) was identified to be another important behaviour for academic and practicing lawyers alike and was displayed fairly widely. This behaviour, to the best of our knowledge, has not been previously identified in studies of information-seeking and was observed to operate at only the combined document and content level.

Updating behaviour is related to 'monitoring' behaviour in the sense that it is possible to ensure a current understanding of legal documents and content through keeping abreast of developments in a particular legal area. Therefore sometimes updating behaviour can be facilitated through monitoring. However, updating is an important and different behaviour in its own right as it does not necessarily involve regular searching or browsing within electronic resources (as with active monitoring), nor does it involve passively receiving updates (as with passive monitoring). Instead, updating can occur on a document-by-document basis. For example, if a lawyer wants to know whether a case is still good law, he or she will often consult a case citator (perhaps forwards chaining in order to facilitate this task). Similarly if a lawyer wants to know whether a piece of legislation is still in force, he or she might consult the electronic version of Halsbury's 'is it in force' to find out.

In addition, updating behaviour is also related to 'forwards chaining behaviour' in the sense that it is possible to ensure a current understanding of legal documents and content through identifying and accessing documents which have subsequently cited the current document. Indeed, the previous example shows that updating behaviour can be facilitated through forwards chaining. However, unlike forwards chaining, updating does not necessarily involve *accessing* documents which have subsequently cited the current document, but merely *identifying the existence* of these documents.

Similarly, updating behaviour is different to Ellis's behaviour of 'verifying,' *"checking the information and sources found for accuracy and errors"* - (Ellis et al., 1993, p.364). This is because, through the behaviour of updating, a lawyer is not checking that a particular legal document or the content within it is accurate or error-free, but whether it is a currently accepted interpretation of the law (i.e. whether it is still good law). Through

doing so, the lawyer is also updating his or her *understanding* of the current state of play in a particular legal area. Therefore updating behaviour is just as much about updating personal *understanding* of the current state of a particular legal issue or area as about finding out whether a particular case or piece of legislation is up-to-date.

As explained by one DR Associate, knowing the 'most up-to-date position on everything' is crucial for lawyers:

> *P5 (DR Associate): That's what's really invaluable as a lawyer. You want to know **the** most up-to-date position on everything. It's so crucial. You don't want to be giving out wrong advice on something.*

This sentiment is echoed by a law librarian, who describes knowing if a legal case is still valid (along with 'knowing the judicial history' of cases – which we discuss in section 5.5.7 on 'history tracking') to be 'vital' for lawyers:

> *A20 (Librarian): Knowing the judicial history, if you like, of a particular case and its relationship to other cases is a vital part of what they're trying to do and they do need to know if a case is still valid - if a decision hasn't been overturned by a later case that has been decided in a higher court, so it is quite an important thing.*

We found that updating behaviour could be performed *directly* by manually checking that the document or content within is up-to-date or good law (for example by conducting manual searches within a electronic legal resource). Updating behaviour can also be performed *indirectly* by using electronic tools to check whether a particular document or the content within it is up-to-date or good law.

Direct document and content updating

Direct document and content updating was fairly widespread across all groups of lawyers and was achieved in three main ways:

1. By conducting manual searches within a electronic legal resource.
2. By looking for mentions of amendments within the document text.
3. By looking for symbols within the document (usually a case) to indicate its positive, negative or neutral treatment.

The first way in which direct document and content updating was achieved was by conducting manual searches within a electronic legal resource, often by typing one of the party names involved in a case or the legislation name into the general search field:

> *[Participant points to current case that was deemed relevant]. P16 (DR Associate): That was in 2002, so I need to know if it's been overruled and there doesn't seem to be anything about it here [points from top to bottom of text on summary page]. So what I would then do is return to my results, modify search and actually type in that case name and it would show me where it has been mentioned in other cases.*

As explained by this BVC student, it is also possible to search for recent cases in order to ensure that the particular case that a lawyer has found it still likely to be good law. The student also explained that 'there's no point citing a case' that is not current law:

> *A33 (BVC student): It's useful to know that what you're reading is likely to still be good law. So maybe you could go on and just search to see if there's been any cases advertised in the last month that may have caused this to change. But it's quite important that you make sure that whatever it is that you're using is up-to-date because you don't want to be relying on things that have been overruled later or things that have changed the law. There's no point in citing a case that is a year old if it's not current*

law. You need to make sure that the law is still the case today for any argument that you use the law for. You'd look a bit of a fool when you go to court if you didn't!

Direct document and content updating was also achieved by looking for mentions of amendments within the document text. As explained by this DR Professional Development Assistant, LexisNexis Butterworths incorporates amendments to legislation within the document, therefore it is possible to know whether a particular piece of legislation is up-to-date simply by looking for (and reading) any amendments:

P17 (DR PDA): If you want a piece of legislation that is up-to-date, so let's say you have an Act that's from 1980 and you want to make sure that it's as up-to-date as possible, then LexisNexis is the place that you would go because they incorporate all the amendments to that Act into the text of the document. And the good thing about it, as you've probably seen here, it will have notes added to it. [Scrolls through PDF document full-text and points at an example note].

A BVC student also mentioned similar updates present in the online version of Halsbury's Laws of England and Wales:

A33 (BVC student): One of the good things that you get is the footnotes on Halsbury's Law. LexisNexis do an update section when they list at the bottom of the commentary all the updates. I don't know if I can remember how, but you used to be able to get the date that it was correct from [pauses] maybe if I click on 'source information' [does so and scrolls down a popup with some lengthy text]. So it was updated on 1st June 2006.

It should be noted, however, that whilst looking at amendments in legislation can help lawyers gain an understanding of whether the version of the legislation that they are looking at is *up-to-date* (i.e. current), the amendments do not necessarily include details about whether the piece of legislation is *still in force* and therefore *still good law.*

A final way in which direct document and content updating was achieved was by looking for symbols within the document (usually a case) to indicate its positive, negative or neutral treatment. As explained by this DR Trainee, LexisNexis Butterworths displays symbols to indicate the treatment of certain cases and provides a help page to explain what each of the symbols mean:

P6 (DR Trainee): I wanted to make sure not just that it hadn't been overturned, but if it had been distinguished in any way, or if there were cases that had followed it - cases with a slightly different set of facts that might have been more relevant to the question I was doing. You can do this in Butterworths. It has little symbols next to the case name which tell you [pauses] I think if there's a blue squiggly line then I think it tells you one thing, it tells you that it hasn't been overturned. If it's got a big red cross next to it, then it means that it has been overturned. There's an online database where you can look up exactly what it means.

Symbols to denote case treatment are also available in other electronic legal resources such as Westlaw.

Indirect document and content updating

We found far less evidence of indirect document and content updating as compared with direct updating. Few practicing lawyers mentioned or demonstrated the use of citator tools and even fewer academic lawyers mentioned or demonstrated the use of these tools. However, when indirect document/content updating behaviour was mentioned or displayed, it always involved using case citator tools within electronic legal resources designed for tracing the judicial history of legal cases.

For updating legislation, lawyers used the electronic version of Halsbury's 'is it in force' (which provides the date a particular provision of an Act came into force):

> *P3 (DR Paralegal): There is a database where you type in the name of a law and there's a button which says 'is it in force?' and it gives you updates in terms of 'yes, it is in force, however it's been amended later on' or 'no, it's no longer in force because x, y and z.' It's quite useful.*

A couple of practicing lawyers also mentioned and demonstrated use of the statute citator tool within the Current Legal Information electronic legal resource, which states whether a particular provision of an Act has been amended or repealed:

> *P17 (DR PDA): The 'housing grants, construction and regeneration' Act [conducts search for Act in CLI statute citator and clicks on Act title hyperlink]. You see CLI is quite useful because it lists some of the cases where that Act has been important and there's a bit more history behind this case because it's a bit older [points to some Statutory Instruments listed as having affected the Act].*

For updating cases, lawyers used citator tools within electronic legal resources (such as the LexisNexis Shepardize tool for US cases and the 'CaseSearch' electronic resource for UK cases). The Shepardize tool was used by a PhD student who had a particular interest in locating US cases and, as he explained, shows the lawyer whether a particular case has been subsequently supported, overruled, repealed etc.:

> *A4 (PhD student): You probably don't know whether a case has been supported or overruled by later cases. 'Shepardize' compiles information about these things so you know whether this law is good authority or not. But even in Lexis UK, we can't use 'Shepardize' because they only cover US cases.*

Whilst the above use of the Shepardize tool was the only mention or use of citator tools amongst academic lawyers, there was greater mention and use amongst academic lawyers, many of whom mentioned a UK equivalent of the Shepardize tool called 'Case Track' (which is a citator for UK cases):

> *P6 (DR Trainee): Another thing that's useful to do with a specific case is put it in the search engine that I was talking about where you can check whether it has been overturned, how it has been treated in any subsequent cases, just to check that it is up-to-date and it is good law. You can also read through and see if it refers to any other cases that you think might be handy, again with the specific question you're trying to answer in mind.*

As this DR Associate explained, using a citator tool can provide more details about the judicial history of a case than a manual electronic legal resource search for one of the party names involved in the case might provide:

> *R: Is that more useful than just searching for the case name and just seeing where it's been mentioned? P18 (DR Associate): It's more targeted because you can look at particular points that have been discussed in a case and track them through or if you were looking at legislation, you could look at where a particular section had been discussed as opposed to the Act as a whole.*

An example of the use of a citator tool is illustrated by the Tax Trainee in the excerpt below, who used the CaseSearch citator within LexisNexis Butterworths to determine whether a particular legal case was still good law:

> *P24 (Tax Trainee): If we look at what we get through CaseSearch [clicks on version of the case with 'CaseSearch' listed as the source in the results list]. CaseSearch tracks cases that have gone before that a case refers to. So if you're investigating a line of cases, you can immediately see [pauses]. For example this case here is an older case, 1995. So if you're looking at how this line goes, this tells you that it's been 'applied' by the case that we're looking at now, so we know this 1995 case is still good law. If it's not, it'll come up in red and say something like 'overruled' or something like that and that*

*will tell you that all of your research on this case is only good up until that year when it was overruled or up to the case that we are looking at now and that we now have to use **this** law. So you're always using the up-to-date stuff and that's what CaseSearch is good for.*

Summary of updating behaviour

Updating behaviour was found to be an important and fairly widely displayed behaviour amongst the lawyers in our study. To the best of our knowledge, this behaviour has not been identified in any previous information-seeking studies and we argue that this may be due to the fact that it is a behaviour pertinent to legal information-seeking in particular. We found that updating behaviour operates at the document and content level, but is just as much about updating personal *understanding* of the current state of a particular legal issue or area as about finding out whether a particular case or piece of legislation is up-to-date. We found that updating behaviour could be performed *directly* by manually checking that the document or content within is up-to-date or good law or *indirectly* by using electronic tools. We found that direct document and content updating was achieved in three main ways: by conducting manual searches within a electronic legal resource, by looking for mentions of amendments within the document text and by looking for symbols within the document (usually a case) to indicate its positive, negative or neutral treatment. Indirect document and content updating was achieved through the use of electronic tools within electronic legal resources to check whether particular legal cases or pieces of legislation are up-to-date or good law.

5.5.7   History tracking (D&C)

Another low-level behaviour subsumed under the broader category of 'selecting and processing' is history tracking which, like updating behaviour, was found to operate only at the document and content level. Document and content history tracking is different from 'updating' behaviour because it involves ensuring a *historical* as opposed to a *current* understanding of legal amendments or changes to documents and content (i.e. an understanding of the treatment a particular case or piece of legislation has received over time). This means that the emphasis of history tracking is not to ascertain the current state of play of a particular legal issue or area, but to determine how the law *has changed* over a particular period of time. Like with updating, this process is just as much about gaining a historical *understanding* of how the law has changed over time as about finding out how a particular piece of legislation has changed. Document and content history tracking was observed amongst the lawyers that took part in our study, although was not a particularly common behaviour. It also almost always involved tracking the history of *legislation* as opposed to cases, perhaps due to the fact that it may be more useful to have a *current* understanding of how a case has been treated in the past than a historical one (since this knowledge is likely to be used to build arguments for cases based on current law). In addition, history tracking was only mentioned or displayed by practicing lawyers, probably due to the fact that our academic lawyers had little need to trace the developments related to a particular case or piece of legislation over time.

Like updating behaviour, history tracking has to the best of our knowledge not been identified in previous information-seeking studies and, again like updating, we believe that this is because it is a behaviour particularly pertinent to legal information-seeking and may not be an important behaviour for other domains. Although, as with updating, history tracking behaviour can theoretically be performed either directly or indirectly, we only observed direct history tracking behaviour amongst our academic and practicing lawyers, probably due to the fact that very little mention or use was made of the citator tools within electronic legal resources that can support this behaviour.

This DR Trainee explained the need to sometimes look at the wording of a piece of legislation that is not currently in force in order to acquire an appreciation of how the legislation has changed over time. In this case, the Trainee wanted to find out whether a bond would be considered to be a 'speciality' (a special type of contract or agreement). Only by tracking the history of the Limitation Act 1980 (which covers the definition of 'bonds')

was he able to find out that, in an old version of the Act, a bond had been defined as a type of speciality. The Trainee then went on to explain the essence of history tracking behaviour:

> *P6 (DR Trainee): The other day I was doing a piece of research on bonds and whether the Limitation Act applied to them. And the question we had to answer was whether a bond would be considered something called a 'specialty' for the purposes of the Limitation Act. We were looking through to find case law and discovered that in the Act which preceded the Limitation Act, the wording of the Act had said 'any bond or speciality,' so the predecessor of the current Act had obviously envisaged that a bond was a specialty. I actually found that out by reading a case which was referring to the old Act and I wanted to check the exact wording of the old Act itself because that was obviously helpful in our understanding of defining a 'specialty' and whether a bond would be included.* **R: So you're tracking changes to help your understanding?** *P6: Yeah, it's tracking changes and if you're looking at what the intention was behind legislations, sometimes it can be useful to look at any previous Acts that came before. Obviously it's not conclusive, because it's the current piece of legislation which is in force, but yeah it's tracking changes I guess. You've probably put your finger on it!*

Another DR Trainee explained that document/content history tracking is a means of examining the 'direction' or trend in which events or laws have progressed in and therefore might continue to progress in:

> **R: So does the tracking changes serve as a way of putting events in context almost?** *P8 (DR Trainee): Yeah, essentially what the partner wanted to do is stand up there and say 'in relation to Cartel behaviour, in the year 2000 the EU Commission issued 80 million Euros worth of fines and in 2006 they have issued 400 million Euros worth of fines, so you can see that this is increasingly important' and that kind of thing. So we needed the figures, we needed the years and we needed to know why things had changed. We needed to know whether that actually was the change because a lot of this was coming from the Partner saying 'it's my feeling that it's going in this direction' and you're like 'right, I need the evidence to backup what his view is' or to say 'well that might be your view, but the facts point us in a different direction.'*

Similarly this Tax Trainee was set a task by one of the Partners of the firm to 'find out what the meanings are of 'avoidance' of 'abuse,' how the courts have viewed it in the past.' The Trainee explained that historical definitions of the terms were needed in order to find out the most current definition of the terms *and* to use the historical definitions to help predict future trends in changes of the legal definitions of the terms:

> **R: Why exactly did you need to get a historical idea of the definitions of these terms?** *P7 (Tax Trainee): Umm. Basically they've sort of moved around because in the 80s or 90s they were rejecting the principle that you could say that tax avoidance was allowed to diverge from the treaties, whereas increasingly they're going into a lot more detail. The Fairs case, for example, is just one line saying 'this can't be done,' whereas now they're getting to about two pages where they're saying 'well it can be done, but certain things have to happen.' .* **R: So is the idea to find out what the situation is now, or to get a historical view to be able to predict trends in the future, or both?** *P7: Probably both. It's for a seminar which a group of people are going to be giving, so he wants to understand how we've got to the position where we are now and basically predict how it's going to move in the future.*

An example of history tracking behaviour is presented by this DR Paralegal, who is examining UK legislation surrounding packaging requirements. She loaded the full-text of the 'The Packaging (Essential Requirements) (Amendment) Regulations 2006' and, as she was interested in how the amendments came about, mentioned the need to refer to the previous versions of the regulations from 2003 and 2004 rather than simply the most recent 2006 version:

> *P3 (DR Paralegal): From my research previously I remember that the 'Packaging Essential Requirements' are dated 2003 and 2004. Subsequently they have been amended in 2006 and that's why*

I'm going to the 'amendment' one and going to full-text to see what they actually contain. What I would really have to do is then go and print the 2003 and 2004 version and see what exactly the amendments relate to.

Aside from accessing several historical versions of a particular piece of legislation, some practicing lawyers also demonstrated history tracking behaviour by accessing consolidated versions of legislation (which include updates and amendments to the legislation in the full-text). This allows lawyers to track the history of a particular piece of legislation without having to find previous versions of that legislation:

P6 (DR Trainee): I think that on CLI [Current Legal Information] you can also get legislation in it's original, un-amended form, whereas on Butterworths you can get legislation, but you get the consolidated version which, obviously, sometimes you want. By 'consolidated version,' I mean that any updates or any new pieces of legislation that have been inserted get put in and you can't see what the legislation originally looked like. Whereas, I think, on CLI you can. I have had to do that a few times.

Summary of history tracking behaviour

History tracking behaviour was displayed, albeit not particularly commonly, amongst the practicing lawyers in our study (and was not displayed at all amongst academic lawyers). As with updating behaviour history tracking has, to the best of our knowledge, not been identified in any previous information-seeking studies. Again, we argue that this may be due to the fact that it is a behaviour pertinent to legal information-seeking in particular. We found that history tracking behaviour operates at the document and content level and, like updating, is just as much about gaining a historical *understanding* of how the law has changed over time as about finding out how a particular piece of legislation has changed. We found that history tracking was only achieved directly, probably due to the fact that the lawyers in our study demonstrated little use of the citator tools within electronic legal resources that can be used to support this behaviour. Direct history tracking behaviour was, therefore, achieved by locating several historical versions of a particular piece of legislation, or by locating a single consolidated version which included updates and amendments to the legislation within the document text.

5.5.8   Analysing (C) and synthesising (C)

The information behaviours of 'content analysing' and 'content synthesising' were also identified by Meho and Tibbo (2003) in their study of the information behaviour of social scientists. The authors, however, do not elaborate much on these behaviours. As with the other behaviours that have not been formally defined in previous studies, we adopt definitions based on the Oxford English Dictionary. Therefore we regard analysing to involve *"examining in detail the elements or structure of the content found during information-seeking"* and synthesising to involve *"combining the elements of the content found during information-seeking into a coherent whole."* We found these behaviours to operate at the content level (which is somewhat unsurprising, as it is not possible to analyse or synthesise anything other than the content within a document). Content analysing and synthesising were not commonly observed in our study, probably due to the fact that the broad task in the observation portion of our study directed lawyers to 'find' information that they require, but not necessarily to process (or even 'manage') it in any way.

Both content analysing and content synthesising involve, in some regards, cognitive processes that cannot be directly supported by electronic resources. For example, the Tax Trainee in the excerpt below is searching the firm's Knowledge Management database in order to find out the answer to a specific legal issue. This issue involved a Partnership that had been set up between two individuals, who each held a 50% stake in the business. The Partnership had acquired some assets on which Capital Allowances were available and now the Partners were interested in knowing how these Allowances would be apportioned between them if they changed their stakes in the business so that one Partner now held 1% and another 99% of the business. The Trainee, whilst

reading some text from one of the documents, performed the mental act of analysing the implications of the document for the issue at hand:

> P13 (Tax Trainee): So I can see here that UK bank is the Partner that's going to change from a 1% interest to a 99% interest and the document is saying that its entitlement to Capital Allowances 'should be established by reference to 99% of the qualifying expenditure incurred by the partnership on plant and machinery' [reads document text verbatim]. So that suggests that what this document is going to tell me is that it **is** going to be possible to vary the Partnership interests.

However, there are often some physical actions associated with analysing and synthesising the content within documents. Content analysing often involves the lawyer writing (usually hand-written) lists of questions to be answered, issues to look out for or points to prove through reading particular content (as described by a PhD student and DR Associate below):

> A12 (PhD student): My supervisor was drafting a report for the House of Lords Constitution Committee, so what I did is I went through all their reports, read them all and summarised [pauses] looked for the questions they asked and the idea was to make a checklist on the basis of the reports they had written.

> P18 (DR Associate): So that enables me on my notepad to write down bullet points of what I need to prove [reads out things to prove from Halsbury's]. I'll put all of these things down and that straight away breaks down everything that I need to know.

Similar hand-written note-taking was performed by the Tax Associate below, who was conducting research for a talk on the impact of new tax rules. The notes made by the Associate, related to 'what the new rules say,' were aimed at supporting his research task of finding information about the *impact* of the new rules:

> P21 (Tax Associate): I'm just making handwritten notes about what the new rules say. I'm going to refer to the legislation that's already there and see exactly what that says and make a note of that. [Looks through paper version of the Finance Act 1996].

As the Associate explained, these notes were useful for expressing his line of argument and serve as a reminder if he should need to 'break off' the current research task and return to it later:

> P21 (Tax Associate): It's rare that I look back and use them, but they just help me express exactly where my thoughts are going and I guess they might be helpful if I need to break off, because I'll know exactly where to go back to. Also, I'm giving a talk, so it might be a useful foundation for that. I don't know exactly what's relevant and not relevant at the moment, so making notes is useful.

None of the electronic legal resources used by the lawyers in our study supported this type of note-taking behaviour.

The end result of content analysing and synthesising is often a piece of written work (such as a report to a client or a research note to a senior member of the firm (such as an Associate or Partner). As explained by this DR Paralegal and DR Trainee, analysing does not simply involve lifting information from a document, but applying that information to the information-seeking problem at hand:

> P1 (DR Paralegal): Then we'd read through all of this document and use the information in our report to the client. We wouldn't just copy and paste the information from the document, but we would be using the results of the search. We would say 'this is the name of the law,' 'this is what it prohibits,' and then we'd move on to another country or issue.

R: What do you do afterwards in order to compile your research note? P14 (DR Trainee): I'd print it all off and there's a kind of structure that you would follow in writing a note. You would write about the task that you've been given, what sources you've consulted and then the law and then how the law applies to the situation that you're talking about. So I would use the information that I've found to write about what the law is and then I'll have to apply it to the situation and case.

This Tax Trainee uses both manual and computer-based methods of synthesising the content of a document that he has deemed to be potentially useful. He explained that, when finding information to form the basis of a talk that one of the firm's Partners is to give in the near future, he printed and highlighted parts of documents that he deemed useful and copied and pasted electronic versions of the documents into a Word document. The parts he copied and pasted are those which correspond to the parts he has highlighted on his printouts:

P7 (Tax Trainee): What I do is I put it in a folder [points to paper folder on desk] and then I'll highlight what I think's relevant and then, for this particular one, I'll just cut and paste into this Word document what I'm doing, what I think are the key provisions [switches to open Word document containing the notes]. But normally, if I'm researching and writing a note then you wouldn't copy and paste chunks of text from what you've found, but as this is for someone else, he wants all sorts of chunky stuff that he can read and form his own opinions. By the end it will just turn into bullet-pointed crib notes for a speech.

Academic lawyers who analysed the content of legal documents also followd a similar pattern of applying the information found to the information-seeking problem at hand in order to form the basis of their written work:

A30 (2nd year LLB student): There will be one or two paragraphs in the judge's decision in which he states the preceding law and preceding authorities and says how these apply to the facts of this case and how it leads to this conclusion. And basically that's what we're looking for. R: Then what would you use those facts to do? A30: We'd apply them to the problems that we were facing if it's a problem question or say that in these circumstances, this and this applies [pauses] this case looks pretty much similar to that, so we can infer such and such [pauses] but if it's a problem question, it normally involves slightly more philosophical discussions of points or statements and will basically use the case report as an illustration either way for or against the statement.

5.5.9   Collating (D, C)

The penultimate behaviour subsumed under 'selecting and processing' to be identified in our study is 'collating' which was found to operate at the document and content levels. As with some of the other lower-level 'selecting and processing' behaviours, collating has not, to the best of our knowledge, been identified in previous studies. This behaviour involves *"the physical act of drawing together documents and/or content for later use"* and although it appears similar on the surface to Ellis's (1989) behaviour of 'ending' (*"the assembly and dissemination of information or the drawing together of material for publication"* - Ellis et al, 1993, p. 365), it is actually quite different. This is because although Ellis's 'ending' behaviour shares a similar definition to our definition of 'collating,' typical 'ending' activities in studies by Ellis and his colleagues involved searching for final pieces of information to fill gaps as opposed to the collating behaviour that we describe in this section. Our behaviour of collating is also likely to bear some surface similarity to Meho and Tibbo's (2003) behaviour of 'synthesising' although we do not discuss this further as Meho and Tibbo do not make reference to synthesising behaviour other than to stipulate that it occurs during their 'processing' stage of information-seeking (and we assume that their behaviour refers to the *cognitive* as opposed to the physical act of drawing together material). Like analysing and synthesising, collating was not commonly observed, again probably due to the nature of the information-seeking task in our study, however some evidence for this behaviour was identified across all groups of lawyers that took part in our study.

As explained by this DR Associate, it is important for lawyers to retrieve information that's in a 'presentable' format and that can be easily transferred into some form of output:

R: You mentioned a bit about the need, sometimes, for quick answers. Are there any other things that are important for lawyers? P5 (DR Associate): So we've said we want to be up-to-date, we want to get information quickly. I think as well, the other point was getting it in a good format - something that's presentable. I think as well, something that you can easily transfer into something. So, for example, you want to take a quote that a Lord has said and put that into a Word document, that's quite useful, you can do that.

The main way in which lawyers collated documents and content was by using facilities within electronic legal resources to print groups of legal documents at the same time. This DR Associate explained the 'print list' facility within LexisNexis Butterworths, which allows lawyers to tick which documents from a list they would like to print and then print all of the documents in a single batch job:

P5 (DR Associate): One good thing about Halsbury's is the print list. You can check different things, you don't have to print as you go if you like. You can just go to the link of all the different databases or different areas you've been looking at. And they have this tree. So it starts off, for example, 'Contracts,' and then it'll go down to all the different types of contracts and there are drop-down menus all the way through, so you can just go down and check what you like and click 'print' and then it all comes together.

This was sometimes followed by highlighting sections of the documents that are deemed useful with a highlighter pen and arranging the documents in order in a paper folder, or adding sticky tabs as markers to denote that the page contains useful content.

5.5.10 Editing (C)

The final behaviour subsumed under 'selecting and processing' to be identified in our study is 'editing,' which was found to operate at the content level. This behaviour has also not, to the best of our knowledge, been identified in previous information-seeking studies. This behaviour involves *"preparing and arranging content for later use by making revisions or adaptations."*

The main way in which lawyers edited content that they had found and deemed relevant was to paste the content into Microsoft Word. As these two 2nd year undergraduate students explained, this allows them to apply formatting to certain parts of the content and delete parts of the content that are not deemed to be useful:

A11 (2nd year LLB student): I would then just copy-paste in Word. I find it easier to work in Word format because you can italicise, you can underline, rather than just reading a mass of text.

A24 (2nd year LLB student): I generally like to highlight it and put cases into Word and print it. I prefer the layout, how it's printed. You can get rid of bits that you are not really interested in.

This 2nd year undergraduate student explained that, for a legal case, he was usually uninterested in any of the content other than the judgement of the case (i.e. the judge's decision) and therefore pasted the content into Microsoft Word and deletes the beginning part of the case report:

A30 (2nd year LLB student): Once we'd saved a copy on our hard disk I think we could edit it. And I think a lot of people do edit case reports by cutting out all of the beginning. For example, in this document you have the headnote and there would often be a brief statement by the judge, the judge's decision, arguments by the lawyers [pauses] and only then do you have the judgment, which is really what most people are concerned about. So most people would just highlight all of these and delete them. [Laughs].

5.6 Distributing documents, content and search queries/results

The fourth and final broad category of 'distributing documents, content and search queries/results' is also a standalone information behaviour and involves handing or sharing out entire documents, particular content or search queries/results to others (almost always colleagues who had made the initial information-seeking request). Theoretically, it is also possible to distribute resources and sources (or at least links to them). However, there was no evidence for distributing behaviour at the resource or source levels. Once again, distributing behaviour was not commonly observed, probably due to the nature of our naturalistic task. In addition, distributing behaviour was only mentioned or displayed by practicing lawyers, for the most part by information support staff such as Practice Development Lawyers/Assistants (although sometimes by Trainees who distributed the results of their research to their superiors). Distributing was also mentioned by an academic librarian, who often received remote enquiries by e-mail and returned the results using the same medium.

5.6.1 Document distributing

Distributing at the document level was often carried out by e-mailing documents that had been found and deemed to be useful or relevant to colleagues. Sometimes this involves a manual process of writing an e-mail and attaching the documents to it:

> *A20 (Librarian): This enquiry was for a lecture which the professor in question was giving in a couple of days. He wanted an Australian case and he was unfamiliar with Westlaw and so he wasn't sure how to find it. So he rang me up and asked me how to get hold of it. And so what I did was to track the case down and e-mail the result to him.*

Other times, as with the DR Practice Development Assistant example below, it involves using e-mail tools within a particular electronic legal resource to send the document to someone (without the need to manually write an e-mail):

> *P10 (DR PDA): The good thing about [the firm's internal Knowledge Management database] is that some of the documents are available, as you can see here [points to an MS Word icon] and you can just double click on it and print it out or send it straight from there to the person that requests it by selecting it [ticks checkbox next to one of the results] and clicking here [clicks on 'send selected item' combo box] and just click on 'go.' [Nothing happens for a while].* **R: What normally happens when you do that?** *P10: You click on 'go' and it will create an email with a link and then you'd have to go into your e-mail and send that to the person.*

5.6.2 Content distributing

Distributing content was also often achieved by sending manual or electronic resource-assisted e-mails. As the DR Trainee below explained, content distributing can involve structuring content into a coherent document (in this case a standard-form research note document) and then distributing that document (in this case, the Trainee e-mailed and printed a hardcopy of the research note to distribute to the senior member of the firm who commissioned the research):

> *P14 (DR Trainee): There's a standard form document that you put [a research note] into, like an internal note and then I would normally give a hardcopy of it to whoever's asked me to do it but also send them a softcopy by e-mail, like a link to it.*

However, content distributing does not always follow the creation of a formal document such as research note and sometimes can simply involve simply pasting relevant sections of content into an e-mail. This DR Practice

Development Assistant, for example, pasted the part of a standard legal form that contains a list of boilerplate clauses (i.e. standard pieces of legal advice) into an e-mail, without transforming the textual content in any way (for example by attempting to analyse the potential suitability of each clause for inclusion in a particular legal contract):

> *P15 (DR PDA): [Reads out text deemed relevant]. Then I would just copy that, send it to the fee-earner and say that this was last updated in June 2005 by [a certain Partner] and then I would go back to [the internal Knowledge Management database] and continue searching to find more documents that have those clauses.* ***R: So would you just copy and paste example clauses from the documents or would you synthesise them together yourself?*** *P15: For pure speed I would just copy and paste the text into an e-mail and say 'please find attached the boilerplate clauses'* ***R: So the fee-earner will bring them all together themselves?*** *P15: Yeah, because normally they've already got a contract and are just looking for suggestions of terminology or they want to see whether the other side have correctly drafted it and it hasn't got any leaks in it or it's wrong and won't protect any of the parties in the document and we need to send back comments to the other side.*

Similarly this Tax Trainee explained how he structured an e-mail reporting the results of his research task to the Associate who had asked for the research to take place. In this case, the research task was rather informal and the Trainee only needed to summarise his findings and refer the Associate to the legal material to support these findings (in this case the Capital Gains Act). Within the e-mail, the Trainee also provided the Associate with hyperlinks to the important articles found on the topic:

> *P23 (Tax Trainee): In the beginning I say 'well you asked me to do this,' so that everyone knows our terms of reference in that if you made a mistake in the first line, then it saves her from bothering to read through all of your research. And then if you've got it right, then obviously she can be more confident that your research is correct. Then a quick summary of what I've found and the section that she should look at -Section 12 of the CGA [Capital Gains Act] and then links to these articles, with a quick summary in the e-mail of what each of the articles is about so that she doesn't have to open each one up in order to see what she's gonna get.*

5.6.3 Search query/result distributing

The final level at which distributing behaviour was found to operate was the search query/result level. Although rare, sometimes information support staff distributed information about the search queries that they used and information about the results returned when undertaking a particular piece of research on behalf of a client (or more usually, a colleague). As with the distributing at the other levels, this was usually facilitated by e-mail (and was always achieved manually, probably because the electronic legal resources used by the lawyers that took part in our study do not offer the facility of semi-automating this process in the same way that many of them allow e-mails to be sent from within the library itself). In this example, a DR Practice Development Assistant explained how he would usually structure an e-mail that contained details of the search queries that he used and results that he found. He explained that he often had an e-mail dialogue with the person who asked for the research to take place, giving them details of the search queries used and results returned at different stages in the information-seeking process. This, he explained, keeps the person who requested the research 'involved' in the process and also allows them to steer the research back on course if they feel that they are not receiving the sorts of results that they require:

> ***R: If you were sending them a document that structures all of that, how would you structure that document?*** *P19 (DR PDA): It's usually just in the e-mail attaching it. I'll usually just start 'blah blah blah, here are the ones you asked for' and then just go through the databases and the order, then the search terms I've used and then put 'see attached the results at the bottom.' Sometimes you can go a bit too crazy on trying to get it really tight and you only give them one or two things, where you might be missing a lot of stuff, so try and keep them involved as much as possible. I usually try to say 'I'm*

searching this now.' So sometimes I send it all in one package, other times I'll say 'I've just finished searching Lawtel, for example, this is what I did, here are the results' and see if they are happy with that. And quite often that gives me a little bit of an indicator as to how I should pitch it if I want to search another database. So it saves times both ways.

No evidence was found of search query and result distributing amongst other types of (non-support) lawyers. As this Tax Trainee explained, the Associate who issues him with research tasks would 'expect you to get the searching right' and therefore would not be 'very interested in the search terms' used:

P23 (Tax Trainee): Particularly because on this occasion, I got the answer and so she's not very interested in the search terms. If I hadn't or I hadn't found it or hadn't spent much time on it, then I suppose she might want to look further on it herself. But they kind of expect you to get the searching right, understandably enough.

5.6.4   Summary of distributing behaviour

Distributing behaviour was found to be displayed mostly amongst the practicing lawyers in our study (and, for the most part, by members of information support staff). As with many of the other lower-level selecting and processing behaviours, distributing has not, to the best of our knowledge, been identified in any previous studies of information behaviour. We found distributing behaviour to operate at the document, content and search query/result levels. However, it is also theoretically feasible that information-seekers might choose to distribute resources or sources (or at least links to them). Most distributing behaviour, at all levels, was achieved by sending e-mails to the client (or more often colleague) that had requested that a particular piece of research be undertaken or by printing and physically handing over hardcopies of documents. When distributing entire documents by e-mail, sometimes lawyers used tools within electronic legal resources to semi-automate the process (by allowing the e-mail to be sent from within the library tool itself rather than from a separate e-mail client). The process of distributing content and search queries/results by e-mail was a manual one, probably due to the fact that the electronic legal resources used to facilitate these levels of distributing behaviour do not have tools to semi-automate the distribution of particular content within documents or search queries/results.

5.7   Summary and reflection

In our study, we identified four broad high-level information behaviours: 'identifying and locating,' 'accessing,' 'selecting and processing' and 'distributing.' We also identified several lower-level behaviours (most of which were originally identified by Ellis and his colleagues and by Meho and Tibbo in information-seeking studies in other disciplines) that could be subsumed under the broader behaviour of 'identifying and locating.' These were surveying, monitoring, searching, browsing, distinguishing, filtering, selecting, extracting and chaining. In addition, we identified many lower-level behaviours that, to the best of our knowledge, have not been identified in previous information-seeking studies. These were recording, updating, history tracking, analysing, synthesising, collating and editing. In this chapter we have described, for each of the behavioural characteristics identified, the ways in which the behaviour was achieved during observation or mentioned during interview. In addition, where relevant, we have also indicated how common each behaviour was amongst the different groups of lawyers that took part in our study (i.e. taught students, research students and staff, contentious Dispute Resolution lawyers and non-contentious Tax lawyers). The key findings in this regard were that some behaviours were only observed amongst practicing lawyers and not academics (e.g. distributing) and some behavioural subtypes only by taught students (e.g. heavily directed surveying). Most behaviours, however, were mentioned and observed across groups of lawyers. Ellis's behaviour of 'verifying' (*"checking the information and sources found for accuracy and errors"*- Ellis et al., 1993, p.364) was not identified in our study. We suggest that this is not because accuracy is not as important for lawyers as it is for physical scientists but, quite

the opposite, that legal documents must be checked thoroughly for accuracy *before* they are made available in electronic or paper form.

It is important to note, however, that the ways in which behaviours were achieved are by no means intended to be exhaustive and we fully expect that lawyers in different academic institutions and in different firms might achieve the same behaviours in different ways (if only due to the fact that they might use different electronic resources that afford certain ways of achieving particular behaviours but not others). It is also important to note that our interpretation of how common particular behaviours are was informed through questioning, sometimes focused on uncovering details about particular behaviours, as well as through observation. Whilst we do not believe that this influenced our findings in a negative manner, our findings relating to how frequently behaviours were displayed are not intended to represent 'hard' data from which any strong quantitative interpretations or comparisons can be made.

Our work makes several theoretical contributions to the field of Information Science. Firstly, it serves to validate Ellis's model in the new academic domain of law and in the new workplace domain of a large multinational law firm. Secondly, it serves to validate Ellis's model through a new research method of naturalistic observation (combined with an in-depth interview element). This is particularly useful due to the fact that all of the previous studies that identified information behaviours (e.g. those by Meho and Tibbo and Ellis/his colleagues) only used semi-structured interviews. This means that previous studies have only been based on participants' *reports* of the behaviour that they display as opposed to observed behaviour with paper-based or electronic resources. Our study has shown that these behaviours *are* actually displayed by lawyers when they use electronic legal resources. In addition, our findings extend Ellis's original model to include behaviours pertinent to legal information-seeking (e.g. updating and history tracking), and broaden the scope of Ellis's original model to cover information behaviours that overlap with information-search and information *use* as well as information *seeking* behaviours. Finally our findings enhance the potential analytical detail of Ellis's original model through the identification of mutually-exclusive pairs of subtypes of behaviour and through the identification of different levels at which many of the behaviours can operate.

Our approach for identifying information behaviours proved to be fruitful, resulting in rich think-aloud data. This data was not only useful for addressing our goal of understanding *what* lawyers do when looking for electronic information, but also for helping us to understand *why* they do it. We believe that our approach was particularly successful because it combined observing a naturalistic information task with asking questions – two approaches that are highly complementary. Indeed, we believe the three most important aspects of our approach, which helped to ensure its success were:

1. Asking participants to think aloud whist performing a broad, naturalistic information task (a task that had the potential to encourage the display of a broad range of information behaviour).
2. Asking participants questions about what they were doing and why whilst they were performing the task.
3. Asking pre and post observation interview questions to get a broader picture on the information behaviours identified fit in to participants' broader information work.

We therefore believe this approach could be used to frame future studies aimed at identifying electronic information behaviour, particularly in other domains. We would also expect that a competent qualitative researcher would be able to use our approach to investigate the *usability of particular electronic resources*.

Chapter 6: Informing the development of two novel evaluation methods

This chapter at a glance…

In this chapter we:

- Discuss the development and early testing of the Information Behaviour (IB) methods – two novel methods for evaluating the functionality and usability of electronic legal resources (based on the findings of our study of lawyers' information behaviour).
- Present the current version of the methods, along with two worked examples showing how the methods can be used to evaluate the functionality and usability of an electronic legal resource.

6.1   Overview

In this chapter we discuss the early development and testing of and present the current versions of the Information Behaviour (IB) methods, two novel methods for evaluating the *functionality* and the *usability* of electronic resources. The IB methods are both underpinned theoretically by the refined version of Ellis's (1989) model of information behaviours that we presented in chapter 5. The IB methods are based on the core premise that information-seeking models from the Information Science domain can provide a useful insight into the behaviour that users display when interacting with electronic resources. In section 6.2, we present an introduction to the IB methods, consisting of an overview, a discussion of the methods' place alongside other functionality and usability evaluation methods, a discussion of the rationale behind the methods and an overview of the information behaviours at the heart of the methods. We then, in section 6.3, discuss the development and early testing of the methods, explaining how they evolved and discussing the insights gained. Next, in section 6.4, we present the current versions of both methods (supported in section 6.5 by the presentation and discussion of two detailed examples of how the methods can be used to evaluate both the functionality and usability of the current public version of the LexisNexis Butterworths electronic legal resource). Finally, in section 6.6, we discuss the benefits and limitations of using the methods.

6.2   Introduction to the Information Behaviour (IB) methods

6.2.1   Overview of the IB methods

In his doctoral thesis, Ellis (1987) presents generic design recommendations for supporting each of his identified behaviours in information retrieval systems. However, Ellis also asserts that the behaviours can inform systems *evaluation* as well as design. Motivated by this assertion, we developed two novel, specialised User Evaluation Methods for electronic resources based on these behaviours. The methods are novel as they are based on the observed information behaviour of lawyers and are theoretically underpinned by an extension of Ellis's behavioural model of information-seeking. The methods are specialised in the sense that they are intended to be used to evaluate *electronic resources* as opposed to other types of interactive system.

The IB methods were developed based on the premise that observed information behaviour can provide a useful structure for evaluating electronic resources, by using the observed behaviours as theoretical 'lenses' to evaluate the *functionality* of electronic resources (i.e. the features provided by the resource aimed at supporting users) and the *usability* of these resources (how easy to use they are). Although both the functionality and usability IB methods are underpinned by the same information behaviour theory, they can be regarded as two separate, self-contained methods in their own right. An IB *functionality* evaluation aims to provide data on the range of

functionality supported by a particular electronic resource (and provides the basis for discussing whether this range is appropriate). An IB *usability* evaluation aims to provide data relating to the difficulties that users face when using a particular resource and how severe and easy to address the evaluator considers the usability issue(s) to be.

Both the IB functionality and usability methods are underpinned by three categories of information behaviours that we observed across electronic legal resources. The first category of behaviours are those we regard as 'core' information-seeking behaviours. All of these behaviours have been observed across disciplines by David Ellis and his colleagues. These include accessing, surveying, monitoring, searching, browsing, chaining, extracting, selecting, distinguishing and filtering. The second category of behaviours are particularly pertinent to legal information-seeking (and have not been noted in other disciplines). The two behaviours in this category are updating and history tracking. The third and final category of behaviours involve the *use* of information that has been found during information-seeking and include the behaviours of analysing, synthesising, collating, editing, recording and distributing.

These three categories of information behaviours feed in to the functionality and usability evaluations in different ways. They are used in an IB *functionality* evaluation as a framework for assessing the functionality provided by own and competitor electronic resources. An 'own resource' is an electronic resource which an evaluator has a direct stake in (for example, it is a resource that the evaluator or the company he or she works for has been involved in developing, evaluating or marketing). A 'competitor resource' is developed and sold by another (possibly competing) firm. Competitor resources are likely to be of interest as they might share similarities with one or more own resources (such as supporting similar information tasks, or having similar functionality support). When assessing 'own resources,' an IB functionality evaluation involves evaluators discussing whether and in what ways an own electronic resource currently supports the information behaviours, at a number of levels identified in our empirical study (i.e. whether and in what ways the resource supports users in working with the *resource* itself, *sources* within the resource, individual *documents*, *content* within a document and *search queries/results*). This also involves evaluators discussing whether they might support user behaviours/levels in additional ways and considering whether it is necessary to continue to support all of the currently supported behaviours/levels. When assessing competitor resources, an IB functionality evaluation involves evaluators exploring the resource to determine whether and in what ways the competitor resource currently supports the behaviours, at each applicable level.

The three categories of information behaviours are used in an IB *usability* evaluation as the foundation of think-aloud tasks that are set to intended or actual users of the electronic resource. These tasks are based on our empirical data of lawyers' information behaviour In an IB usability evaluation, evaluators set a number of behaviour-focused tasks to users, who are asked to perform the tasks whilst thinking aloud. Evaluators then identify usability issues from the resultant think-aloud data and make summary judgements on how severe they consider the issues to be (i.e. whether or not they need immediate attention) and the amount of effort they consider to be required to address the issues.

Although the methods were developed based on empirical data of *lawyers'* information behaviour, we hypothesise that the IB methods can be used to evaluate both legal and non-legal resources. This hypothesis is based on evidence that information behaviour has been found to be similar across domains (see Ellis 1989, Ellis 1993, Ellis et al. 1993, Ellis and Haugan 1997) and has not been found to have changed substantially since Ellis's studies in the 1980's and 1990's even though electronic information-seeking is now far more common (see Meho and Tibbo, 2003). Indeed, the IB methods are *extensible* and *customisable* and can therefore be tailored to include any additional or alternative behaviours relevant to a new, non-legal, domain. For example, currently the behaviours used as the basis of the IB evaluation methods do not include behaviours identified in other empirical studies but not in our own (such as 'verifying' and 'networking'). However it is possible to incorporate additional relevant behaviours into both an IB functionality and IB usability evaluation. For example, when evaluating the functionality of electronic resources designed to support Physical Scientists' information work, evaluators might wish to assess functionality support for verifying behaviour (i.e. support for "*checking the information and sources found for accuracy and errors*" - Ellis et al., 1993, p.364). Similarly

when conducting an IB usability evaluation of an electronic resource designed for Physical Scientists, it is possible to set one or more tasks focused on verifying behaviour (e.g. 'check one of the documents you have found for accuracy and errors'). It is therefore possible to customise the IB methods to new domains *by including or excluding certain behaviours*. In effect, this means that when applying the IB methods to a new domain, it is only necessary to change the *theory base* of the methods (i.e. the information behaviours to be used to frame the evaluation). It is not necessary to change the methods themselves. Indeed, we hypothesise that the IB methods can be applied to a wide range of domains with only slight modification to the theory base. As some identified information behaviours share conceptual similarities (e.g. selecting and distinguishing), care must be taken when customising the methods to ensure that all behaviours used as part of the theory base are clearly and precisely defined. This should help minimise the chance of users of the methods misunderstanding what the behaviours entail (and therefore making functionality suggestions or defining usability tasks related to an incorrect understanding of a particular behaviour).

The IB methods both address a 'niche in the market' of evaluation methods as they have roots grounded in information theory (and, as discussed previously, are empirically grounded). The IB methods are specialised in the sense that they aim to evaluate the functionality and usability of *electronic resources* and not interactive systems in general). The methods aim to provide a bridge between the domains of Information Science and Human-Computer Interaction, by providing users with the opportunity to conduct functionality and usability evaluations that are *highly structured* but are also *flexible* in the sense that they can be tailored to particular domains and/or foci. Part of this structure is provided by the theoretical underpinning of the information behaviours and levels which provide a framework for assessing electronic resource functionality (as part of an IB functionality evaluation) and for setting behaviour-focused tasks (as part of an IB usability evaluation). Part of the structure is also provided by supporting forms, which can be used to record the output of the evaluations (see appendices 6, 7 and 8 for examples of these forms). Flexibility is provided by the *extensibility* and *customisability* of the methods, which can be tailored to facilitate the evaluation of particular information behaviours or levels of interest.

In the sections that follow, we discuss the IB methods' place alongside other evaluation methods, along with the rationale behind the methods. We then discuss the information behaviours at the core of the methods, followed by a discussion on the early development and testing of the methods. We then describe how to carry out both an IB functionality and IB usability evaluation, supported by two worked examples. Finally, we discuss the benefits and limitations of using the IB methods.

6.2.2   The IB methods' place alongside other evaluation methods

Functionality and usability are two aspects of an interactive system that can be evaluated with the aim of making design suggestions. Evaluating the *functionality* of a resource involves in some way examining the features provided by the resource aimed at supporting users. Evaluating the *usability* of a resource essentially involves examining how easy to use it is (the International Organization for Standardization defines it as *"the effectiveness, efficiency and satisfaction with which specified users achieve specified goals in particular environments"* – ISO 9241 available from www.iso.org).

Surprisingly little has been written about *functionality*-based evaluation methods. Mack and Nielsen (1994) highlight that these types of evaluations, known as 'feature inspections,' *"focus on the function delivered in a software system: for example, whether the function as designed meets the needs of intended end users"* (p. 6). Mack and Nielsen (1994) also highlight that feature inspections can include design as well as evaluation of features within a system. However, there is almost no literature detailing existing feature inspection methods. This is with the exception of work by Bell (1992), who details a method for designing programming languages that are easy to write that is aimed at evaluating both the *facility* provided by a programming environment (which Bell describes as *"the ability to solve problems easily,"* p. 7) and the *expressiveness* of the environment (which he describes as *"the ability to state solutions to hard problems simply,"* p.7).

Existing methods in the HCI domain provide different ways to examine the usability of interactive systems and many of these methods have a specific focus which helps to make them useful in a specific way. For example, Cognitive Walkthrough (Polson et al., 1992; Wharton et al. 1994) allows those using the method to assess the learnability of an interface, with a focus on the user's cognitive processes and perception. Similarly CASSM (Concept-based Analysis of Surface and Structural Misfits, Blandford, Green et al., 2008) is a method that can be used to highlight mismatches between how users conceptualise aspects of an interactive system and how the system supports user concepts. Some of these user evaluation methods have been applied to a digital library/electronic resource context. For example, Blandford et al. (2004) used a range of evaluation methods (Heuristic Evaluation, Cognitive Walkthrough, Claims Analysis and CASSM) to evaluate various digital libraries. Similarly, Blandford, Green et al. (2008) present a worked example of how CASSM can be applied to think-aloud data of postgraduate HCI and Library and Information Studies students using a range of electronic resources. Blandford and her colleagues have also tailored Claims Analysis to a digital libraries context (see Blandford et al., 2006; Blandford, Keith et al., 2007).

As highlighted by Blandford and Green (2008), evaluation methods can be broadly classed along three dimensions:

1.   Whether or not they are carried out with the active involvement of users (those that involve users are known as *empirical* methods and those that do not as *analytical* methods).
2.   Whether or not they are carried out with a running system.
3.   Whether or not they are carried out in a realistic context of use.

According to Nielsen (1993), usability evaluation with 'real users' *"is the most fundamental usability method and is in some sense irreplaceable, since it provides direct information about how people use computers and what their exact problems are with the concrete interface being tested"* (p. 165). Similarly Landauer (1995) has described user testing as the 'gold standard' for evaluation. However, user testing can be resource intensive and this has resulted in the use of evaluation methods that do not require user involvement (such as those described in the previous paragraph). The IB functionality method is not carried out with active end-user involvement per se, however evaluators may seek to use the method armed with usage data of their resources in order to help them reason about increasing or reducing functionality. Therefore this method can be classed as an *analytical* method. The IB usability method, on the other hand, is primarily an *empirical* method as it involves analysing think-aloud data of <u>users</u> performing certain information behaviour-focused tasks. The IB usability method is also partly *analytical*, as it involves <u>evaluators</u> *identifying usability issues* from the think-aloud data and *deciding on how severe and easy to address they are*. The mixed nature of the IB usability method allows it to benefit from the use of rich user data (where it would otherwise be difficult for evaluators to predict how users are likely to behave with an electronic resource) and from a theoretical underpinning to drive task setting and analysis (where it would otherwise be difficult to analyse this rich data in a structured way).

As the IB usability method involves observing real users performing behaviour-focused tasks, it requires the use of a *running system*. An IB functionality evaluation can be supported by (but does not require) a running system, particularly as evaluators may be familiar with most of the functionality provided by the resource under evaluation. Although underpinned by information behaviour theory, an IB functionality evaluation does not involve a *realistic context of use* (as it does not actively involve users). On the other hand, whilst the think-aloud tasks that are part of an IB usability evaluation might be performed in an artificial setting, the tasks themselves (also based on empirically observed behaviour) *do* aim to ensure a realistic context of use.

We also believe the IB methods address a clear 'niche in the market' of evaluation methods. Whilst there are other methods that are underpinned by theory (such as Polson et al.'s Cognitive Walkthrough), we are unaware of any other evaluation methods with roots grounded in *information* theory. In addition, the IB methods are specialised in the sense that they aim to evaluate the functionality and usability of *electronic resources* and not interactive systems in general. Whilst Blandford and her colleagues have applied and tailored various evaluation methods to an electronic resource context, we are unaware of any other evaluation methods *developed especially*

to evaluate electronic resources. We also believe that the IB methods are novel because they are empirically grounded.

6.2.3 The rationale behind the IB methods

The IB methods were developed based on the premise that information-seeking models from the Information Science domain, particularly our refined version of Ellis's information-seeking behaviour, can provide a useful insight on users' behaviour when interacting with electronic legal resources. Usability experts, or other stakeholders in an electronic resource, can use the IB methods to assess the *functionality* provided by or the *usability* of own or competitors' products. The resultant insights can then be used to improve the design of these products.

The rationale behind the IB methods is closely related to Norman's (1986) notion of 'bridging the gulf' between user and system (see figure 9). The IB *functionality* method aims to support people with a stake in a resource to 'push' the resource closer to its users by supporting them in examining whether the resource supports an appropriate range of user behaviour. The IB *usability* method aims to 'push' electronic resources closer to users by supporting stakeholders to identify usability issues that, if addressed properly, should lead to improved support for user behaviour.

This is with the broad aim of improving the design of interactive systems so that they 'speak the user's language' as opposed to forcing users to learn the system's language. Through the IB methods and the resultant functionality and usability insights that they provide, we seek to help bridge the gulfs between a system and its users by making the system more closely match user goals and therefore become more user-centred.
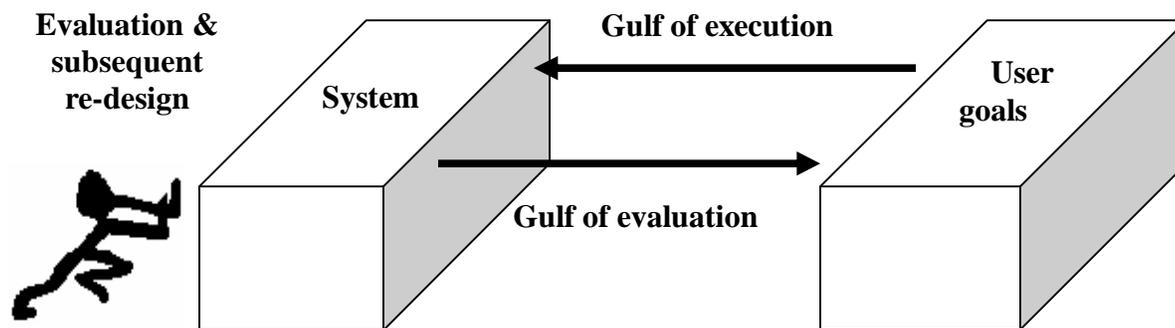
Figure 9: How evaluation and subsequent re-design can bridge the gulfs of evaluation and execution by 'pushing' the system closer to the users and their goals. Diagram adapted from Norman (1986).

6.2.4 The information behaviours at the core of the IB methods

At the core of both the functionality and usability IB methods are a series of information behaviours that have been observed empirically amongst groups of academic and practicing lawyers as part of this thesis. Many of these behaviours have also been shown to apply to other domains and fields, including several physical and social science domains (see Ellis 1989, Ellis et al. 1993, Ellis and Haugan 1997, Meho and Tibbo 2003) and the field of English Literature (Smith 1988, also reported in Ellis 1993). In this section, we provide an overview of these behaviours and recap on the concept of 'levels,' which is particularly important for understanding the IB methods. Tables 5, 6 and 7 in chapter 5 list and define the individual and groups of information behaviours that are used as a basis for evaluating the functionality of electronic resources (and as a basis for defining behaviour-focused tasks to be used to evaluate the usability of electronic resources). They also highlight the various *levels* at which it is possible for users to perform (and therefore electronic resources to support) each behaviour. These

levels were also identified as part of our study on academic and practicing lawyers' information behaviour. However, none of the levels are specific to legal resources.

One of these levels is the 'resource' level. Many of the behaviours can be performed at this level (for example, it is possible to *search* an Internet search engine in order to locate resources). However, as the IB methods are aimed at evaluating electronic resources themselves, most of the ways of supporting behaviours at the resource level are beyond the scope of the methods. This is with the exception of 'accessing' behaviour (as it is highly important for designers to provide an easy means of accessing an electronic resource). The other levels include the source level (which refers to behaviours that are concerned with an information source or sources within a particular electronic resource), the document level (which refers to behaviours that are concerned with an individual document or documents within a particular information source), the content level (which refers to behaviours that are concerned with content within a particular document) and the search query/result level. Many of the behaviours can operate at multiple levels. Some behaviours, such as surveying and monitoring, operate at a combined 'document and content' level. This level is used when it is difficult, impossible or undesirable to separate whether a particular behaviour is performed on the document, or the content within it.

The concept of levels is important for structuring both IB functionality and usability evaluations. An IB functionality evaluation involves using the information behaviours discussed in this section to frame a functionality evaluation of an electronic resource, by examining whether and in which ways each behaviour is supported by the resource *at each applicable level*. 'Applicable' levels are those at which the behaviour was found to operate (see chapter 5). An IB usability evaluation involves setting behaviour-focused tasks to intended or actual users of an electronic resource, asking them to think aloud whilst performing the tasks (i.e. to verbalise their thoughts, actions and feelings) and analysing the resultant think-aloud data. Depending on the focus of the usability evaluation, these tasks may be aimed at performing a particular behaviour *at a certain level* (for example, one of the possible tasks that can be set as part of a custom IB usability evaluation is 'browse to see whether a particular *source* is available in the electronic resource (for example by browsing through a list of sources that the resource contains'). This task aims to encourage users to perform browsing at the *source level* as opposed to the more common document level (where the user browses to locate particular *documents*, not sources).

6.3   Development and early testing of the IB methods

In this section we explain how the IB methods have evolved, through a process of iterative development and testing. We begin by discussing our starting point and provide brief details about the early versions of the methods and why changes were made. We then discuss a series of user and developer pilot studies which helped to shape the methods (and, in particular, the IB usability method).

6.3.1   The need for user evaluation methods grounded in information-seeking theory

As we have discussed earlier, the broad motivation of the IB methods was to address a 'niche in the market' of usability evaluation methods – to develop one or more HCI methods aimed at evaluating electronic legal resources that were grounded in Information Behaviour theory. This niche is referred to as an opportunity or need by Blandford and Green (2008). Our aim was to feed empirical findings from our studies on academic and practicing lawyers' information behaviour into an HCI method or methods that were theoretically grounded in a particular information-seeking model (in this case the refined version of Ellis's 1989 behavioural model that is presented in this thesis). In our studies, we identified several information behaviours performed by lawyers when looking for electronic information and we soon realised the potential for usability experts other stakeholders in an electronic resource to use these behaviours as lenses for which to evaluate the functionality and usability of electronic resources (electronic legal resources in particular). The core need of the IB methods

was clear from the outset: they should allow designers to in some way evaluate electronic resources with regard to the information behaviours identified.

6.3.2 Our starting point for addressing the need

Blandford and Green (2008) note the exploratory nature of identifying possible approaches for addressing a need, highlighting that when we *"start with an opportunity, or resource, the exploration is typically less concerned with finding appropriate theory and more with finding ways to develop theory into a method"* (p. 5). Our starting point for addressing the need was to theorise about exactly *how* the behaviours might feed in to a usability or functionality method. Early versions conceptualised the IB functionality and IB usability methods as only one method (which combined both functionality and usability aspects and consisted of a list of streamlined questions similar to those asked in the final IB functionality method). More specifically, for each part of an electronic resource, the initial version of the method required an evaluator to ask themselves whether the part of the resource currently supported each information behaviour or not. If it *did not* currently support the behaviour, the evaluator would then ask in what ways it might be possible to support the behaviour in that particular part of the resource. If the part of the resource *did* currently support the behaviour, the evaluator would ask in what ways the electronic resource *currently supports* the behaviour, *how well* it supports the behaviour and how it might be possible to *improve support* for the behaviour.

In order to answer these questions, we envisaged users of the method stepping through an electronic resource screen by screen. However, when we attempted this in practice, we found the process to be highly time-consuming and more difficult than first envisaged. We also found questions surrounding 'how well' behaviours are supported and how support might be improved to be rather vague and difficult to answer in practice. In addition, the particular streamlined questions asked appeared to make the evaluation process tedious, rather than even though they were intended at making the it simple for users to internalise and implement the method. This serves to highlight the importance for evaluation methods to strike a delicate balance between richness and simplicity, a point also raised by Garzotto and Perrone (2007).

6.3.3 The change from a 'theory driven' to 'experientially driven' development approach

As our 'theory driven' approach to development did not seem to have led to a method that was easy to apply in practice, we switched to a 'experientially driven' approach which involved taking a step back from the previous version of the method as asking ourselves the question 'what is an intuitive way of using the information behaviours identified in our study to evaluate electronic resources?' To help answer this question, two think-aloud studies of two electronic legal resources were conducted. One of these was a self-think aloud study of Westlaw. The other involved asking Ann Blandford - a Professor of Human-Computer Interaction who had been involved in supervising the development of the methods, to think-aloud whilst using Justis. The procedure for conducting both think-aloud studies was purposely under-defined, but involved attempting to use the information behaviours identified in our study of academic and practicing lawyers to frame an evaluation of the electronic legal resource. As we did not yet have a clear idea of the structure of the method, we could not specify exactly *how* this was to be accomplished. Instead, it was left to the think-aloud participants (i.e. myself and Ann) to use the behaviours to frame an evaluation in a way that was most intuitive. We both verbalised our thoughts as we performed the evaluation and the sessions were audio recorded and transcribed in note form to aid the further development of the method.

A big change in how we conceptualised the method came as a result of the Justis think-aloud session. Ann found herself using the information behaviours not as a framework for evaluating the *functionality* provided by Justis, but as a framework for evaluating its *usability*. Indeed, almost all of the think-aloud data she generated highlighted usability issues with Justis but, under the IB method as it stood, this data would not have been captured at all. This resulted in a shift in conceptualising the IB method – a shift away from conceptualising it as one combined method for evaluating resource functionality and usability and towards conceptualising it as two

separate methods that share the same theoretical underpinning of information behaviours. This also suggested the potential for the method to be applied by intended or actual *users* of electronic resources, not just by stakeholders in the resource, such as developers or HCI experts.

The switch to an experientially driven approach also led to a change in the scope, where the focus shifted away from supporting both electronic resource evaluation and re-design and towards supporting only *evaluation* (i.e. supporting users of the method in identifying ways in which the resource supports or might support certain user behaviours and identifying usability issues, but not determining precisely how to support behaviours at the interface level or how to address the usability issues identified).

The development process continued by trying to apply successive versions of the IB functionality and usability methods to different electronic legal resources. This led to iterative improvements for both methods. An important part of the development process of the IB functionality method included conducting a functionality survey of six of the electronic legal resources commonly used by the lawyers in our empirical study of their information behaviour. These resources included LexisNexis Butterworths, LexisNexis Professional, Westlaw, Kluwer Arbitration, Justis 5 and HeinOnline. The survey was conducted in order to identify ways in which an electronic resource might support each information behaviour, at all applicable levels and to match these ways with the information behaviour displayed by the academic and practicing lawyers in our earlier empirical study. The 'matching' process was conducted by reviewing the transcripts of user think-aloud sessions from our study and noting which of the ways that the electronic resources were found to support each behaviour were also demonstrated by one or more of the lawyers in our study. The matching process not only served to ensure that the ways of supporting particular information behaviours were *supported* by one or more electronic legal resources, but that there was *some empirical data* to suggest that lawyers might perform these information behaviours in these ways. Also, despite the fact that the task given to participants in our empirical study had the potential to discourage participants from displaying behaviours that were not directly related to finding information, we were able to gain a general impression of *how common* each way of performing a behaviours might be for lawyers. This enabled us to make a judgement on how important it is for electronic legal resources to support each of these ways.

It also became clear whilst conducting the survey that the 'ways' identified in both the survey and displayed by the lawyers in our earlier empirical study might also provide a useful framework for devising behaviour-focused user think-aloud tasks, which could be used as a springboard for identifying usability issues associated with the resource. The functionality survey (supported by our empirical data) was therefore used as a basis for both the supporting documentation for conducting an IB functionality evaluation (presented in appendix 3) and many of the think-aloud tasks that intended or actual users of the resource are asked to perform as part of an IB usability evaluation (i.e. the more behaviour-focused tasks that are set when conducting a 'recommended' or 'custom' evaluation). All of the think-aloud tasks are presented in appendix 5.

6.3.4   A series of three pilot think-aloud sessions with users of electronic legal resources

The next stage in the iterative method development and testing process involved conducting a series of three pilot user think-aloud sessions with one Trainee Solicitor and two final year Bachelor of Laws (LLB) students in order to determine the tasks and procedural details to be used as part of an IB usability evaluation (i.e. procedural details for both obtaining and analysing user think-aloud data). The three pilot participants had all previously taken part in our empirical study of lawyers' information behaviour and also kindly agreed to participate in this pilot study.

All three participants were asked to rate their perceptions of their experience using electronic legal resources in general and the version LexisNexis Butterworths under evaluation as either 'not experienced,' 'somewhat experienced,' or 'very experienced.' Their responses are summarised in table 9 below:

Pilot participant	Self-rated experience using electronic legal resources in general	Self-rated experience using the version of LexisNexis Butterworths under evaluation
Trainee Solicitor (Pilot participant 1)	Somewhat experienced	Not experienced
1st final year LLB student (Pilot participant 2)	Very experienced	Not experienced
2nd final year LLB student (Pilot participant 3)	Somewhat experienced	Not experienced

Table 9: The experience level of each of our three pilot participants in using electronic legal resources in general and the version of LexisNexis Butterworths under evaluation as rated by the participants themselves.

Although no attempt was made to recruit novice users of LexisNexis Butterworths as pilot participants, we believe that the fact all three participants rated themselves to be 'not experienced' with the resource worked in our favour, as the participants encountered many difficulties that might not have been demonstrated if participants with more experience in using the resource had been recruited and therefore this yielded useful think-aloud data for the purpose of identifying usability issues.

The first pilot session was with a Trainee Solicitor working for a small London firm of Solicitors. This participant was a vocational Legal Practice Course student when he participated in our previous study as participant number A28. The participant was asked to think aloud whilst performing many of the 'recommended' tasks in appendix 5 (the tasks he was asked to perform are denoted by the symbol ❶). The second pilot session was with a final year Bachelor of Laws (LLB) student (participant A5 in our previous study). This participant was asked to perform the same tasks as those performed by the first pilot participant, along with some additional, more prescriptive, tasks. The tasks he was asked to perform are also presented in appendix 5, denoted by the symbol ❷). The third pilot session was with another final year Bachelor of Laws student (A3 in our previous study). This participant was asked to perform only prescriptive tasks, related to particular behaviours and levels. The tasks she was asked to perform are presented in appendix 5, denoted by the symbol ❸ and did not overlap with any of the previous tasks that had been piloted. The procedure used in all three pilot studies was largely the same and has been formalised as a participant instruction sheet in appendix 5. In short, the pilot participants were asked to think aloud whilst using the resource to perform the tasks they had been allocated and were given the opportunity to ask questions before commencing each task so that any unclear task wording could be identified. Some minor changes were made to the wording of tasks as a result, in order to ensure maximum clarity.

All three pilot sessions resulted in rich think-aloud data which, upon analysis, highlighted a number of serious and not so serious usability issues. These issues are summarised in appendix 4, for each pilot participant. The user pilot sessions served to illustrate that a wide range of different types of tasks would be suitable for using as the basis of an IB usability evaluation. These tasks ranged from broad and naturalistic (but only slightly behaviour-focused) tasks, as piloted with the first two participants, to fairly prescriptive, highly-behaviour focused (but only slightly naturalistic tasks), as piloted with the third participant.

The pilot sessions were also useful in shaping the procedure for analysing the resultant think-aloud data. We had envisaged that it would be useful to categorise usability issues under the headings of awareness (whether the electronic resource makes it clear that it supports ways of performing the behaviour), action (whether the resource makes the user aware of the steps required to achieve each way of performing the behaviour), knowledge (whether the resource makes the user aware of the knowledge required to achieve each way of performing the behaviour) and feedback (whether the resource makes the user aware of their progress as try to perform the behaviour and aware of whether the behaviour has been successfully achieved). It became clear, however, from the first pilot session that this would not be a useful coding scheme in practice because the categories did not always appear to categorise the essence of certain usability issues, there was often ambiguity and overlap involved in categorising issues under these categories and the categories did not appear to add much value to the method overall. Therefore during the piloting process, the requirement to use these categories was removed and in the current version of the IB usability method usability issues are identified, but not categorised.

In a similar way, we had envisaged that it would be useful to ask people analysing the user think-aloud data to identify *positive* as well as negative usability issues (recommended by Bellotti et al., 1995, who highlight the importance of preserving positive aspects of design and suggest that purely negative feedback is unlikely to endear an HCI method to developers). However, we found the process of labelling an issue as 'positive' to be much more difficult than expected. For example, the participant spent the majority of the session unaware that whilst he was making selections to search for particular types of legal materials (e.g. legislation) in one combo box, he was submitting his search using a search button that was not associated with the combo box (i.e. the searches he was submitting were ignoring the selections he had made in the drop-down combo box because the combo box was designed to be used in conjunction with a different search button). Towards the end of the session, he appeared to realise his mistake (although he did not mention this explicitly). Becoming aware of his mistake might be regarded as a 'positive' usability observation. Alternatively, in the context of his previous difficulties using the search screen, this might be regarded as a 'negative' usability observation because the resource did not make the participant aware of his mistake sooner. Therefore we also removed the requirement to identify positive as well as negative usability issues from the data.

Overall, we deemed the user pilot sessions to be successful. They demonstrated that a variety of different types of tasks could be set for users to perform in order to yield useful data for analysing the usability of an electronic resource. This resulted in the IB usability method becoming extensible by allowing users of the method to conduct either a 'core' IB usability evaluation (where users are asked to perform a small number of broad information-seeking tasks), a 'recommended' IB usability evaluation (where users are asked to perform a range of behaviour-focused tasks) and a 'custom' IB usability evaluation, where the tasks can be selected from a bank of behaviour and/or level-specific tasks. The pilot sessions also demonstrated that our previous procedure for identifying usability issues from the think-aloud data did not work well in practice and the procedure was later simplified and improved.

6.3.5 A pilot think-aloud data analysis session with an electronic resource developer working in the Human-Computer Interaction department of a London university.

After conducting a set of user pilots to determine the tasks and procedure to be used for obtaining think-aloud data as part of an IB usability evaluation, we then conducted a pilot data analysis with a developer of electronic resources, who worked in the HCI department of a London university. This was with the aim of determining the usability, usefulness and potential for future use of the IB usability method (and suggesting ways of improving the data analysis procedure), before we presented both methods to a group of stakeholders working for LexisNexis Butterworths. The developer pilot session involved presenting an edited video clip of the three user think-aloud sessions discussed in the previous section to a developer of electronic resources, and asking him to identify usability issues from the data. Note that the developer pilot only sought to assess the *analysis* part of the usability method (i.e. identifying usability issues from user think-aloud data) and not the process of collecting the think-aloud data in the first place (which was informed by the user pilots, not the developer pilot).

The developer that agreed to participate in our study had only limited experience developing resources in the legal domain. However, the developer had been working on a funded academic research project that involved designing digital libraries for humanities users for the past two years and therefore had some experience in developing electronic resources in general. The developer also had several years of HCI experience, including in conducting usability evaluations.

We did not feel that it would be feasible to pilot the IB *functionality* method with the developer as the he only had limited experience using the resource and therefore would only have been able to evaluate the resource as if it were a 'competitor' resource. Whilst this may have been a useful exercise, we did not believe that the benefits of piloting only a limited aspect of the functionality evaluation would outweigh the drawbacks (most significantly, the need for the developer to gain a thorough knowledge of the concepts of behaviours and levels and how each behaviour and level might apply to electronic legal resources). However, it was still necessary to provide the developer with the opportunity to gain experience with the functionality supported by the resource in

order to enable him to use the IB usability method to identify usability issues. This was facilitated by introducing the concepts of information behaviours and their associated levels to the developer in a fairly light-weight and informal fashion and asking him to explore the electronic resource in relation to the behaviour(s) of his choice. This provided a structured way for the developer to assess the functionality of the resource (and therefore would provide some limited feedback on the how useful behaviours might be for framing functionality evaluations).

Before beginning the exploration, the developer was presented with a 'crib sheet' detailing the behaviours, a definition of each and a list of levels at which each behaviour or set of behaviours can operate. The developer was invited to ask questions related to the information on the crib sheet. The developer was then encouraged explore the resource in relation to behaviours/levels of his choice until he felt that he had gained sufficient experience with the resource to enable him to identify usability issues from video clips of user think-aloud sessions with the resource. He chose to explore the resource for around 30 minutes and, after this time, the researcher asked him four semi-structured interview questions surrounding the usefulness of the concepts:

1.   What previous approaches have you used to evaluate the functionality of electronic resources you have designed, if any?
2.   How useful do you think the concepts of information behaviours and levels would be as part of a 'language' for evaluating the functionality of the electronic resources that you develop?
3.   If you think these concepts would be useful (or somewhat useful), how would you suggest they are incorporated into your functionality evaluations? If you do not think the concepts would be useful, explain why.
4.   What, in your opinion, are the potential *benefits* and *drawbacks* of using the concepts of information behaviours and levels as part of functionality evaluations?

After the structured exploration of the resource, the developer was asked to watch an edited video clip of the two academic lawyers from the user pilot sessions using the resource, lasting around 20 minutes. The video clip consisted of three of the tasks from the second and third user pilots (which we now regard as the 'core' tasks of the IB usability method). These tasks were:

1.   Gain access to (i.e. log in) to the resource. This was performed by the academic lawyer from the third user pilot.
2.   Find out which parts of the resource you have access to.
3.   Think of some legal information you currently require or have recently required for your work and demonstrate, using the resource, how you might go about finding it. These two tasks were performed by the academic lawyer from the second user pilot.

Although editing the think-aloud video clips is not normally part of an IB usability evaluation, we followed this process because none of the three user pilot participants had been asked to perform *all three* of the tasks and therefore it was necessary to 'cut and paste' a user attempting each of these tasks into one video clip. No other changes were made to the video data.

The developer was asked, whilst he watched the video of the users' think-aloud sessions, to fill out a usability form similar to that presented in appendix 8. The version of the form used in the pilot session did not, however, have a column for recording 'reflections on the usability issues identified' Also, the form did not allow own usability observations (i.e. observations that were made as a result of watching the think-aloud session, but were not directly related to the think-aloud participant's comments or actions). In addition, the developer was instructed to fill out the form *whilst watching* the think-aloud session (as opposed to a mix of whilst and after watching the session, as with the current version of the form).

After the developer finished conducting the 'core' usability evaluation of the electronic resource, he was asked a number of semi-structured questions relating to the usefulness and usability of the IB usability method, along with questions about the developer's likely future use of the method. The semi-structured interview questions

were identical to the focus group questions (used later in the evaluation of the method with a group of usability experts working for LexisNexis Butterworths) presented in appendix 12. However, we also did not ask questions related to the learnability of the IB usability method as, at this early stage, we were more interested in the method's usability and potential utility (and wanted to focus on how easy to learn each of the methods were once we had 'polished' versions ready to evaluate).

The developer pilot led to a number of insights related to using the concepts of information behaviours and levels to frame functionality evaluations or discussions and identifying usability issues from user think-aloud data (i.e. data arising from users conducting behaviour-focused tasks). These insights are discussed below.

Insights related to using the concepts of information behaviours and levels to frame functionality evaluations or discussions

In his interview, the developer suggested that he found the language' of 'behaviours' and 'levels' to be 'very useful' as a framework to enable users of the IB functionality method to focus on the functionality of products 'from a critical perspective.' This, in his opinion, would allow developers or other users of the method to avoid feeling as though they had 'forgotten something' when stepping through a resource to assess functionality and would avoid evaluators 'getting hung up' on small functionality details or 'forgetting to consider' some functionality.

The main issues surrounding the developer's use of the framework to structure his exploration of the resource concerned his interpretation of behaviours and, in particular, their associated levels. For example, although the developer was given the opportunity to ask any questions about the definitions of the behaviours/levels that were presented to him, he assumed that 'gaining an overview of documents' (i.e. monitoring at the document level) meant gaining a broad overview of the structure of a document itself not, as intended, locating a document that provided a broad overview of a particular legal *area*. These types of interpretation issues highlighted the need to provide concrete examples of all behaviours and associated levels, perhaps in the form of video clips or screenshots and text captions to explain what each screenshot is intended to illustrate. Indeed, the developer asserted that the biggest potential drawback of using the concepts of behaviours and levels to frame functionality evaluations or discussions was that evaluators might have difficulty understanding the framework (which he referred to as an 'abstract structure') and 'what all the things mean.' Overall, however, the developer claimed that this was unlikely to be a big issue, as he assumed that most electronic evaluators and usability experts would be 'used to working with abstract structures' such as the concepts of information behaviours and levels.

Insights related to identifying usability issues from user think-aloud data

The developer found the IB usability method to be 'very useful' in helping him to highlight usability issues that might inform design/re-design. This, he claimed during his interview, was due to the structure provided by the analysis approach. The developer explained that a previous approach for analysing user think-aloud data which he carried out involved highlighting usability 'themes' identified from transcripts of user think-aloud sessions using electronic resources and then 'coding' parts of the transcripts based on the themes identified. The developer suggested that the IB usability method compared favourably to the 'thematic coding' approach as it provided a useful structure for examining usability issues without having to develop a coding scheme from scratch, which could sometimes 'become rigid very quickly' and 'restrict the usability issues identified.'

The developer noted that although he had only analysed 20 minutes of video data using the IB usability method, the method seemed balanced in terms of highlighting severe vs. not so severe issues and issues that seemed easy to address vs. those that seemed more difficult to address. The developer commented that he *would* use the method when evaluating electronic legal resources in future, although he does not conduct such evaluations often. He suggested that rather than take the method off-the-shelf, he would most likely 'integrate it' with the

thematic coding approach. He also suggested that it may be possible to use the method *alongside* the broader coding approach in order to ensure that broader usability themes and individual usability issues are highlighted.

The process of completing the form provided with the IB usability method was found by the developer to be straightforward overall and, according to the developer, the form made it clear what information he should write in each column. The particularly difficult aspects of conducting the usability evaluation, as highlighted by the developer, were:

- It was sometimes difficult for the developer to pinpoint the exact usability issues that user actions or comments were pointing to on the fly as this was asking a 'deeper question' that required some reflection (which he commented was difficult to do on-the-fly).
- It was sometimes difficult for the developer to decide whether or not particular comments or actions in the video clip should be noted on the form. The developer generally chose to take an inclusive approach, where he *did* make a note of these comments or actions. He commented, however, that he tended to wait a few minutes to do so, because the issues would most likely be mentioned or 'crop up again' on the video clip within a short time period. The types of things that the developer was unsure whether to note down included:

  o Usability issues where the task was not directly supported by the resource (e.g. finding out which sources the user had access to). This can be considered as either a usability issue, a functionality issue, both, or neither.
  o Usability issues where the user did not know how to go about a particular task, even when the task *is* possible. This was one of two instances where the developer chose not to make a note of any usability issues as he believed it would be 'counterfactual to make usability judgements' in this case. The developer suggests that, in this case, it is only possible to make a note that the user is 'a long way away from where they should be.'
  o Inferred user misconceptions or inferred usability issues (e.g. the user initially thinking that the Table of Contents menu bar was a results list, the user not using the Table of Contents menu bar often for navigation purposes, but frequently chaining through mentioned documents in the current full-text, which is not supported other than the provision of standard hyperlinks).
  o Small-scale usability issues (e.g. the user not knowing it is possible to stretch the Table of Contents bar/the bar not sizing properly to fit the width of the text).
  o Issues that were unlikely to be addressable (for whatever reason). For example the lack of hyperlinks to some case reports that the participant believed would be available in a leading competitor's resource if they were not available in the current resource. This was the second case where the developer chose not to make a note of the usability issue.
  o It was sometimes difficult for the developer to pinpoint exactly when in a video clip the user actions were taken or comments were made, as sometimes usability problems built up over time or where 'themes' of problems rather than individual, isolated problems. The developer likened this to not knowing whether to highlight a few individual sentences in a transcript of a large part of the text to illustrate a particular usability issue. Therefore it was not always easy to associate user actions or comments with underlying usability issues.

The particularly easy aspects of conducting the usability evaluation, as highlighted by the developer, were:

- Identifying the screen within the electronic resource that a usability issue related to (or most closely related to). According to the developer, this was 'easy 90% of the time.' The developer was, at first, unsure what name to assign to each screen and suggested using the word 'page' as opposed to screen on the firm. The developer, then, commented that on occasions a usability issue appeared to either relate to more than one screen, or to no screens at all (i.e. when it was part of a wider usability issue). The developer gave the example of the pilot participant who conducted a search then, when only one result was returned and the full-text was immediately displayed, became disorientated and confused as to what the Table of Contents in the left-hand menu bar would allow him to do). The participant first assumed it

was a results list, then a list of cases mentioned in the currently displayed case report. This could be regarded as being related to the main search page, or the document view page, or both, or neither.

- Identifying the severity of usability issues and determining the amount of effort required to address usability issues. Although these judgements required some reflection, the developer commented that it was 'surprisingly easy' to make these as snap judgements on-the-fly. This contrasts with the findings of a study by John and Packer (1995), who asked a Computer Scientist to record his thoughts when learning and applying the Cognitive Walkthrough method. The Computer Scientist noted the difficulty involved with deciding on severity ratings for usability problems. Perhaps the relative ease with which the developer identified the severity of usability issues might be explained by the fact that he had several years of both software development *and HCI* experience. The developer also commented that it was often necessary to re-visit these judgements later as they were made as the video was playing back. The developer commented that sometimes, after watching more of the clip, a severity judgement may need to be revised. The developer chose to leave the judgement columns blank until he was sure what judgement to make (i.e. after he was sure he had seen as much of the video clip that related to the usability issue as was available).

Relating to the layout of the usability form, the developer made some suggestions for improvement. The developer suggested moving the 'approx. time in video clip' column so that it can be filled out after the usability issue has been noted. This was with the aim of 'changing the perception' of the role of the column by encouraging evaluators to make a note of any time reference/index that enables them to jump to the appropriate point in the video clip. The developer also suggested that removing the need to articulate usability issues in writing on-the-fly would 'increase the chances' of users of the method being able to successfully analyse user data when watching a live user performing tasks (as well as when reviewing pre-recorded user think-aloud data). The developer also claimed to find it 'tempting' to write down usability issues that he identified from watching the video clip, but that were not directly related to the user's actions or comments. This suggested the need for a 'personal observations' section on the form. All of these changes have been incorporated into the current version of the form in appendix 8.

The developer commented that both the greatest benefit and the biggest drawback of the IB usability method was its highly-structured nature. The structured usability form, according to the developer, helps to contribute to a 'thorough analysis.' However, the highly-structured nature of the method can also be regarded as constraining, and, according to the developer who took part in our pilot, might potentially make it difficult for evaluators to 'see the wood from the trees.' This, according to the developer, is because the method encourages 'local' or individual usability issues to be listed, but does not provide support for grouping these issues together globally (i.e. reflecting on the issues identified and summarising them so that user actions/comments that refer to the same usability issue are grouped together). The developer did, however, comment that the ability of the method to highlight individual usability issues was very useful and should not be removed.

In general, the developer was positive about both the IB functionality and usability methods. He used the IB usability method with hardly any instruction and found the method, overall, to be 'easy to use.' As the developer only analysed three short tasks using the method, it is not possible to comment on the volume or quality of the usability issues that he identified. Nor would it be useful to make any comparisons between the usability issues the developer identified and those identified by the researcher involved in developing and piloting the method. This is because, as demonstrated by Hertzum and Jacobsen (2001) and discussed briefly earlier, an 'evaluator effect' exists when conducting usability evaluations with different evaluators (different evaluators identified substantially different sets of usability issues from four videotaped usability sessions). It would also not be useful to make any comparisons between the severity and 'amount of effort required to address issue' judgements. Evidence for this claim can also be found in Hertzum and Jacobsen (2001), who asked four evaluators to select the ten 'most severe' from those they had identified but found that none of the 'severe' issues appeared on all four evaluators' list of top-ten issues.

Overall, we also deemed the developer pilot to be successful. It provided encouragement for using the behaviours as theoretical lenses on the *functionality* of electronic resources, which led to the further development

of the IB functionality method. The developer pilot also suggested promise for the IB usability method. The developer that the method was both useful and usable and that he would use the method in the future, by integrating it with the ways he currently assesses the usability of the electronic resources he develops. The developer pilot highlighted a number of small changes that could be made to the usability form and suggested three important changes in the analysis procedure:

1. Emphasising in the guidance materials that an inclusive approach towards noting user actions/comments and usability issues should be taken (i.e. that if in doubt about whether to note down something from the user think-aloud session, it is preferable to note it down and cross out the notes afterwards if no longer required rather than avoid noting it down).
2. Allowing users of the method to fill out the 'usability issues identified,' 'severity of issue' and 'amount of effort required to address issue' columns of the form either *whilst* or *after* reviewing the think-aloud session in order to allow for the reflective nature of this information.
3. Allowing users of the method to record own usability observations from watching the video clip that were not directly related to the user's actions or comments.

6.3.6   Summary of the development and early testing of the IB methods

In this section, we have discussed the development and early testing of both the IB functionality and IB usability methods. We believe that this is important as there are very few documented accounts of how evaluation methods have been developed and, in particular, how they have evolved over time. The process of developing and testing the IB methods began by identifying a need - the need for a method or methods to allow evaluators to assess the functionality and usability of electronic resources with regard to the information behaviours identified in our empirical study of lawyers' information behaviour. Our initial attempts to address the need were less than successful, however a change from a 'theory driven' to an 'experientially driven' development approach yielded useful benefits. A process of iterative application and refinement of both methods followed. Next we conducted a series of pilot user think-aloud sessions in order to determine the tasks and procedural details to be used as part of an IB usability evaluation and a developer pilot analysis session in order to determine the usability, usefulness and potential for future use of the IB usability method and suggest ways of improving the data analysis procedure. The session suggested useful ways for improving the IB usability method. The developer pilot also suggested the potential utility of using information behaviours as lenses for evaluating the *functionality* of electronic resources.

6.4   Description of the current version of the Information Behaviour (IB) methods

We now briefly describe the functionality and usability IB methods. Detailed guidance for conducting a functionality or usability evaluation is presented in appendix 2. The IB methods are also supported by a number of resources, including the behaviour definition and examples document provided in appendix 3. This document lists the definition of each of the information behaviours at the core of the method, along with illustrative ways that electronic resources might support these behaviours at each applicable level. The document also contains screenshots illustrating some of the examples using Justis, an electronic legal resource. Furthermore, the IB methods are supported by a set of forms that are used to record the output of the evaluation and the user think-aloud information sheet and list of behaviour-focused tasks that evaluators can choose to present to users in order to obtain think-aloud data (which they will, in turn, identify usability issues from). The forms for recording the output from an IB functionality evaluation are presented in appendices 6 and 7. The form for recording the output from an IB usability evaluation is presented in appendix 8. The think-aloud information sheet and task list to be used to structure user think-aloud sessions as part of an IB usability evaluation are presented in appendix 5.

The information behaviours that we have discussed previously feed into the functionality and usability IB methods in different ways. They are used in the *functionality* IB method as a framework for assessing the functionality provided by electronic resources. An IB functionality evaluation involves users of the method discussing whether and in what ways an electronic resource currently supports the information behaviours, at a number of different levels (i.e. whether and in what ways the resource supports users in working with the *resource* itself, *sources* within the resource, individual *documents*, *content* within a document and *search queries/results*). Where relevant, an IB functionality evaluation also involves evaluators discussing whether it might be possible to support user behaviours/levels in additional ways and considering whether it is still necessary to support all of the behaviours/levels that the resource currently supports.

The information behaviours are used in the *usability* IB method as the foundation of think-aloud tasks that are set to intended or actual users of the electronic resource. In the usability method, evaluators set a number of behaviour-focused tasks to users, who are asked to perform the tasks whilst thinking aloud. This involves the users verbalising their thoughts, actions and feelings whilst performing the tasks using the specified electronic resource (just as in a conventional HCI think-aloud session). The evaluators then identify usability issues from the resultant think-aloud data and make summary judgements on how severe they consider the issues to be (i.e. whether or not they need immediate attention) and the amount of effort they consider to be required to address the issues. The process of identifying and making summary judgements on usability issues that are identified from think-aloud data is not unique to the IB usability method. Indeed, this is often a standard part of user testing. It is the theory-based task-setting element that makes the IB usability method unique, as the behaviour-focused tasks are used to frame the think-aloud session – providing a structure to the tasks that aims to encourage the display of a broad range of information behaviours (or particular behaviours/levels of interest).

The IB methods can be used by anyone with an interest or stake in an electronic resource. However, we recommend the methods are used by people with a basic grounding in usability evaluation and without a strong bias towards the existing design of the electronic resource under evaluation. This means that whilst the method can, in theory, be used by the developers of a particular resource themselves, this is not advisable as it is likely to be difficult to avoid attachment to particular system functionality or other related issues (such as power relationships within the firm or the evaluation team itself). This is the case with many HCI evaluation methods.

The process of planning and conducting both an IB functionality and usability evaluation is framed by Blandford, Adams et al.'s (2008) PRET A Rapporter Framework (PRETAR) – a framework for structuring user-centred evaluation studies, including evaluation studies of information retrieval systems. The broad process involves:

1. Defining the **purpose and boundaries** of the evaluation. More specifically this involves: deciding which electronic resource to evaluate, deciding which type of evaluation to carry out (i.e. a functionality evaluation, a usability evaluation or both), deciding which parts of the resource to evaluate and deciding which behaviours to evaluate the resource in relation to. The size of the resource is likely to influence the decision on which parts to evaluate (it may not be possible to evaluate large resources in their entirety). The decision on which behaviours to evaluate the resource in relation to is likely to be influenced by the *focus* of the evaluation. For example, it may not be within the scope of the resource to support wider information-use behaviours such as analysing, synthesising, recording, collating, editing and distributing information. Therefore, in such cases, certain behaviours may be excluded from the evaluation.
2. Deciding on the **practicalities** of the evaluation. This involves considering issues such as when in the design and evaluation cycle the resource should be evaluated, who should participate in the evaluation, how much time should be devoted to the evaluation and how the evaluation should be recorded.
3. Considering the **ethical issues** surrounding the evaluation. This involves considering issues surrounding keeping participant data as anonymous as possible and respecting participants' confidentiality and privacy. It may also involve considering how participant data will be used and, if applicable, disseminated.
4. **Conducting the evaluation itself** and **recording the output**. This stage is discussed in detail below for both the functionality and usability methods.

5. **Communicating the findings** from the evaluation. There are many varied ways of communicating the findings of an IB evaluation, ranging from using them as the basis of formal reports to using them as a basis of informal presentations or discussions.

6.4.1 Conducting an IB functionality evaluation

An IB functionality evaluation involves deciding whether the information behaviours, at each applicable level, are currently supported by the electronic resource and then:

- For levels of a behaviour that the resource currently supports, determining in which way(s) the resource currently supports and in which additional ways it might support the behaviour at this level.
- For levels of a behaviour that the resource does not currently support, determining in which way(s) the resource might support the behaviour at this level.
- Considering whether there are any behaviours/levels that it may no longer be necessary to support and, for these behaviours/levels, discussing the potential arguments for and against ceasing support.
- Considering any *current ways* the resource supports any of the behaviours/levels that may no longer be necessary and, for these ways, discussing the potential arguments for and against ceasing support.

This functionality evaluation can be supported by reference to or exploration of a running version of the resource under evaluation (whether this be a full running version or a limited functionality prototype). Table 10 lists the behaviours and each applicable level at which resource functionality should be assessed.

	Resource level	Source level	Document level	Content level	Search query/result level
Accessing	Currently supported? (Y/N)	Currently supported? (Y/N)	Currently supported? (Y/N)		
Surveying		Currently supported? (Y/N)	Currently supported? (Y/N)		
Monitoring		Currently supported? (Y/N)	Currently supported? (Y/N)		
Searching		Currently supported? (Y/N)	Currently supported? (Y/N)	Currently supported? (Y/N)	
Browsing and extracting		Currently supported? (Y/N)	Currently supported? (Y/N)	Currently supported? (Y/N)	
Chaining		Currently supported? (Y/N)	Currently supported? (Y/N)		
Selecting, distinguishing and filtering		Currently supported? (Y/N)	Currently supported? (Y/N)		
Updating			Currently supported? (Y/N)		
History Tracking			Currently supported? (Y/N)		
Analysing				Currently supported? (Y/N)	
Recording		Currently supported? (Y/N)	Currently supported? (Y/N)	Currently supported? (Y/N)	Currently supported? (Y/N)
Collating			Currently supported? (Y/N)	Currently supported? (Y/N)	
Editing			Currently supported? (Y/N)		
Distributing			Currently supported? (Y/N)	Currently supported? (Y/N)	Currently supported? (Y/N)

Key: Core information-seeking behaviours ☐   Law-specific behaviours ▨   Information use behaviours ▨

Table 10: Behaviours and levels to be considered in an IB functionality evaluation.

It will not always be appropriate to determine ways that an electronic resource *might* support particular behaviours/levels or to consider whether particular behaviours, levels or ways of supporting them are necessary. This is especially the case when evaluating resources where users of the IB functionality method only have an

indirect stake in the resource (for example, when evaluating an electronic resource developed by a competitor firm). In such cases, a 'cut down' functionality evaluation can be conducted that involves exploring a running version of the resource to determine *whether it supports* the behaviours/levels in table 10 and if so, *in which ways* it currently supports them.

The procedure for conducting both a competitor and own resource IB functionality evaluation is discussed in detail in appendix 2 and the form for recording the detailed output of an IB functionality evaluation is presented in appendix 7.

6.4.2   Conducting an IB usability evaluation

An IB usability evaluation involves asking intended or actual users of an electronic resource to think aloud whilst using the resource to perform a number of tasks and analysing the resultant think-aloud data. This essentially involves conducting a conventional HCI think-aloud session, where users are asked to verbalise their thoughts, actions and feelings and their verbal protocols are recorded (with minimal researcher intervention). The general think-aloud process is documented in both HCI textbooks (such as Dumas and Redish, 1999) and in articles (such as Boren and Ramey, 2000). Both present guidelines for conducting think-aloud sessions. An IB usability evaluation differs from a conventional HCI think-aloud study in that it is *solely based on the behaviour-focused tasks* that users are asked to perform.

There are three sets of tasks that it is possible to ask users to perform in an IB usability evaluation, 'core,' 'recommended,' and 'custom.' In a *core* IB usability evaluation, participants are asked to perform two tasks related to 'accessing' the electronic resource under evaluation and then the broad (and only somewhat behaviour-focused) task of finding information currently or recently needed for their work. This task is the same as the one set to participants in our empirical study of lawyers' information behaviour and was included as part of the IB usability method as setting this task to the lawyers in our study resulted in rich naturalistic data. The broad nature of this task encourages (but does not guarantee) the display of a wide range of information behaviours. Although a 'core' IB usability evaluation is highly naturalistic it does not, however, encourage the display of *particular* behaviours. We recommend a 'core' IB usability evaluation as a 'quick and dirty' way of acquiring user think-aloud data that highlights usability issues.

The tasks set as part of a *recommended* evaluation are more behaviour-focused than the core tasks but not as naturalistic. They are based on *common ways* that the lawyers in our empirical study performed the full range of information behaviours that were identified. A 'recommended' IB usability evaluation is a way of acquiring rich and behaviour-focused think-aloud data, again with the potential to highlight usability issues. We advise conducting a 'recommended' evaluation in most cases. The tasks set as part of a *custom* evaluation are more prescriptive and less broad than those set in a core or recommended evaluation. This makes them highly focused on particular information behaviours, but only somewhat naturalistic. These tasks are based on *less common ways* that the lawyers in our study performed particular behaviours than the tasks in a 'recommended' evaluation. Some of the 'custom' tasks also encourage the display of behaviours at particular *levels* that were less commonly displayed by the lawyers in our study (i.e. the source, content and search query/result levels).

'Core' tasks should always feature in an IB usability evaluation. 'Recommended' tasks should feature alongside the core tasks unless financial and/or resource constraints make this impossible. 'Custom' tasks should be used to tailor IB usability evaluation where there is a need to focus on particular behaviours and/or levels at which the behaviours can operate. It is therefore possible to mix and match custom tasks that aim to encourage demonstration of particular behaviours/levels of interest.

In a 'core' IB usability evaluation, participants are asked to perform the three information-seeking tasks in table 11:

'Core' tasks in an IB usability evaluation:

1. Gain access to the electronic resource
2. Find out which parts or sources within the resource you have access to
3. Think of some information that you currently need or have recently needed to find for your work and demonstrate, using the electronic resource, how you might go about finding it.

Table 11: The three information-seeking tasks that think-aloud participants are asked to perform as part of a 'core' IB usability evaluation.

In a 'recommended' evaluation, participants are asked to perform the three core tasks listed above, *plus any of the tasks in table 12 that are currently supported by the electronic resource*. The recommended tasks are based on the full range of behaviours that were displayed by the lawyers in our empirical study (and were found to be commonly supported by our survey of the functionality of electronic legal resources). These tasks are, therefore, designed to encourage the demonstration of a broad range of behaviours. We use the word 'encourage' as it is not always possible to predict exactly how participants will carry out the tasks. Note that the recommended tasks only encourage demonstration of behaviour at the *document* level as this was the most common level at which they were displayed in our empirical study. Also note that some of the tasks in tables 12 and 13 have been customised for evaluating *legal* resources, however similar tasks can be set to evaluate resources from other domains.

'Recommended' tasks in an IB usability evaluation:

Gain an overview of an area by:
- Trying to gain a basic understanding of the law relating to a particular legal area (e.g. Breaches of contract).
- Trying to gain an appreciation of the importance of a certain legal journal author's role in a particular legal area.
- Trying to locate a legal journal article written by an author who has published many articles or many highly cited articles.

Gain a current or historical understanding of the importance of a document by:
- Trying to find out whether a particular case is still good law.
- Trying to find out what amendments have been made to a particular piece of legislation over a certain time period.
- Trying to find out whether a particular piece of legislation is currently in force.
- Trying to locate a historical version of a particular piece of legislation (i.e. a previous version that has since been amended).

Maintain awareness of developments in an area by:
- Trying to find out whether there have been any recent developments in a particular legal area (e.g. Discrimination law).
- Trying to set up an alert so that you can be informed every time new documents are added to the system that match particular search terms (e.g. when new documents that match the term 'discrimination' are added).
- Trying to set up an alert so that you can be informed every time there are new developments in a particular legal area.

Return to any one of the tasks where you found useful documents and:
- Determine which sections of a document that you have found are important to you.
- Keep a softcopy (downloaded or saved) record of a document that you have found.
- Keep a hardcopy (printed) record of a document that you have found.
- Download two documents into a single file (i.e. you should end up with one file saved on the computer that includes the text of two separate documents, e.g. two different legal journal articles or two different sections of a particular piece of legislation).
- Keep a softcopy or hardcopy record of *part* of a document that is important to you (e.g. print or download only certain parts of a case report).
- Distribute a document that you have found, by e-mail, to a fictitious colleague from within the electronic resource.
- Store a document on the server of the resource (i.e. save a copy of the document to a personalised area on the electronic resource itself, so you can access it again quicker in future).

Table 12: Tasks that think-aloud participants are asked to perform as part of a 'recommended' IB usability evaluation. In addition to these tasks, participants are also asked to perform the three core tasks in table 11.

In a 'custom' IB usability evaluation, the choice of tasks that users are asked to perform will vary depending on the focus of the evaluation. These tasks are designed to encourage *specific information behaviours* as identified in our empirical study. Table 13 lists possible custom tasks relating to 'chaining' behaviour (i.e. *"following chains of citations or other forms of referential connections between material"* - Ellis, 1989, p.179).

In a custom evaluation, tasks can also be set to encourage *behaviours at a particular leve*l. For example, three tasks relating to surveying at the *document* level are presented as part of the set of recommended tasks in table 12 (listed under the 'get an overview of an area by:' heading). However, it is possible to an additional (or alternative) task related to surveying *sources*, such as 'try to found out which sources contain information about a particular legal area.' A full list of custom tasks (for a range of different information behaviours and applicable levels) is presented in appendix 5 and a more detailed discussion about the difference between types of IB usability evaluation is provided in appendix 2.

Custom tasks related to 'chaining' behaviour:

Try to follow a hyperlink or other form of connection from:
- A legal case to a previous case or a particular piece of legislation mentioned in the case report.
- A piece of legislation to other pieces of legislation mentioned in the text of the Act.
- A legal journal or commentary article to a case or piece of legislation mentioned in the article.

Try to follow a hyperlink or other form of connection from a document to other documents that have been written since this document and have subsequently mentioned it. Specifically, try to:
- Find a particular case, then find out which more recent cases have mentioned it (if any).
- Find a particular piece of legislation, then find out which more recent pieces of legislation have mentioned it (if any).
- Find a particular legal journal article, then find out which more recent articles have mentioned it (if any).

Table 13: Custom tasks related to chaining behaviour.

In an IB usability evaluation, users are audio and screen recorded as they think-aloud whilst performing the information tasks and the resultant think-aloud data is reviewed in order to identify usability issues. Whilst reviewing the think-aloud data, the evaluator(s) keep a record of any user actions, user comments, or personal observations that might suggest a usability issue and makes a note of what they believe to be the underlying usability issues identified from the actions/ comments/observations. The evaluator(s) also record details of the screen(s)/page(s)/part(s) of the resource that the actions, comments, or observations relate to and makes judgements on the severity of and the amount of effort required to address the usability issue. The procedure for conducting an IB usability evaluation is discussed in more detail in appendix 2 and the form for recording the detailed output of an IB usability evaluation is presented in appendix 8.

6.5 Worked examples of carrying out an IB functionality and usability evaluation

We now present two short worked examples to illustrate the use of the IB functionality and usability methods. The examples involve evaluating LexisNexis Butterworths (LNB), an electronic resource widely used by legal professionals worldwide. The *purpose* of both the functionality and usability evaluations was to illustrate the use of the methods to an audience likely to be unfamiliar with them. As this resource has many customisable options (e.g. dedicated search screens for various legal practice areas), one *boundary* that we set was to evaluate only those parts of the system available as default in the publicly available, uncustomised, academic version of the resource, which we accessed in late 2007. As we did not have a direct stake in LexisNexis Butterworths, there was no need to consider *practicalities* such as the amount of time to devote to the evaluations, or the place of the evaluations in the product's development cycle (we only had access to the publicly available version of the

resource). We also had access to a single evaluator and to several audio and screen recordings of lawyers using LexisNexis Butterworths to perform a range of the tasks that feature in an IB usability evaluation. We chose to use a small clip of a Trainee Solicitor performing 'updating' behaviour as this clip illustrated a number of potential usability issues. The *ethical issues* surrounding the evaluation mainly involved gaining permission from LexisNexis Butterworths to evaluate their resource and to report the evaluation, complete with screenshots, in this thesis. It was also important to obtain permission from the Trainee Solicitor to print the data arising from his think-aloud session. This data is presented in appendix 9 (and summarised on the IB usability form in appendix 8). We *communicate the findings* from the IB functionality and usability evaluations below.

6.5.1   Example IB functionality evaluation of an electronic legal resource

To illustrate part of an IB functionality evaluation, we evaluate LNB in relation to 'browsing and extracting' behaviours. Browsing involves *"semi-directed searching for <u>sources</u>, <u>documents</u> or <u>content</u>."* Extracting can often work hand-in-hand with browsing and involves *"systematically working through a particular resource to identify <u>sources</u> of interest, a particular source to identify <u>documents</u> of interest and/or a particular document to identify <u>content</u> of interest."* As well as from the definitions, it can be noted in table 10 that browsing and extracting can operate at three levels: the source, document and content levels. As a reminder, browsing and extracting behaviours are intended to be analysed together. We now evaluate LNB in relation to browsing and extracting behaviours, for each of these three levels.

An IB functionality evaluation usually involves asking ourselves *in which ways* an electronic resource supports a certain behaviour or behaviours at a particular level and in which *additional* ways *might* it support them. It also usually involves considering whether there are any behaviours, levels or ways of supporting them that *may no longer be necessary.* We do not discuss potential functionality reduction and arguments for and against support in this example as we feel we would only be able to do so if we had a direct stake in the resource (and were armed with knowledge about the use of the various parts of the resource's functionality). We do, however, discuss *additional ways* in which the resource *might* support each behaviour/level (even though this is not normally necessary if the evaluator does not have a direct stake in the resource under evaluation). However, we do so purely to illustrate how the IB functionality method can lead to suggestions for additional functionality support. See appendix 7 for a form that outlines the questions to be considered in different circumstances as part of an IB functionality evaluation. This form can also be used to help record the output of the evaluation. We now turn to ask in which ways the resource *currently supports* and in which *additional ways* it might support browsing and extracting behaviours at each of the *source*, *document* and *content* levels.

In which way(s) does the resource *currently support* browsing and extracting at the <u>source</u> level? It is possible to browse to locate (and extract) sources in LNB by using the dedicated 'browse' functionality provided by the resource. As illustrated by the radio buttons labelled 'a' in figure 10, it is possible to browse for and extract sources in a number of ways. The first two are by publication type (this is illustrated in the part of figure 10 labelled 'b', where the publication type 'legal journals' has been selected and the relevant sources displayed) and by area of law (i.e. the legal areas that the source is deemed to cover). The other two allow browsing and extracting of sources that contain business and news-related and industry related materials. As illustrated by the drop-down combo boxes in figure 10 (labelled 'c'), it is also possible to filter the list of sources to include only those that fit specified criteria (e.g. sources covering only UK law) and then browse the filtered list of sources.

Figure 10: Browsable list of sources in July 2007 public release version of LexisNexis Butterworths, listed by 'publication type.'

In which *additional ways* might the resource support browsing and extracting at the _source_ level?
Although LNB provides comprehensive support for browsing and extracting at the source level, the resource might support this behaviour/level further by allowing legal professionals to browse by other aspects of source meta-data that they deem to be important. For example, by when the first or most recent document in the source was published (i.e. browsing by source coverage dates).

In which way(s) does the resource *currently support* browsing and extracting at the _document_ level? As well as supporting browsing and extracting at the source level, LNB also provides comprehensive support for browsing sources in order to extract *documents*. Clicking on the 'browse' hyperlink next to any of the listed sources in figure 10 allows users to view documents contained within the source (albeit after drilling down several more levels). Another way that LNB supports browsing and extracting at the document level is illustrated by the 'browse TOC/index' sidebar in figure 11 (labelled 'a'). This sidebar, presented adjacent to the full-text of documents in LNB, facilitates browsing other documents within the currently selected source (which in figure 11 is the 'Journal of International Economic Law'). This is achieved by clicking on one of the article titles listed in the expanded tree. It is also possible to use the 'TOC/index' sidebar to browse for and extract journal articles in other issues of the selected journal series.

157

Journals > WTO for Trade and Development Pos...

Search | Practice Areas | Sources | My Research | Results ?

Previous Next

Hide TOC / Index

View Full Text

1 of 1   Search Source

Browse TOC / Index   a

☐ Issue 3, September
☐ Issue 2, June
  ☐ Articles
    ☐ Negotiations on Safeguards and Subsidies in Service
    ☐ Smooth as Silk? A First Look at the Post MFA Textile
    ☐ The WTO, Science, and the Environment: Moving To
    ☐ Independence under Fire: Extra-legal Pressures and
    ☐ When do the Activities of Private Parties Trigger WT
    ☐ The EU—ACP Economic Partnership Agreements and
    ☐ Compliance by WTO Members with Adverse WTO Dis
  ☐ Reviews
☐ Issue 1, March
☐ 2006, Volume 9
☐ 2005, Volume 8
☐ 2004, Volume 7
☐ 2003, Volume 6
☐ 2002, Volume 5
☐ 2001, Volume 4
☐ 2000, Volume 3
☐ 1999, Volume 2
☐ 1998, Volume 1

WTO for Trade and Development Post-Do...

Journal of International Economic Law
JIEL 2007 10 (553)
1 September 2007

WTO for Trade and Development Post-Doha

Seung Wha Chang

Professor of Law at Seoul National University. Currently Visiting Professor of Law at Harvard University (2007) and Yale University (2008). E-mail: changsw@snu.ac.kr.
© Oxford University Press 2007

Abstract The WTO Members' negotiations under the Doha mandate on special and differential treatment ('SDT') and development issues have made little progress. The gap between developed countries and developing countries in this regard seems too wide to be bridged. This gap originates from a fundamental difference in their basic stances on cross-cutting issues. In principle, without prejudice to currently available SDT under individual WTO Agreements, developing countries should be given flexibilities in implementing WTO rules, which may result in the rebalancing of rights and obligations of WTO Members but only when policy measures at issue can contribute to particular developing countries' development needs and no alternative less-trade restrictive measures are reasonably available. After reviewing WTO Members' discussions since the adoption of the Doha Ministerial Declaration, this article suggests a 'measure-specific ex ante approach' for a workable solution to bridging the gap between developed countries and developing countries on SDT and development issues.

Introduction

Currently, developing countries account for more than three-fourths of the WTO membership. With more than a dozen years of WTO history, developing country Members of the WTO have expressed their discomfort with the end-results of the implementation of the Uruguay Round and subsequent negotiations under the WTO. It is generally understood that most developing countries, in particular, the Least Developed Countries ('LDCs') have been unable to harvest benefits from the

Internet

Figure 11: Table of Contents/Index sidebar in July 2007 public release version of LexisNexis Butterworths, which allows users to browse documents from the current source.

In which *additional ways* might the resource support browsing and extracting at the <u>document</u> level? Despite strong coverage for browsing and extracting at the document level, there are additional ways that these behaviours might be supported at this level. Figure 12 illustrates the 'keywords' listed at the top of a legal case in LNB. The resource currently facilitates browsing to other documents listed as part of the keywords (in the case of figure 12, the 'Employment Rights Act 1996'). However, there is also scope to allow browsing and extracting by keyword and other items of document meta-data. For example, it may be possible to allow users to click on a keyword and be presented with other documents that are related to the keyword (whether those documents be cases, or other types of document such as legislation or journal articles). There may also be scope to provide, for legal case reports, the facility to browse to all other cases involving a particular party in the currently displayed case (for example, the 'Secretary of State for Trade and Industry' in figure 12), all other cases with the same judge presiding or all other cases with the same counsel. Similarly, for legal journal articles, there may be scope to find all other articles written by the same author as the currently displayed article.

Family - All subscribed Family so... ☐ > ("unfair dismissal" age) (3) > Document (3/3)

Search | Practice Areas | Sources | My Research | Results ?

View Full Text   Narrow Search [ ] Go

3 of 3   Next Steps Edit Search Go

☐ Secretary of State for Trade and Indu...   Create Alert

Secretary of State for Trade and Industry v Rutherford and others

Discrimination – Sex discrimination – Employment – Indirect discrimination – Disparate impact – Justification – Tribunal finding provisions indirectly discriminatory – Correctness of decision – Employment Rights Act 1996 – Treaty of Rome, art 141 – Council Directive 76/207EEC.

EAT/1029/02/SM, (Transcript)

Figure 12: Keywords listed at the top of a legal case report in July 2007 public release version of LexisNexis Butterworths.

In which way(s) does the resource *currently support* browsing and extracting at the <u>content</u> level? As with the other levels, LNB also provides a considerable amount of support for browsing and extracting at the *content* level. As illustrated in figure 13, users' search terms are automatically highlighted in bold in the full-text of a document and it is possible to cycle through each occurrence of the search terms by clicking on the small arrows in the bottom-right-hand-corner of the screen (labelled 'a').
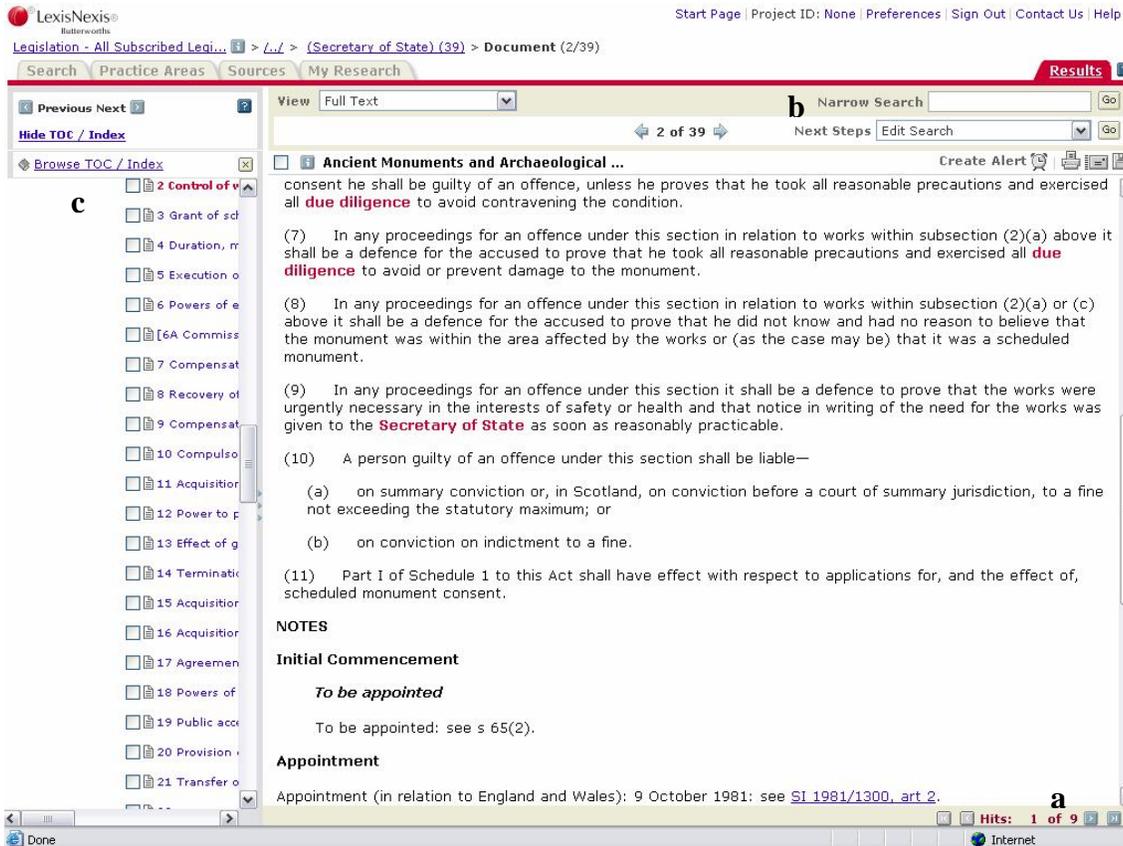
Figure 13: Document in July 2007 public release version of LexisNexis Butterworths with search terms automatically highlighted in bold.

In which *additional ways* might the resource support browsing and extracting at the <u>content</u> level? Also as illustrated in figure 13, it is possible to enter additional search terms in the 'narrow search' field (labelled 'b'). This serves to highlight additional words or phrases in the document text. There is scope to provide similar functionality for highlighting particular words or phrases in the current document, without having to refine the original search. Similarly figure 13 illustrates the TOC/index sidebar (labelled 'c'), which allows users to 'jump' to particular documents within the current source (in the case of figure 13, users can jump to other Sections of the currently displayed Act of Parliament). There is scope, however, to extend this functionality to also allow users to 'jump' to particular *section headings within the current document* (for example, to jump to the 'notes' section heading in figure 13).

As we have previously highlighted, the suggestions for increasing the range of functionality supported by LexisNexis Butterworths are to be treated as illustrative as they are not grounded in marketing or usability data. Therefore it follows that developers at LexisNexis Butterworths, armed with such data, may either decide that the functionality serves a useful purpose and/or make the resource more usable, or adds unnecessary complexity to the resource. It also follows that increasing the range of functionality supported will not necessarily lead to a holistically more usable resource and care must be taken to ensure that a balance is struck between the functionality supported by the resource (and the resultant complexity arising from this support) and the overall

usability of the resource. This provides a case for conducting both functionality and usability evaluations of resources. In the next section, we shift our focus from functionality to usability and present an example of an IB usability evaluation of LexisNexis Butterworths.

6.5.2 Example **IB usability evaluation of an electronic legal resource**

We now present another example, this time of analysing think-aloud data as part of an IB *usability* evaluation. Note that although the process of conducting an IB usability evaluation also involves setting tasks to users and collecting the data, we only discuss the data *analysis* process in this example. The think-aloud data is based on a Trainee Solicitor, who worked for a small London law firm, using the same electronic resource as in the previous example (LexisNexis Butterworths). The Trainee was asked to use the resource to perform one of the tasks that is set as part of a 'recommended' IB usability evaluation related to 'updating' behaviour. This task was to try and 'find out if a particular case is still good law.' In order to complete the updating task, the Trainee chose to examine a case called 'White vs. White.'

The Trainee had only used LNB a few times before, as part of a Legal Practice Course. The trainee's 'novice' status proved useful in yielding usability data that we hypothesise might not have been obtained from a more experienced user.

In this example, we provide a description of how the Trainee Solicitor went about the updating task. This serves as a summary of the written transcript of the Trainee performing the task, presented in appendix 9. Throughout our description, we also discuss the usability issues arising from the actions and comments made by the Trainee whilst performing the task. The usability data is summarised on the IB usability form presented in appendix 8. The usability data on the form can be cross-referenced to highlighted sections of the transcript. Each highlighted section of the transcript in appendix 9 is marked with a letter, which corresponds to a row (i.e. a usability issue) on the usability form in appendix 8.

It is important to note that there is an element of subjectivity involved in identifying usability issues related to complex electronic resources such as LNB. Therefore different evaluators are likely to identify different usability issues (and potentially differ on their subjective ratings of the severity and ease of addressing each issue). This 'evaluator effect' is evidenced by Hertzum and Jacobsen (2001), who asked four evaluators to select the ten 'most severe' from those they had identified but found that none of the 'severe' issues appeared on all four evaluators' list of top-ten issues. It therefore follows that this example evaluation should not be regarded as the sole or most authoritative interpretation of the Trainee Solicitor's think-aloud data, but as *one possible interpretation* of the data.

Description of how the Trainee Solicitor went about performing the updating task and discussion of the resultant usability data

The Trainee Solicitor began by stating that he wanted to find out whether an ancillary relief case, fought by the parties White and White, was still good law. He clicked on the 'cases' tab, and mentioned that he was about to type the party names in the 'enter search terms' field, as opposed to the 'case name' field (which he soon realised would have been a more appropriate choice). This suggests a potential usability issue as the 'case search' page did not make it immediately clear to the user in which segmented field to enter their search query. The Trainee's comments and actions related to this issue are marked with the letter 'a' in the full transcript of the think-aloud session in appendix 9 and in table 14 (which presents only the comments and actions which suggest a usability issue, not the full transcript). Each usability issue is also summarised in a row on the usability form in appendix 8. As the Trainee noticed the correct field within a couple of seconds, this issue is not deemed to be severe. Also as the Trainee quickly noticed the error, this does not suggest that the segmented fields are poorly labelled. As it is not immediately clear what caused the Trainee difficulties, this issue would likely take a large amount of effort to address.

a	I'm gonna click on 'cases' and under 'enter search terms' I'm going to put [pauses]. No, I'm going to put it in the 'case name.'
b	[Participant types in 'White v W,' then pauses]. [Reads caption underneath 'case name' field]. 'To find Smith v Jones, enter Smith **and** Jones.' So I'm going to use the word 'and' rather than 'v.'
c	Here we go. 'Save search,' 'create alert,' 'search the source,' 'find related cases.' Hmm. I'll click on 'find related cases' and see what that brings up.
d	'How do I work with my search results?' I'll click on that as it might help me in some way. [Reads out sub-headings]. No, that doesn't help me. Let's search. [Searches within TotalHelp for 'current cases' and clicks on a result entitled 'results page' then soon after closes the help page]. I've closed that down because it just seems to be frustrating me.
e	Ok, the tutorial seems very slow and cumbersome, so [closes tutorial window]
f/g	[Enters 'white and white' into the 'get a specific document' part of the LexisNexis Butterworths homepage with 'find a case' selected in the combo box. Participant then presses the grey 'go' button next to the field he has filled out, but after a few seconds presses the main red 'search' button. [An information box is displayed, which says 'Please enter a search term'].
h	[Participant types 'White' in 'citation' field]. I'm not looking for a citation. [Participant then enters 'White and White' in the case name field].
i	View [clicks on 'view' combo box again and reads out options]. 'List,' 'expanded list,' 'White v White 2000.' [Selects White v White 2000]. Oops, I clicked on something.

Table 14: Comments and actions which suggest a usability issue, extracted from the full-transcript in appendix 9 of a Trainee Solicitor 'trying to find out whether a particular case is still good law.'

The Trainee continued by typing 'White v W' in the 'case name' segmented field, then read a caption underneath the field that instructs users on the correct syntax to use when searching for cases (i.e. 'and' instead of 'v'). Whilst this suggests that the required syntax was not made immediately clear (letter 'b' in appendix 8 and table 14), this issue is also non-severe. It could be addressed with little effort by allowing search fields to 'intelligently' accept a variety of syntax.

Next, the Trainee conducted the search and, upon receiving 1,700 results, filtered the results set to only display legal case reports and again to display only cases from the 'Family Court Reports' source. This brought up the required White vs. White case and the Trainee began by reading part of the case aloud. The Trainee then wondered 'how [he could] check to see it the case has been updated or not' and clicked first on the 'view' drop-down combo box, reading the options aloud, then on the 'next steps' combo box (pictured in figure 14). Again, he chose to read the options aloud and selected the 'find related cases' option to 'see what that brings up.' This suggests the actions facilitated by the 'next steps' combo may not be as transparent as they could be, an issue that we deem to be quite severe and moderately difficult to address. This relates to letter 'c' in appendix 8 and table 14.
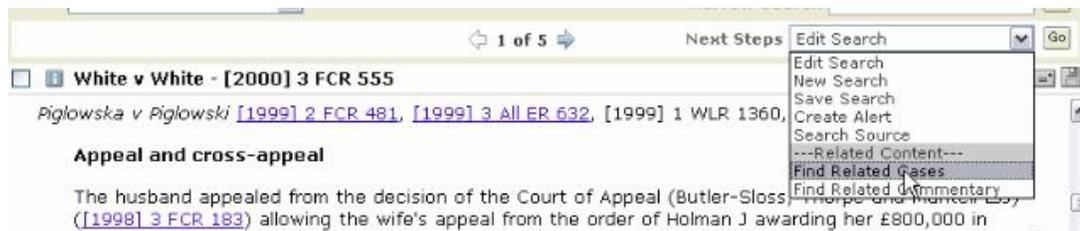
Figure 14: The 'Next Steps' drop-down combo box in July 2007 public release version of LexisNexis Butterworths.

The Trainee was then presented with a list of 'related cases,' which he scrolled through. He then selected one of the cases in the list, Wood vs. Rost and noticed a reference and hyperlink to the White vs. White case, which he clicked on. Feeling 'a bit stuck,' he clicked on 'help' but did not find anything that could assist him to perform the task at hand and after a few minutes closed the help page as it '[seemed] to be frustrating' him. He then tried viewing a tutorial, but after a few minutes closed the tutorial too, commenting that it 'seems very slow and

cumbersome.' These are both usability issues that we deem to be quite severe, if only because they caused the Trainee frustration. We deem them to be moderately difficult to address, as although updating tasks are important for lawyers, it is a considerable challenge for help systems to provide the required assistance in a useful format without patronising or frustrating users. These issues relate to letters 'd' and 'e' in appendix 8 and table 14.

After closing the tutorial screen, the Trainee decided to 'try searching again' and selected 'find a case' from the 'get a specific document' combo box (see figure 15). However, the Trainee did not notice the 'go' submit button turn grey after he submitted the search and, possibly assuming he had clicked on the wrong submit button, clicked on a 'search' button that was not related to the current search field. Figure 15 illustrates the Trainee clicking on the unrelated 'search' button and receiving a popup error message. This suggests that the relationship between search fields and their associated search buttons may not be as clear as it could be. This usability issue (letter 'f') is, in our opinion, very severe and would require moderate effort to address (perhaps by hiding or greying-out submit buttons that are not relevant to the search fields currently in use). The Trainee not noticing the 'go' button turn grey can be regarded as a separate usability issue (letter 'g'), associated with the resource providing potentially unclear feedback. We deem this issue to be quite severe and suggest it would require moderate effort to address.

Figure 15: Popup error message displayed by July 2007 public release version of LexisNexis Butterworths when a search is submitted without anything in the 'enter search terms' field.

The Trainee then dismissed the popup box, returned to the 'case search' page, and briefly entered a party name in the 'citation' field, before correcting the error (which was identical to an error made towards the beginning of the task – letter 'h' in appendix 8 and table 14). Like related usability issue 'a,' we deem this issue to be non-severe, but difficult to address. Next the Trainee ticked 'House of Lords' in the 'court' selection box and, after conducting the search, filtered the results in the same way as earlier in the task, briefly pausing to tick the tick-box beside the 'White vs. White' case and to press 'view tagged' to check his hypothesis that it would 'be the same as if [he] just clicked on it.' The Trainee's hypothesis was confirmed when the full-text of the White vs. White case was displayed. The Trainee then read the 'next steps' combo box options aloud for the second time and then selected 'White v. White 2000' from the 'view' combo, which re-loaded the full-text of the case. This seemed to confuse the Trainee, who stated 'oops, I clicked on something.' This suggests that the effect of selecting the current document from the 'view' combo may be unclear. This is a usability issue (letter 'i') that we deem to be non-severe and moderately difficult to address.

6.6 Benefits and limitations of using the IB methods

We now discuss the benefits and limitations of the IB functionality and usability methods. These benefits and limitations have not been empirically tested as they are grounded in specific features of the methods. Our discussion is followed by a summary of the methods.

6.6.1 Benefits of using the IB functionality and usability methods

We believe the key benefits of using the IB functionality and usability methods are related to those aspects of the methods that make them novel. For example, because the methods are based on *empirically observed* information behaviours, they allow users of the methods to take a truly user-centred as opposed to system-centred focus on improving usability and providing an appropriate range of functionality. Similarly because the methods *specialise in facilitating the evaluation of electronic resources* and not other types of interactive system, they allow evaluators to focus on functionality and usability-related issues that are specific to or particularly problematic with these types of systems. For example, issues to do with 'accessing' behaviour are particularly important for developers of electronic resources to consider, as these resources are not always accessed in a straightforward manner like an e-commerce or other website might be. Similarly there are few other types of interactive system where it might be important to consider providing the functionality to restrict access at a variety of levels (e.g. access to the resource itself, access to particular sources within the resource, access to particular documents). One should note, however, that addressing access issues may be not necessarily be within developers' control (see Bates, 2002; Blandford, Keith et al., 2007).

Other benefits, which we have already discussed, include the *extensibility* and *flexibility* of the methods. To recap, the IB methods can be tailored in two main ways. Firstly particular behaviours and levels can be included or excluded from both functionality and usability evaluations depending on the current domain or focus. Secondly, evaluations can be conducted on entire resources or particular parts of a resource. Modularity, scalability, flexibility and customisability are all highlighted as important features of methods by Garzotto and Perrone (2007). We have also previously discussed the *highly structured* nature of the methods. This feature of the methods, combined with their extensibility and flexibility allows users to strike their own balance between having the necessary structure and guidance to perform a successful evaluation with the flexibility to tailor the evaluations to meet their needs. There is also potential for customisation to particular domains to, over time, enrich the behavioural theory base at the heart of the methods. We are unaware of any existing evaluation methods with similar enrichment potential.

There are further benefits still, associated with both of the separate methods. Because an IB functionality evaluation allows users of the method to consider ways that particular behaviours *might be supported* at certain levels, the method supports the making of novel as well as incremental design improvements, within certain limits. In addition, because an IB functionality evaluation allows users of the method to consider whether it is *necessary* to continue to support particular behaviours/levels, it does not propagate the (often incorrect) assumption that supporting a greater range of user behaviours will result in a better, more usable system. Because the behaviour-focused tasks used in the IB usability method were generated as a result of empirical data (and by surveying how a range of electronic legal resources support information behaviours at various applicable levels) the tasks (and therefore the usability method in general) are easily customisable and updatable. In addition, new ways of supporting particular behaviours/levels that are discussed as part of functionality evaluations can feed in to the tasks provided in usability evaluations. Future empirical work or surveys of electronic resources, perhaps in another domain, also has the potential to update and customise the methods.

6.6.2 Limitations and scope of the IB functionality and usability methods

Although there are a number of benefits associated with both the functionality and usability IB methods, there are also a number of potential limitations associated with them. Many of these limitations are related to the scope of the methods.

Whilst some solutions to usability observations may become apparent simply by noting that an issue exists, the scope of the IB usability method is restricted to *identifying* as opposed to addressing usability issues and the scope of the IB functionality method is restricted to examining the *range of support* provided for particular information behaviours in an electronic resource (as opposed to adding or removing functionality). Although there is an ongoing debate with regard to what extent evaluation should directly inform design, discussed in Wixon (2003), we argue that the IB methods would become too complex and difficult to use if they were intended to help users of the method decide exactly how to address usability observations or make *binding* decisions on whether, and if so in which ways, to support a particular behaviour at the interface level. Although we do not believe that evaluation and design efforts should be separated completely, our IB methods are pitched primarily as evaluation rather than as design tools. That is not to say that both methods cannot provide useful feed-in to future design discussions. For example, the usability issues identified and the broad ways in which an electronic resource *might* support particular behaviours/levels might be used as a basis for future design discussions. Indeed the outputs of both the IB functionality and usability methods can be used informally to support design discussions, or more formally in conjunction with design tools such QOC (see MacLean et al., 1991) to help stakeholders make interface-level design decisions.

The IB functionality method aims to help stakeholders in an electronic resource consider ways that information behaviours can be supported, at particular levels. However a challenge for evaluation methods is to promote use of the method at a consistent and appropriate level of abstraction, so as to ensure that the output of the method is useful to all who use it, not just those who instinctively use the method at the level of abstraction that the creators intended. This challenge is discussed further in Blandford, Keith and Fields (2006) and Blandford, Gow et al. (2007). The supporting documentation for the functionality IB method (see appendix 3) includes *illustrative ways* that each behaviour or set of behaviours might be supported at particular levels. For example, one way of supporting the behaviour of 'recording' at the document and content level (i.e. making a record of documents and content) is by providing the facility to download entire documents or parts of documents. These examples aim to give an indication of the intended level of abstraction that the IB functionality method should be used at when evaluators consider new ways of supporting user behaviours. This can also be regarded as an important scope constraint for the functionality method, as the examples are neither so general so as to provide little guidance in how a behaviour might be supported by an electronic resource, nor so detailed so as to be prescriptive of the exact form that the support might take at the interface level. Bellotti et al. (1995) explain that providing such examples in evaluation methods *"provides something concrete to ground understanding and to reflect upon, but is not meant to dictate design – having understood the principle at stake, a designer can tailor an example for a specific solution"* (p. 12).

Another challenge for evaluation methods is to support the making of *novel* as well as *incremental* design improvements. The IB usability method indirectly supports users of the method in making incremental design improvements by facilitating the identification of usability issues associated with the current system (which, if addressed effectively, should lead to an incremental improvement in the usability of the electronic resource concerned). The IB functionality method also indirectly supports users in making novel design improvements by allowing them to consider *additional ways* in which the resource might support information behaviours at particular levels. However, it only provides support for revolutionary design improvements within those strict limits. It does not directly take into account the fact that new ways of performing information tasks might be supported by future electronic resources that change the way that users perform behaviours with these resources (or even change the fundamental behaviours that they perform). Instead, it relies on the assumption that the empirical basis for using the behaviours as part of the method will remain valid (or at least that subsequent research will help to maintain the validity of the method). Therefore despite evidence from our study of lawyers'

information behaviour and from Meho and Tibbo's (2003) study of Social Scientists that similar behaviours are performed nowadays to those originally identified almost twenty years ago by David Ellis and his colleagues, we cannot be absolutely certain that information behaviours observed in a particular domain will remain the same over time. In short, the IB functionality method supports novel as well as incremental design, but only to a limited extent. The heavy research investment required to ensure the empirical basis of the methods remains valid is also an important limitation of both the IB functionality and usability methods.

Finally, the IB functionality method allows users of the method to consider new ways in which user behaviour might be supported. However, just as introducing new ways of supporting a certain behaviour at a particular level has the potential to lead to improved usability overall (e.g. by providing support for a behaviour that has so far been neglected by the resource), it also has the potential to have a *detrimental* impact on the overall usability of the resource (e.g. by helping to create a feature-overloaded resource that remains too complicated to use). This is why the functionality method does not assume that support for a greater range of behaviours and levels will necessarily lead to a more usable electronic resource, an argument also made by Mack and Nielsen (1994) who point out that several chapters in their handbook of Usability evaluation Methods *"refer to the need for evaluations to focus on the usefulness of interface function, and not simply the usability of the interface, as an implementation of that function"* (p. 6). The IB usability method allows evaluators to make usability observations associated with particular behaviours and levels. However, this also has the potential to have a detrimental impact on the overall usability of the resource. For example a design intervention aimed at addressing some of the usability observations identified might unknowingly introduce new usability observations, or the design intervention may seem to be an improvement on paper but not improve the usability of the resource in practice. This is why both the functionality and usability methods can be used at various points and at multiple times during the design process (and we encourage iterative use of the methods to complement an iterative design process).

6.6.3   Summary of the IB methods

The IB methods allow both direct and indirect stakeholders in an electronic resource to evaluate its functionality and usability, by using a number of empirically identified information behaviours and levels at which these behaviours can be performed as springboards. These behaviours are used to consider whether and in which ways a particular resource currently or might support each behaviour, at each applicable level, aimed at ensuring that the resource supports an appropriate range of functionality. It is also possible to use these behaviours as 'lenses' on the usability of the resource – by having intended or actual users step through the resource and attempting to perform tasks related to each behaviour (and potentially at particular levels). This is with the aim of highlighting usability issues related to each behaviour that, if addressed effectively, can lead to an improvement in the usability of the resource.

One aspect that makes the IB methods novel is the specialised nature of the methods, which were developed to facilitate the evaluation of electronic resources as opposed to other types of interactive system. Another aspect is the extensibility and flexibility provided by the methods, which allows evaluators to tailor various aspects of both the IB functionality and usability methods to meet their particular needs whilst simultaneously providing enough structure to help ensure that rich data is obtained from the resultant evaluations.

6.7   Chapter summary

In this chapter, we have presented the IB functionality and IB usability methods – two methods for assessing electronic resource functionality and usability based on the findings of our study of lawyers' information behaviour. An IB functionality evaluation involves determining whether and how an electronic resource supports certain information behaviours at particular levels, along with additional ways that these behaviours might be supported at a particular level. An IB functionality evaluation also involves considering whether it is

still necessary to support all of the behaviours/levels that the resource currently supports. An IB usability evaluation involves setting a range of (usually behaviour-focused) tasks to intended or actual users of an electronic resource, asking them to think aloud whilst performing the tasks, and filling out a usability form listing user comments/actions or personal observations that might suggest a usability issue, details of the underlying usability issue, and judgements on how severe the issue is and how much effort is required to address it.

The IB methods were developed in an iterative manner, employing first a 'theory driven' and then an 'experimentally driven' development approach. In this chapter, we have discussed a series of three pilot think-aloud sessions with users of electronic legal resources (referred to as our 'user pilots') and a pilot think-aloud data analysis session with an electronic resource developer (referred to as our 'developer pilot'). In the next chapter, we evaluate the IB functionality and usability methods with a small group of stakeholders (including usability experts) working for LexisNexis Butterworths – a large electronic legal resource development firm. This is with the aim of ascertaining how usable, useful and easy to learn they think the methods are and how likely they are to use them in the future. Many of the findings from this evaluation session have already been incorporated into the versions of the IB methods presented in this chapter (and in particular into the supporting documentation in appendices 2 and 3 and the usability and functionality forms in appendices 6, 7 and 8).

Chapter 7: Evaluating the evaluation methods

This chapter at a glance…

In this chapter we:

- Present the methodology and findings relating to our formative evaluation of the Information Behaviour methods.
- Reflect on the success of the evaluation session.

7.1 Overview

This chapter details the methodology and findings relating to our formative evaluation of the IB functionality and usability methods with a group of stakeholders working for LexisNexis Butterworths – a large electronic legal resource development firm. We begin in section 7.2 by highlighting the aim of the evaluation session, discussing several 'success measures' for HCI evaluation methods in the context of the IB functionality and usability methods. Next, in section 7.3, we detail the format and content of the evaluation session followed, in section 7.4, by a discussion of the methodology used to evaluate each of the methods. In section 7.5, we discuss the findings from these two evaluations and improvements to the methods based on these findings. This is followed, in section 7.6 by a reflection on the formative evaluation session.

7.2 Aim of the evaluation session

The aim of the formative evaluation session was to determine the success of the IB functionality and usability methods. However, there are a number of possible ways of determining the 'success' of evaluation methods as summarised by Blandford and Green (2008). These include examining the:

- **Validity of the method** – ensuring that the method, as far as possible, supports the analyst in correctly predicting user behaviour or identifying problems they will have with the system, minimising the number of false positives (issues that are identified as problems that actually are not) and misses (failures to identify actual problems that are within the scope of the method).
- **Reliability of the method** – the extent to which different analysts applying the same method will achieve the same results.
- **Productivity of the method –** the number or severity of problems that the method helps to identify. Blandford and Green emphasise that, in recent years, this requirement has been 'relatively de-emphasised.'
- **Usability of the method** – the extent to which the method fits within design practice.
- **Learnability of the method** – how easy it is to learn the method. Blandford and Green suggest that *"how easy it is for a practitioner to pick up and work with a method will be a strong determinant of take-up and use in practice"* (p. 4).
- **Insights derived from the method** – whether or not the method yields insights that will help to improve design.

We now discuss each of the above 'success measures' in relation to both the IB functionality and IB usability methods.

Regarding *validity*, it can be argued that the IB methods both have some *inherent validity* due to the fact that they are based on empirically observed behaviour. We argue that, for the IB usability method, the validity of the method is only *dependent on evaluators' interpretations of the user think-aloud data*. However, we hypothesise that these 'interpretations' will vary depending on the complexity of the electronic resource under examination and a host of other factors (such as the evaluator's familiarity with the resource and their level of experience in evaluating electronic resources). We also hypothesise that evaluator interpretations of think-aloud data will vary depending on the richness of the data that the method is used to analyse. We suggest that analysing rich think-aloud data can result in the identification of usability issues of varying complexity – ranging from relatively straightforward surface-level interface issues to issues pertaining to the information-seeking task (which also manifest themselves at the interface level, but are far deeper-rooted than the surface-level issues). In light of the expectation of subjectivity amongst evaluators, we did not believe it would be appropriate to place heavy emphasis on the validity of the IB usability method.

A related argument can be made for the measures of *reliability* and *productivity*. The somewhat subjective nature of the data analysis process, which is dependent on the think-aloud data to be analysed, suggests that the IB usability method should not strive for inter-rater reliability (although, of course, we should still see similarities in the usability issues identified across evaluators if they are analysing the same think-aloud data). The subjective nature of the data analysis process also suggests that the method should not simply focus on problem count or severity – due to the fact that deep-rooted issues might be interpreted in different ways by different evaluators. It is not simply the case that some issues (such as the more deep-rooted ones) might be missed by certain evaluators, but that different evaluators might have a different interpretation of the underlying causes of an issue that has manifested itself at the interface level. To complicate matters further, issues might be related to one another. All of these factors would make it difficult (and arguably not very meaningful) to measure the productivity of the IB usability method.

The measures of *validity, reliability* and *productivity* are also not particularly suitable for determining the success of IB functionality method. This is because the functionality method aims at supporting functionality discussion in relation to particular behaviours and levels as opposed to identifying functionality-related issues with an electronic resource. As these measures are relatively 'issue-centred,' they do not have much meaning in the context of the IB functionality method. In addition, regarding *validity*, the IB functionality method does not necessarily require evaluators to obtain 'the same results' when evaluating an electronic resource – whilst it is expected that different evaluators who are examining the same electronic resource are likely to identify functionality support for similar behaviours and levels, the creative nature of the IB functionality method (which requires evaluators to think of *additional ways* that behaviours might be supported at particular levels and to discuss the *arguments for and against* support) suggests that the output from one group of evaluators is unlikely to be the same as that from another group of evaluators.

Our main aspiration for both the IB functionality and usability methods was to create a method that stakeholders with an interest in electronic legal resources would *actually want to use.* This relates to the final three 'measures' discussed by Blandford and Green – *usability*, *learnability* and *insights derived*. The fit of the method into design practice (the method's *usability)* was extremely important for us in order to help 'complete the loop' between our empirical observations of lawyers using electronic legal resources and the improvement of these resources. We could only hope to improve these resources (and therefore make lawyers' information work easier) through successful take-up of the IB methods amongst developers and other stakeholders. It also follows that take-up might well be influenced by the *learnability* of the method and the *insights derived* from the methods (i.e. whether or not the method yields insights that will help to improve design or how *useful* the methods are to those who choose to apply them). All three of these measures are related to our main aspiration of creating methods 'that stakeholders with an interest in electronic legal resources would *actually want to use*' and are the measures which we used in this formative evaluation study to determine the success or failure of the IB methods. It therefore follows that our evaluation of the methods aimed to find out whether they were **useful**, **usable**, **learnable** and whether they were **likely to be used** in future. These aims for our evaluation are mirrored as important by Buckingham Shum and Hammond (1994), who assert that the successful uptake of evaluation

methods by practitioners *"depends on how easily they can be understood, and how usable and useful they are"* (p. 1).

The 'measures' of usefulness, usability, learnability and likeliness of future use were also used in the development of CASSM – an HCI method for assessing the 'quality of fit' between how users think about activities and the way systems represent these activities. The development of the method is discussed in Blandford and Green (2008). Although Blandford and Green regarded learnability as an integral part of usability and therefore did not consider it separately, we have essentially employed the same success criteria as that used to evaluate CASSM (albeit for our method-specific reasons described above rather than to follow a trend). The concepts of usefulness, usability, learnability and likeliness of future use were used to structure both the focus group sessions and short questionnaires issued to participants during the evaluation session. We discuss the collection of this and other data from the evaluation session in section 7.4. Now, however, we turn to discuss the format and content of the session.

7.3   Format and content of the evaluation session

The formative evaluations for both the IB functionality and usability methods were conducted on the same day as part of a full-day tutorial presented to UK and European LexisNexis Butterworths staff. The tutorial, aimed both at teaching them the IB methods and at collecting evaluation data, was conducted at the firm's London offices (although participants were recruited from across the UK and Europe). The firm was provided with a brief summary of both methods and an agenda for the tutorial. This information was used to help our main contact at the firm decide who to invite to the session. We did not prescribe the job roles that participants should have, but instead allowed anyone with a potential interest in learning to use the methods to attend the session. The only stipulation was that participants should already have a basic grounding in usability evaluation. This resulted in the attendance of participants with a variety of job titles and usability evaluation experience (see table 15 below).

Anonymous participant number	*Job role*	*Participated in IB functionality method tutorial and evaluation session?*	*Participated in IB usability method tutorial and evaluation session?*	*Filled out short questionnaire related to the sessions attended?*	*Handed in IB usability form (from usability practice session)?*
P1	User Experience Consultant (usability expert)	✓	✓	✓	✓
P2	Product Manager (for LexisNexis Butterworths - the resource under evaluation)	✓	✓	✓	✓
P3	Taxonomies Product Consultant	✓	✓	✗	✗
P4	Business Development Manager (and former User Experience Consultant)	✓	✗	✓	✓
P5	Product Manager (for LexisNexis Butterworths)	✓	✗	✓	✗
P6	User Experience Consultant (usability expert)	✓	✗	✗	✗
P7	Quality Analyst (for non information-seeking products)	✓	✓	✓	✓
P8	Quality Analyst (for non information-seeking products)	✓	✓	✓	✓
P9	Product Manager (for LexisNexis Butterworths)	✓	✓	✗	✓
P10	User Experience Consultant (usability expert)	✗	✓	✓	✗

Table 15: Details of participants in the IB functionality and IB usability method evaluation sessions.

The tutorial was split into two autonomous components. In the morning, participants were taught the behavioural theory underlying the IB methods (which is necessary for conducting an IB functionality evaluation, but not for conducting an IB usability evaluation) and how to conduct an IB functionality evaluation. They were then given the opportunity to practice evaluating the functionality of one of the electronic legal resources currently under development (a new version of LexisNexis Butterworths). In the afternoon, participants were taught how to conduct an IB usability evaluation and given a similar opportunity to practice analysing the user think-aloud data collected as part of our user pilot studies (discussed in chapter 6). Note that the tutorial did not include a practice on setting user think-aloud tasks and collecting the resultant think-aloud data as we did not deem this to be a feasible or practical component of a half-day tutorial.

Most participants attended both the morning (IB functionality method) and afternoon (IB usability method) teaching and practice sessions (see table 15). We do not believe that P10 was disadvantaged by not attending the morning theory session, as detailed knowledge of the information behaviours and the levels at which they can operate is not necessary for conducting an IB usability evaluation (and certainly not necessary for analysing pre-collected think-aloud data, which is what the practice task involved doing).

The content of the teaching sessions was strongly based on the 'guidance for conducting IB functionality and usability evaluations' presented in appendix 2. The tutorial format was chosen with the need to make the method learnable in mind and included a mix of theory and hands-on practice exercises (with two main practices that were recorded as part of the data collected to evaluate the IB methods). Guidance was also drawn from a tutorial by Blandford, Fields and Keith (2003) on the 'Usability Evaluation of Digital Libraries.' This tutorial was similar in nature to our own and involved teaching a variety of HCI tools and techniques (such as scenarios and personas) to a group that was previously unfamiliar with them. The most important guidance that we drew from this tutorial concerned building plenty of concrete examples and opportunities for practice into our tutorial session (as Blandford et al.'s tutorial was heavily example-based and included a series of hands-on exercises). We felt this would help make our tutorial more concrete and therefore the method easier to learn.

The broad format of the tutorial involved:

1. Providing participants with an introduction to the behavioural approach for evaluating electronic legal resources.
2. Introducing the information behaviours in small groups, each illustrated by a video-clip example. This was followed by asking participants how an electronic legal resource they are familiar with supports each behaviour that had been taught to them and a practice involving the participants identifying the information behaviours displayed in a series of short video clips.
3. Providing a quick overview of the IB functionality method and then giving participants the opportunity to identify how a competitor resource supports one of the behaviours they had just been taught (i.e. without having learned the details of the IB functionality method yet). This aimed to provide participants with the 'building blocks' upon which to build further concepts, whilst grounding the tutorial in a concrete exercise as early on as possible. This exercise was conducted as a group, with the lead developer of the IB method acting as both facilitator and demonstrator. This involved the participants asking him to perform certain actions on the projector, using the competitor resource, so that they could determine in which ways it currently supports 'surveying' behaviour.
4. Introducing the levels at which information behaviours can operate, enforced by video examples of how these multiple levels of 'surveying' behaviour are supported by an example electronic legal resource (Justis). This was followed by asking participants to identify the levels of 'recording' behaviour displayed in a series of short video clips.
5. Discussing how to conduct an IB functionality evaluation in detail (see appendix 2 for detailed content), followed by a 1 hour practice session using the method to evaluate the functionality of an electronic resource currently under development in the firm – the new version of LexisNexis Butterworths). This is discussed in detail in section 7.4.

6. Providing a quick overview of the IB *usability* method and then giving participants the opportunity to watch a short video clip and identify usability issues from it. This was followed by a plenary discussion about the issues found.
7. Discussing how to conduct an IB usability evaluation in detail, followed by a 1 hour practice session using the method to analyse pre-recorded think-aloud data of lawyers using an electronic legal resource previously developed by the firm (the July 2007 public release version of LexisNexis Butterworths). This is also discussed in detail in section 7.4.

Participants were periodically given the opportunity to ask questions (and slides were presented to them at appropriate points during the tutorial asking them if they had any questions). In practice, the participants asked most of their questions during the tutorial itself, interrupting when they required clarification or wanted to ask any other type of question.

The tutorial was conducted in one of the firm's meeting rooms, which included a projector which was used to display the tutorial slides (and to display a competitor resource used in one of the IB functionality method practice sessions). Participants came equipped with their own laptops and a pair of headphones (to be used for the IB usability method practice session, which involved analysing a video clip of user think-aloud data). Whilst we did not have much control over the meeting room layout, the round-table arrangement proved to be adequate for both the teaching and practice elements of the tutorial. We also do not believe that the room format had any effect on the focus groups that were conducted (these focus groups are described in detail in section 7.4).

As the tutorial was strongly based on the information presented in this chapter and the guidance materials in appendices 2 and 3, we do not believe that it would be necessary to provide face-to-face training in order for others to use the methods successfully (although we recognise that face-to-face training may well be beneficial). To support those who may be interested in using the IB methods in future without face-to-face training, we have made the PowerPoint slides from the tutorial and many of the practice examples/exercises available online at www.uclic.ucl.ac.uk/people/s.makri/IBMethods.html.

7.4   Evaluation session methodology

Aside from teaching the stakeholders the IB methods, the one-day tutorial was devised with the aim of *evaluating* the success of both methods (i.e. determining whether the methods are useful, usable and learnable in the eyes of the stakeholders, and whether the stakeholders were likely to use the methods in future). To achieve this aim, we collected a variety of complementary data:

*   **Output from the participants using the IB methods** (during the one-hour practice sessions built into the tutorial). This included audio-recordings of discussions that ensued as the participants used both the IB functionality and IB usability methods. The output also included the IB usability forms that participants used to record usability-related data (see appendix 8 for the form).
*   **Focus group data**, where the participants were asked questions on the usefulness, usability, and learnability of the methods and the likelihood of them using the methods in future (see appendices 11 and 12 for the focus group questions).
*   **Summary questionnaire data**, again focused on the usefulness, usability, and learnability of the methods and the likelihood of them using the methods in future.

The ethical issues surrounding the collection of this data, including informed consent issues, are discussed in section 7.4.3.

Whilst it is not completely possible to separate the tutorial content and format and our teaching skills from the perceived learnability of the methods (or indeed from the perceived usefulness, usability and likelihood of future use of the methods), the data collected aimed to isolate each of these 'success measures' as far as possible. Specifically, focus group and questionnaire questions were specifically targeted at each measure (e.g. 'how easy

was it to learn the functionality method?' and 'on the following scale, how likely/unlikely are you to recommend that you and your colleagues use the IB functionality method in the future?') In the next section, we discuss the methodology for evaluating the IB functionality and usability methods in more detail.

7.4.1   Methodology for evaluating the IB functionality method

Methodology for collecting output data from participants using the method

As part of the morning tutorial session, participants were given the opportunity to conduct a 'mini-practice session,' applying the IB functionality method as a group to a leading competitor resource. This mini-practice involved them determining whether the resource supported 'surveying' behaviour and, if so, in which ways. This mini-practice was not only devised to help participants get to grips with the IB functionality method as quickly as possible, but also to enable them to ask questions about the method now that they had some experience applying it. The mini-practice also served to prepare participants for the format of a larger, one-hour practice, which was conducted at the end of the morning session. This larger practice session, which followed the same format as the mini-practice, involved all 9 participants present for the morning tutorial session conducting an IB functionality evaluation of a large information-based product that they were currently developing. To support the session, a limited-functionality prototype version of the electronic resource was made available on the shared projector screen, along with the current 'live' version of the product (in case it was necessary to examine any functionality that would be present in the final version of the product but was not available in the prototype version). The electronic legal resource was chosen in advance, in consultation with the participants. The only stipulation made was that the resource should be information-based.

Whilst there was no need to restrict the *boundaries* of the evaluation to particular parts of the resource (particularly since the prototype version was already a cut-down version of the intended final product in the sense that it did not feature pages that could be customised to different legal practice areas), time restrictions meant it was necessary to conduct the IB functionality evaluation only in relation to the *core* set of information behaviours. Regarding the *practicalities* of the evaluation, the lead developer of the IB method (who was also giving the tutorial) acted as both demonstrator and facilitator. The demonstrator role involved performing tasks using the prototype version of LexisNexis Butterworths (and the live version where necessary) under the direction of the group. The facilitator role involved prompting the group to answer the IB functionality questions in relation to each core behaviour and applicable level (see appendices 6 and 7 for the corresponding forms). As it would have been impractical for the demonstrator-facilitator to fill out these forms during the evaluation (and having one of the participants fill the role of note taker would have slowed down the evaluation and possibly resulted in it over-running), it was decided to audio-record the session instead. The facilitator role also involved ensuring fair and equal participation amongst participants. During the mini-practice session, the participants expected that the demonstrator-facilitator would guide the functionality evaluation more (e.g. by suggesting functionality to explore using the prototype in order to answer the IB functionality questions). However, within a couple of minutes it became clear to participants that the demonstrator role was restricted to acting as a 'puppet' when exploring the resource (rather than guiding them through the resource) and the facilitator role was restricted to prompting the participants to answer questions (rather than *assisting* them in answering those questions). In the larger practice session, the participants had got used to the demonstrator-facilitator role and took greater ownership of the IB functionality evaluation from the outset.

Methodology for conducting focus group and issuing summary questionnaires

Following the one-hour IB functionality evaluation practice session, the participants who had taken part in the session were issued with a short summary questionnaire with four questions, each aimed at addressing each of our success measures and to be answered on a five-point Likert scale. The questions were:

1.   How **easy/difficult to learn** do you consider the IB functionality method to be?

2. How **easy/difficult to use** do you consider the IB functionality method to be?
3. How **useful** do you consider the IB functionality method to be?
4. How **likely/unlikely** are you to recommend that you and your colleagues use the IB functionality method in the future?

Care was taken to make the questions and scale as neutral and unbiased as possible and the questionnaire was issued before the focus group session as we believed that the likelihood of the focus group session influencing individual questionnaire responses was more likely that the questionnaire influencing the focus group discussion. The questionnaire was not issued with the hope of obtaining statistically significant results, particularly as the number of participants in both the morning and afternoon tutorial sessions was small. Instead, we regarded the questionnaire as a useful 'quick check' on the success of both the IB functionality and usability methods. Whilst there are potential halo-effect implications attached to the questionnaire, participants were urged to answer anonymously and truthfully. We believe that complementing the questionnaires with the collection of focus group data helped to reduce the potential halo effect (and, indeed, participants' summary questionnaire responses appeared to correspond well to the focus group comments that they made).

The focus groups lasted around half-an-hour and were facilitated by a colleague who had experience in conducting focus groups, but was not involved in developing the IB methods (indeed, the facilitator had no knowledge at all of the method). Although the facilitator had met some of the participants before as part of an unrelated tutorial, we do not believe this had any effect on the focus group discussion. We believe that having an impartial facilitator was important for avoiding any potential halo-effect bias in the focus group data and the data collected, which included a mix of positive and negative comments about both methods, seemed to support this assertion. We had considered issuing detailed questionnaires and conducting semi-structured interviews as an alternative to focus groups, but decided that focus groups were likely to be the best option. We decided they would be more likely to yield richer data about the method than if questionnaires were issued (even if the questionnaires included mostly open questions). We also decided focus groups would be more practical than interviews, especially since the tutorial was only a day long and therefore individual interviews would either need to be extremely short, conducted by a team of interviewees or conducted another time. None of these options were desirable.

7.4.2 Methodology for evaluating the IB usability method

Methodology for collecting output data from participants using the method

As in the morning practice session, participants in the afternoon session were given the opportunity to conduct a mini-practice session, which involved identifying usability issues from a three minute clip of a law student 'trying to find out if there had been any amendments to a particular piece of legislation.' The usability issues identified were discussed in a plenary session and participants were encouraged to ask further questions about the IB usability method. As in the morning mini-practice, this practice session served to prepare participants for the larger one-hour IB usability practice session.

The one-hour IB usability practice session involved the seven participants that attended the afternoon tutorial session analysing one of two video clips, which were compilations of screen recordings of a series of behaviour-related information tasks that our pilot users (two undergraduate law students and one Trainee Solicitor) performed as part of our three user pilot sessions. Although the video clips were edited in order to be compiled into two longer clips, none of the content was removed or edited in any other way. In order to give tutorial participants the opportunity to make usability-related observations from a variety of different types of information tasks, each of the two video clips included tasks related to a variety of the information behaviours at the heart of the IB methods. These tasks were chosen from the illustrative examples in appendix 3 which, in turn, were based on the information behaviour we observed as part of our empirical study. The first video clip included screen-captured audio and video of the lawyers performing the following *recommended* tasks:

1. Try to find out whether a particular case is still good law.
2. Try to find out whether a particular piece of legislation is currently in force.
3. Try to find out whether there have been any recent developments in a particular legal area.
4. Try to set up an alert so that you can be informed every time there are new developments in a particular legal area.
5. Download two documents into a single file.
6. Keep a softcopy or hardcopy record of *part* of a document that is important to you (e.g. print or download only certain parts of a case report).

The second clip included screen-captured audio and video of the lawyers performing the following *core and custom* tasks:

1. Try to gain access to the electronic resource.
2. Try to find out which parts of the electronic resource you have access to.
3. Think of some information that you currently need or have recently needed to find for your work and demonstrate, using the electronic resource, how you might go about finding it.
4. Try to find out which sources contain information about a particular legal area.
5. Try to set up an alert so that you can be informed every time new documents are added to the system that match particular search terms.
6. Try to conduct a more advanced search that is restricted to a particular legal area.
7. (a) Try to follow a hyperlink or other form of connection from a legal case to a previous case or piece of legislation mentioned in the case report and (b) Find a particular case, then find out which more recent cases have mentioned it (if any).
8. Try to find a particular legal journal article, then found out which more recent articles have mentioned it (if any).
9. Try to determine what information is provided whilst you are *browsing* that might help you decide what documents might be relevant or which documents to click on and read.

Each compiled video clip included tasks performed by all three user pilot participants. This was a necessary rather than intentional feature of each compiled clip, as each pilot user performed different tasks during the pilot (and therefore there was not enough task overlap to feature only one user per video clip). Upon reviewing the compiled and original video clips, we did not think that this necessary edit would have a sizable effect on the IB usability practice session. There were few repeated usability issues or issues that had arisen gradually across tasks. Therefore we do not believe that separating the tasks across compiled video clips will have altered evaluators' subjective comments on the severity or ease of addressing the vast majority of usability issues (and we believe this will have had an almost negligible effect on the remaining issues). See chapter 6 for further details on the user pilots and how the above tasks were devised.

As a result of conducting our earlier developer pilot, we estimated that there would only be enough time to analyse 30 minutes of screen-recorded data in an hour practice session. Therefore although we included a variety of information tasks in each video clip, care was also taken to ensure that both compiled clips were each under half an hour in duration and that one clip did not last considerably longer than the other. The first (recommended) clip lasted 27 minutes and the second (core and custom) clip lasted 28 minutes. Most participants were able to analyse one entire clip, although some participants ran out of time before they could review the final few minutes of each clip (see appendix 4 for details). This suggests the need to allow slightly more than an hour for any similar session conducted in future (or to reduce the length of both compilation clips).

Participants in the IB usability evaluation practice were asked to work either individually or in pairs. Two pairs were formed and both were audio-recorded as they conducted the evaluation and recorded the output on the IB usability form (see appendix 8 for form). One pair (participants P1 and P7) analysed the 'recommended' video clip, whilst the other pair (participants P4 and P6) analysed the 'core and custom' pair. Pairs were chosen to ensure that one member of the pair had a detailed knowledge of the product under evaluation (LexisNexis Butterworths), whilst the other did not. This, once again, was to encourage discussion amongst each pair.

The two pairs were placed in separate rooms in order to ensure that they did not disturb the other with their discussions. The aim of recording the pairs was to capture any comments made about the method itself and to identify any difficulties that the pairs encountered. The remaining three participants conducted the usability evaluation individually and were each randomly assigned a video clip (participants P9 and P10 reviewed the 'recommended' clip and participant P2 reviewed the 'core and custom' clip). Wearing headphones, each individual worked in the same room as the pair that was analysing the *other* video clip to them and were not permitted to collaborate with others. All participants were given around an hour to analyse their video clip compilation and to record the output. Although only one IB usability form was requested from each pair, participants P1 and P7 decided to each fill out a separate form, even though they watched and discussed the video clip as a pair. Participants P3 and P10 did not return their forms at the end of the afternoon and therefore no output data was collected from these participants. Participants were encouraged to pause and re-wind the video clips as required and to focus primarily on performing a thorough analysis of the data rather than analysing the video clip in its entirety before the end of the hour session.

As measures of validity, reliability and productivity were not deemed to be appropriate measures of success for the IB methods (and because the number of participants in the IB usability evaluation was too small to be able to report quantitative results with any statistical confidence), we did not attempt to 'score' participants based on the volume of issues identified. Nor did we have sufficient data to form an opinion on how many evaluators would be necessary to identify most or all of the usability problems present in each video clip. The data analysis of the IB usability forms was therefore somewhat exploratory and involved 'grouping together' similar usability issues in order to get an impression of the types of issues identified by evaluators, those which the lead developer of the IB usability method had identified previously from the video clip which the evaluators did not and those identified by the evaluators from the clip which the developer had not identified. Our findings are discussed in section 7.5.2.

Methodology for conducting focus group and issuing summary questionnaires

Our methodology for conducting the focus group and issuing the questionnaires examining the usefulness, usability, learnability and likeliness of future use of the IB *usability* method was identical to that employed when evaluating the IB functionality method and is discussed in the previous section. Identical questionnaire questions were asked when evaluating the IB usability method to those used when evaluating the functionality method. Almost identical focus group questions were also asked when evaluating both methods. The questions only differed when they were referring to specific features of one of the methods (see appendices 11 and 12 for the list of questions used in each of the two focus group sessions).

7.4.3   Ethical issues

There were a number of important ethical considerations when conducting both the tutorial and the associated data collection. Firstly a Non-disclosure Agreement was signed to assure that we would not disclose private details to others, including in this thesis. This included details about the functionality provided by the prototype version of LexisNexis Butterworths evaluated as part of the one-hour IB functionality evaluation. This restriction, however, did not prevent us in evaluating the method by using it to evaluate the prototype. In this thesis we do not report the output of the functionality evaluation, only comments related to the functionality method itself.

As with all of the other studies reported on in this thesis, we obtained written informed consent from each participant. The informed consent form differed, however, from the somewhat standard forms used when interviewing and observing lawyers. The form listed each type of data that would be collected (i.e. recordings of practice sessions, recordings of focus group sessions and collection of IB usability forms) and allowed participants to either give or decline permission for the collection of each type of data. It was made clear on the

form that all participants in the group functionality practice session and in the focus group sessions will have given permission to be recorded and that those who gave permission to be recorded in pairs as part of the IB usability practice session would be paired with someone who had also given permission to be recorded. Participants were also informed that their participation in the tutorial would not be dependent on them giving permission for the above data to be collected and that they would not be penalised in any way if they did not give such permission. In particular, participants were informed that they would be given equal opportunity to practice applying the method as taught in the tutorial regardless of whether or not they gave permission for their data to be collected.

A designated employee from the firm was also appointed to review and, if desired edit, any material arising from the study that we intended to disseminate to an outside audience (including this thesis). Participants were made aware of this on the informed consent form. More conventionally, the informed consent form also asked participants to sign to signify that they understand what each part of the study involves, were aware that any details that could be used to identify them or the firm would be anonymised, that the audio and screen recordings would be disseminated and stored in accordance with the Data Protection Act 1998, that they could ask to review, modify or delete any date they provide for the study at any time and without penalty, and that the study had been granted full ethical approval from the UCL Psychology Department Ethics Committee (approval number PhD/2007/009). As discussed in chapter 6, the lawyers in the user pilot sessions had all given permission for their audio and screen recordings to be reviewed by third-parties (in this case the tutorial participants) in order to improve the product that they were using.

7.5   Findings and related improvements to the IB methods

We now discuss our findings from the practice sessions, focus group sessions and from the summary questionnaires issued to tutorial attendees. In section 7.5.1 we discuss our findings related to the IB *functionality* method, followed in section 7.5.2 by our findings related to the IB *usability* method.

7.5.1   Findings related to the IB functionality method and related improvements to the method

Findings from the IB functionality practice session

Several interesting observations were made during the IB functionality practice session. An important observation, with implications for spurring change to the IB functionality method, was that whilst facilitating the session, it soon became apparent that it was not useful to ask the group to rigorously consider arguments for and against supporting existing particular behaviours/levels. This was especially the case when recent decisions had been made to support particular behaviours/levels in order to fill gaps in functionality support. We asked the group whether future versions of the method should not require evaluators to consider arguments for and against support of behaviours/levels that were already supported, but instead encourage evaluators to do so only if any of them feel that the support is unnecessary. This suggestion for change was warmly received and one member of the group emphasised the 'need to occasionally question' whether particular ways of supporting a behaviour/level was necessary.

Although for the vast majority of the discussion, participants seemed to have understood the behavioural theory underpinning the functionality evaluation, there was some confusion about the definition of behaviours and levels on a couple of occasions. On these occasions, members of the group identified that they were now analysing a behaviour at a different level to when they began their discussion, but this illustrated the potential danger of IB functionality evaluations going 'off-course.' Similarly, on a couple of occasions there was some difficulty in determining how the resource supported certain behaviours (primarily 'chaining' and 'selecting/distinguishing/filtering') at the *source* level. This led to the discussion merging to a discussion on support at the 'document' level and suggests that more emphasis should be placed on helping evaluators

understand some of the more rarely-displayed levels of particular behaviours (and perhaps additional video examples of these should be provided).

Although the lead developer of the IB methods had issued participants with definitions of each information behaviour and examples of how each behaviour might be supported by electronic resources at each applicable level, there was almost no use of these supporting materials. This was despite strong encouragement that they should be used to support IB functionality evaluations. This has important implications for the IB functionality method if it is to successfully transfer into design practice. As the group conducted the evaluation without reference to any guidance materials, they were reliant on the understanding of the behaviours/levels (and the boundaries between them) that they had gained from the morning tutorial session. This suggests that the method should be bundled with enough *accessible* materials to ensure evaluators have enough grounding in behavioural theory to conduct successful IB evaluations. This suggests the need to bundle the theory with concrete examples, perhaps in the form of a short video tutorial, so that evaluators have a good chance of performing successful evaluations even if they make no reference to the written supporting examples (see appendix 3 for the examples).

In spite of the observations discussed above, the IB functionality evaluation progressed well and without facilitator intervention (other than to prompt participants with the relevant questions to answer and to suggest moving to the next behaviour/level). Participants had little difficulty in answering the questions when prompted and the resultant group functionality discussion appeared to be structured by these questions, but not constrained by them. We found this to be encouraging.

Findings from the IB functionality focus group session

We now detail findings from the thirty minute focus group session aimed at examining the usefulness, usability, learnability and likeliness of future use of the IB functionality method. Both this and the IB usability method focus group were facilitated by a colleague who was not involved in the development of the methods (and indeed had almost no knowledge of the methods at all). Almost all participants contributed in the focus group sessions, although to varying degrees. In general, the 'User Experience Consultant' participants (P1, P6 and P10) contributed most to the discussions. The facilitator aimed to ensure that all participants had the opportunity to voice their opinions, but did not ask individuals to contribute – all questions were asked to the group in general. We begin by discussing the findings related to the usefulness of the method, followed by the usability and learnability of the method. Finally, we discuss findings related to how likely the participants were to use the method in future. The findings are illustrated with participant comments. […] denotes that part of the comment has been omitted during write-up.

Findings relating to the usefulness of the IB functionality method

Participants found value in the IB functionality method, particularly as a structured way of supporting functionality discussions. As asserted by User Experience Consultant P1:

> P1: I think it is potentially useful. […] It is a structured way of facilitating a discussion about features instead of talking about features in general. So from that standpoint it was useful and out of that discussion came some new ideas like 'how can we access our products from documents?' We'd previously thought about it a little, but because of that structure it allowed us to really zoom in on it and talk about it in a structured way.

Product Manager P2 highlighted that the functionality method *"shows you what you're doing quite well and what gaps there might be that you can improve on."*

When asked how the IB functionality method compares to previous approaches used to evaluate the functionality of their products, participants with different job roles gave different responses (suggesting that different approaches were used across teams). For Product Manager P5 the 'real thing' related to the usefulness of the method (which we interpret to mean the real thing that makes the method useful) was that his team *"don't really have a way of looking at the functionality of our products."* Quality Analyst P8 and Product Manager P9 mentioned using 'GAP' analysis, which involves *"looking at how we compare to our competitors and seeing what gap there is there"* (P9). Product Manager P9 explained, however, that this form of competitor analysis is *"from a slightly different standpoint"* as compared to the user-centred IB functionality approach. As highlighted by participant P9: *"[The IB functionality] method obviously identifies a gap because it's saying these are the tasks that a user needs to do. Whereas what we've done a lot of is to say 'this is where we are in terms of our product development as we're going into a market, or in some cases new markets and you look at what you need to do to survive in that market.'"*

A similar point was made by Product Manager P2, who stated that he currently used *"a sort of checklist of feature functionality covering the entire scope of the product"* to identify functionality that his products do not currently support *"and then stacking that against what Westlaw does, for example and what other legal resources do as well."* The participant suggested that the IB functionality method may provide benefits over the checklist approach due to the fact that it is more user-focused:

> P2: I think this sort of ties in with the user need, doesn't it as opposed to a kind of heuristic approach which is kind of isolated from the user need which is kind of like 'ok, we win 3-2 on this measure.' [Group laughter ensues]. It doesn't really say anything else.

Overall, participant comments were positive concerning the usefulness of the IB functionality method. Product Manager P2 suggested that they *"could certainly use it in terms of looking at LexisNexis Butterworths and in terms of comparing it to our main competition,"* whilst Usability Consultant P1 highlighted that *"it showed us some gaps already. Some potential gaps."* P2 also thought the functionality method was *"useful as well because of its high level categories."* As explained by the Product Manager, findings from functionality evaluations *"always get filtered down in reports and so it is useful to have high-level summary for management."* The participant continued to assert that *"it would be useful to be able to present findings to management – to senior management, in this kind of manner. So it will make it more accessible I think."*

The functionality method, however, was not regarded by participants as being particularly useful for identifying opportunities for *reducing* functionality. As Product Manager P9 commented:

> P9: Ideas for new things were naturally coming out, but we weren't saying 'oh well, we don't need that.'

Similarly User Experience Consultant P1 commented:

> P1: In terms of the method itself, I didn't think we were going down that path. Because the way the questioning was, was 'is it there?' you know, 'what else could you do?' So it seemed like we were even trying to go beyond what we already had. So it didn't seem like the method had a mechanism built in to say, you know, 'take that out.'

Although the IB functionality method *did* include a 'mechanism' for questioning the need for currently supported functionality, we believe that making this an optional question to consider (and therefore not having the facilitator prompt participants to consider functionality reduction for each behaviour/level) may have discouraged participants from considering functionality reduction. This suggests a possible improvement to the functionality method. Since it is not always appropriate to consider functionality reduction for each behaviour/level (and doing so seemed to make the functionality evaluation seem a little pedestrian), but we believe it would be useful to consider functionality reduction in general, future versions of the method might

leave functionality reduction questions to the end. For example, the following questions might be asked at the end of the evaluation:

- Are there any *behaviours/levels* that it may no longer be necessary to support? For any behaviours/levels which you are considering ceasing support for, what are the potential arguments for and against ceasing support?
- Are there any *ways that you currently support any of the behaviours/levels* that may no longer be necessary? For ways of supporting a particular behaviour/level which you are considering ceasing support for, what are the potential arguments for and against ceasing support?

As highlighted by Product Manager P2, functionality reduction is *"a much more complicated thing to actually do, from both a product management perspective and for meeting market expectations."* Therefore we cannot make any strong assertions on the potential value of incorporating a mechanism for functionality reduction into the IB functionality method without further testing the method. The issue of functionality reduction is complex. We hypothesise that evaluators will be able to answer the new functionality questions listed above. However this is by no means certain, even if the evaluators are familiar with the functionality provided by the electronic resource and are armed with sufficient user data to answer them.

We suspect that it will take many attempts at modifying the IB functionality method to deal with functionality reduction issues satisfactorily. Indeed, we even regard our belief that 'it would be useful to consider functionality reduction in general' as a hypothesis that requires testing. We therefore view the issue of functionality reduction as an important challenge for future research and regard our work on the IB functionality method as a small, but nonetheless useful step forwards in addressing this challenge.

Findings relating to the usability of the IB functionality method

In terms of how easy the IB functionality method was to use, User Experience Consultant P1 commented that *"it was easy enough. There are certainly more complicated methods out there."* However, he was keen to emphasise that the current paper form version of the method might be replaced by an electronic version, such as an Excel spreadsheet, to cut down on the required paperwork:

> P1: It could be slimmer, particularly in relation to these worksheets. You essentially need like 30 of these and I don't see that happening. I see more like just the table at the top in like an Excel spreadsheet and that's where you put in the answers, right in there. I can't see us going through 30 different sheets personally.

Quality Analyst P7 commented that *"it was fairly easy as a method,"* but that many of the information behaviours at the heart of the method were not applicable to the products that he develops, which are not information-based. He suggests that his team might *"use a slimmed-down version of it. Take away some of the things that are not relevant to some of our products and use a slimmed-down version."* This comment suggests the possibility of examining whether the IB methods can, in fact, be applied to non-information-based resources (a possible direction for future work that is discussed further in chapter 8). Not only would this provide an insight as to whether it is possible to apply the methods to these types of resource in practices, it would also provide an insight into the benefits and limitations associated with doing so.

User Experience Consultant P1 identified some blurred boundaries between some of the behaviours (as highlighted earlier from our reflections on the functionality evaluation). He noted that as the behaviours 'overlapped so much,' it made the process of evaluating electronic resources in relation to separate behaviours 'a little vague' at times:

> P1: It was a little vague for me at times. To understand either the level. Well, actually, the level was easy. It was more the behaviour. And I think it's because they overlapped so much. Whether we're

talking about 'surveying' or 'accessing' or 'monitoring.' [...] It may not matter in the end, but at times I wasn't grasping where the boundaries were, particularly with some of the bigger ones like 'surveying.' [...] Levels was okay and I think once we started getting into it, it became clear too. Yeah, it was fairly straightforward.

Again, this might suggest the need to include additional (and perhaps video-based) examples of how current electronic resources support particular behaviours, particularly as none of the participants referred to the examples in the supporting documentation that they were issued with.

Findings relating to the learnability of the IB functionality method

Participant comments suggested the IB functionality was 'pretty easy' to learn. Product Manager P2 suggested this was *"because it does actually mirror what users do using the application."* User Experience Consultant P1 suggested this was because the method had *"familiar things in it."* P1 also suggested that "*sometimes it was stating the obvious, which is good. It was like 'oh, accessing, you've gotta access the product.' Of course. But that's not a bad thing, that's a good thing."*

Regarding the learning experience with regard to the IB functionality method, Product Manager P5 was positive about the example-based format of the tutorial, where the method was learnt *"through applying it to our own product."* Product Manager P9 was also positive about the example and practice-based tutorial format:

> *P9: For me there was the right balance of theory and actually showing us examples and practicing using LexisNexis Butterworths at the end. I think that worked well. So you felt you understood that it was actually founded in something, you know the theory. But there was enough practical stuff in the presentation as well, which helped.*

When asked about what was less positive about the learning experience, User Experience Consultant P1 suggested that he *"thought we could've gone through the theory quicker, but that's just my personal take."* P1 elaborated on this comment shortly afterwards:

> *P1: Because it was so obvious and familiar, you could have said 'here are the behaviours,' 'here are the levels' and 'here's an example of each' and then I'll buy it. I personally think we could have gone through that quicker.*

The above comments by User Experience Consultant P1 are interesting given that he was the participant that also flagged the potential blurred boundaries between behaviours (which we believe can only be effectively addressed through concentrating on the theory and providing *more* rather than *fewer* examples). Product Manager P5 suggested, in a similar vein, that *"maybe we could have got to the practical parts earlier."*

Both of these comments highlight the delicate balance between providing users of the IB functionality method with enough theory, examples and practice to enable them to perform the evaluation successfully whilst minimising the amount of time and effort required to learn the method. We believe there is a danger of the IB functionality method being perceived as too *'obvious and familiar,'* (to quote User Experience Consultant P1) precisely because there are not always clear boundaries between information behaviours. We do not have a quick fix for this issue, but suggest that it is something to continue to bear in mind when producing future revisions of the method and the accompanying teaching materials.

Findings relating to the likelihood that the IB functionality method will be used in future

Product Manager P2 asserted that 'it makes sense' for his team to use the IB functionality method as it is more user-focused than the current GAP analysis approach that he uses:

P2: It makes sense for us to use it because at the moment our functionality analyses, our GAP analysis, has got these high-level categories to be used to divide the functionality but we are doing just that – we're dividing the functionality, we're not dividing the user behaviour. I think it would help if we did that because it would tie-in with what the user-experience guys do and it might give us an earlier warning where we do have functionality issues. Our current approach is good to spot where there are gaps in future functionality, but that doesn't tell you anything about usability.

Similarly Product Managers P5 and P9 both suggested that the IB method might supplement their existing functionality checklist approach. P9 did not, however, *"feel it would be replacing anything that we do."* Product Manager P5 asserted that the IB functionality method *"has potential and could be very useful,"* suggesting that *"introducing something like that would be good because it's something new and would be changing your ways to find the right fit."* We interpret this to mean that participant P5 regards the method to be useful because it provides a novel approach to evaluating functionality. Product Manager P9 also mentioned that he was likely to use the method when developing new electronic legal resources, at the requirements stage:

*P9: I kind of like the idea of imagining you have a new product, particularly if you have a new product that you're starting. I think it could be really useful. You're not **comparing** it to anything else, you're just seeing, you know, 'I've got a new product, and the product I'm building.' [Pauses]. Almost it's like you reach a certain stage of your requirements you can ask 'what have I missed?' or 'how am I stacking up against those requirements.' I can envisage scenarios where it's used just by itself as well.*

A particularly interesting comment about the likelihood of future use of the IB functionality method was made by Quality Analyst P7, who asserted that not only would it be necessary to 'tailor the method' so that it could be used for non-information-based products, but also that the method might be conducted by *individuals* instead of in a group (although these individuals might invite other team members to review the output):

P7: I think, realistically, it's going to be one person completing this template looking at their own product and competitors' and maybe bringing in others to have a look at it, but I don't think you'd get a group of people sat in a room completing this template. It would just take too long.

Although no other participant commented directly on the possibility of conducting IB functionality evaluations individually, the Excel spreadsheet approach suggested by User Experience Consultant P1 might well be suited to individual as well as group evaluations. This suggests the potential for giving users of the method the choice of whether to conduct the functionality evaluation individually or in groups (although we might recommend input from other members of the team if the evaluation is to be conducted individually).

Findings from the short questionnaire on the IB functionality method

The short questionnaire issued to participants shortly before the focus group session mirrored the generally positive comments made during the focus group (see figure 16). As we have touched on previously, due to the small number of participants in both evaluation sessions, we cannot use the questionnaire data alone to make strong assertions about the success of the IB methods. However, we *can* use the questionnaires in conjunction with the other evaluation data collected to build a picture of the tutorial attendees' opinions of both the IB functionality and IB usability methods.
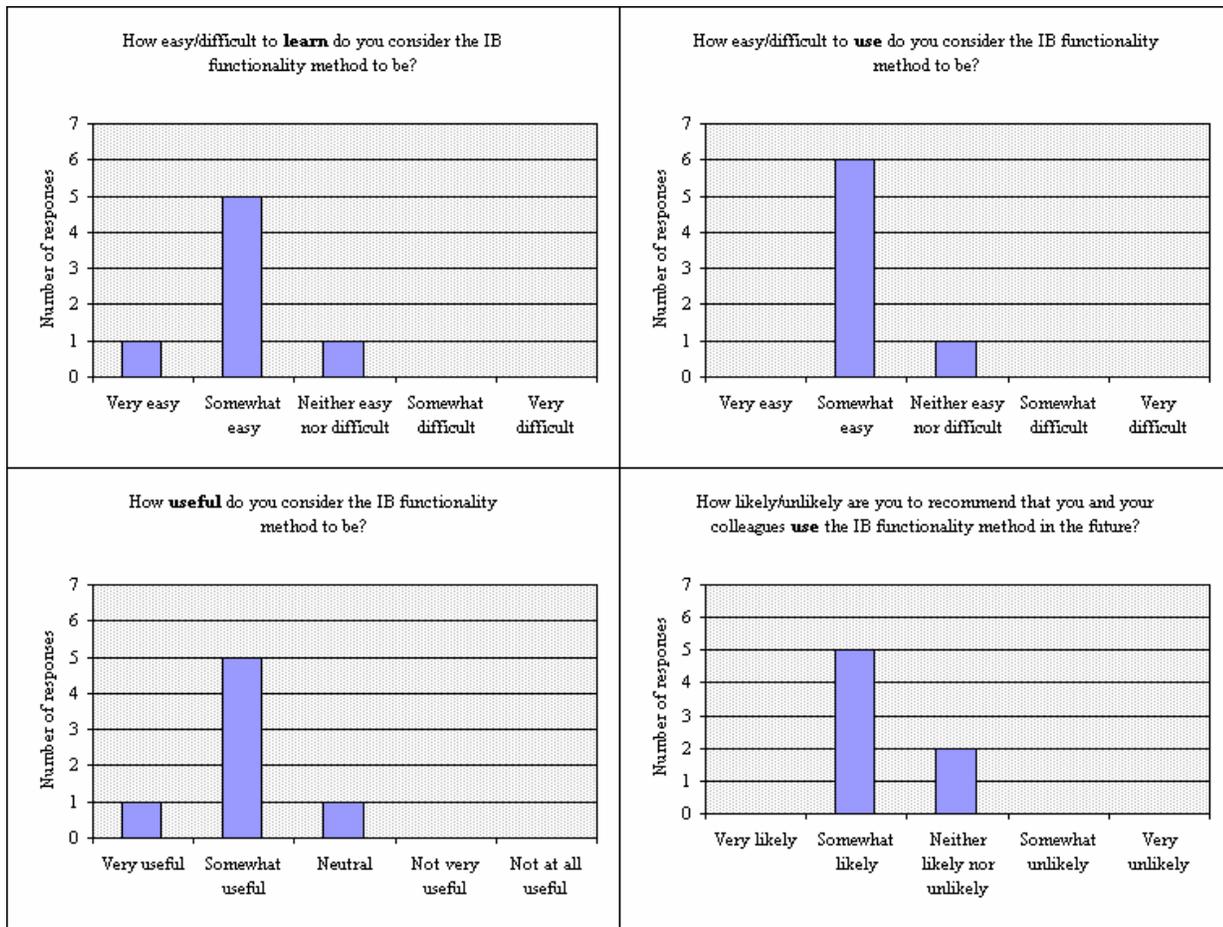
Figure 16: Summary of the questionnaire responses received about the usefulness, usability, learnability and likelihood of future use of the IB functionality method.

7.5.2 Findings related to the IB usability method and related improvements to the method

Findings from the IB usability practice session (and from the resultant IB usability forms)

The IB usability forms were filled out with broadly similar information (see appendix 4 for a consolidated list of usability issues identified by both the lead IB usability method developer and the six participants who handed in completed usability forms). However, different participants defined the term 'usability issue' in different ways (example comments in this column of the form included 'didn't really complete task,' 'usability of sub-tabs,' 'not clear what checkboxes' and 'conceptual – clarity between alerts and saved searches'). Most participants filled out the usability form in its entirety (although one participant, P7, did not fill out the screen/pages/parts of the resource column, whilst the other member of the pair, participant P1, did).

Only light use was made of the 'reflections' column. Participant P2 used the column only to record that he believed two issues (the participant not knowing what sources he had access to and the user looking for articles relating to the current document eve though this is not possible using the resource) were 'training issues.' Similarly pair P4 and P8 deemed that knowing how to set up an alert (and the difference between an alert and a scheduled search) was 'probably a training issue.' However, in general, all forms and both recorded pairs seemed to attribute usability issues to the electronic resource, rather than suggest that they were user errors (or indeed 'training issues'). As highlighted by P10, there was a general attitude that *"if it's not clear to them, it's a problem in our design."*

The pairs in our study worked well together, discussing the usability issues identified and how they might complete the usability form. The pairs both drifted into some limited design discussion on how the usability issues identified might be *addressed*, but this did not seem to send the evaluation off-course. Both pairs also discussed the severity and ease of addressing the issues that they identified. Participant P7 tended to default to the User Experience Consultant (P1's) subjective opinions, but in general both pairs reached a consensus based on discussion. As both pairs included a User Experience Consultant and a Quality Assurance Tester without detailed knowledge LexisNexis Butterworths, functionality-related questions were often posed to the Consultant in order to help the Tester understand some of the usability issues. This suggests the feasibility of conducting IB functionality evaluations in pairs as well as individually.

As might be expected, participants tended to identify similar issues to others analysing the same video clip and to those identified from the same clip by the lead IB usability method developer. However, participants tended to use highly abbreviated sentences to describe even complex issues. For example, refer to recommended task 1 ('try to find out whether a particular case is still good law') in the consolidated list in appendix 4. The lead developer describes the issue in detail: *"Participant used the 'general search' field (rather than segmented fields) for all searches and did not use any Boolean operators. Perhaps the other search possibilities were not made prominent at the interface."* Corresponding descriptions from the participants (extracted from their completed IB usability forms) include:

- *P1 (paired with P7): Can't find segments (visibility). Didn't see.*
- *P7 (paired with P1): No filters on search terms (title) etc. (Can't find them). Design issues.*
- *P9: Can't find guided search form for cases. Search form tables not visible or function not obvious to user.*

As can be noted from the examples above, different short-hand was often used by different participants and much of this shorthand is likely only to make full sense to the person that wrote it. For example technical terms such as 'can't find segments' was used by the User Experience Consultant (participant P1) in one of the two pairs, whilst the other member of the pair (participant P7) who was relatively unfamiliar with LexisNexis Butterworths used different language, referring to the segmented fields as 'search terms (title) etc.' as opposed to 'segments.' This suggests that it may well be useful for evaluator(s) to report back on the issues identified to other relevant members of their team, perhaps demonstrating them using the electronic resource under evaluation. This will help ensure that the usability form output is as useful as possible to the entire team, not just the evaluator(s).

Most usability issues identified by the lead IB usability method developer were also noted on the usability form by at least one of the participants in our study. There were, however, some exceptions. These ranged from highly-subjective issues that we might not expect participants to regard as 'usability issues' in their own right (such as *'Participant assumed that 'old, obscure cases are more likely to be on Westlaw now' and suggest that 'perhaps there isn't much on here as there used to be'*) to more standard usability issues (such as *'how to use the 'Get a specific document' combo box is unclear. Participant was unaware of the necessity to type over the example text in grey'*). See appendix 4 for a full list of usability issues identified only by the lead developer (and not by any of the participants).

Usability issues identified by participants that had *not* been identified by the lead developer were, broadly speaking, related to query formulation issues. These were not perceived by the lead developer to be 'usability issues' per se, but were still difficulties that the users in the think-aloud clips faced nonetheless. Two examples from participants' usability forms include:

- *P7 (paired with P1): Initially can't find info. needed. Query formulation issues.*
- *P9: Issues formulating search. Producing either too few or too many results. System not providing sufficient support in search query formulation.*

A small number of other unique usability issues were also identified by participants, although it was not possible to interpret many of these issues with confidence as the participants often left the 'approximate time in video clip' column blank. A couple of the issues identified *were* identified by the lead developer, but during other parts of the video clip. The remaining issues identified by participants and not by the lead researcher (i.e. excluding query-formulation-related issues and those identified by the developer during other parts of the video clip) were:

- *P9: Needs to search again to find a case previously viewed. Recent documents function not found by user.* We interpret this to mean that the participant was unaware of the 'recent documents' functionality which would allow him to quickly return to a document that he had previously viewed without having to search for it again. The lead developer did not regard sub-optimal task choices as 'usability issues' per se., which explains why this issue was not identified by the developer. This comment relates to 'recommended' task 1.
- *P7 (paired with P1): Had to add search item before results were valid.* Meaning unclear, even after reference to appropriate approximate time-stamp in video clip. Relates to 'recommended' task 1.
- *P1 (paired with P7): > 3000 results returned.* We interpret this to mean that the electronic resource interrupted the search as too many results were returned. The lead developer did not regard this to be a usability issue (although the very fact that the resource did not support the user in narrowing the search might well suggest that this *could* be a potential issue). This comment relates to 'recommended' task 3.
- *P1 (paired with P7): Didn't really complete task.* We interpret this to be referring to the fact that the Trainee Solicitor in the video clip, when asked to 'find out whether there have been any recent developments in a particular legal area' completed a series of *search* tasks rather than a more classic monitoring task. The developer did not regard this to be a usability issue. This comment relates to 'recommended' task 3.
- *P1 (Paired with P7): Not clear what checkboxes.* Meaning unclear and no time reference point provided.
- *P7 (Paired with P1): 'Last 10' label unclear.* Meaning unclear and no time reference point provided. This comment relates to 'recommended' task 4.
- *P7 (Paired with P1): 'Tenancy' and 'Tenant' should be treated as similar terms, but are very different.* Meaning unclear. This comment relates to 'recommended' task 4.

It should be noted that all of the issues listed above relate to the 'recommended' tasks, which were less prescriptive and therefore, we believe, provide greater scope for interpretation. All usability issues arising from the more prescriptive 'custom' tasks that participants had time to analyse (tasks 4-7 from the 'core and custom' clip) were identified by both the lead developer and by at least one participant. This *may* suggest that analysing custom tasks could be less subjective than analysing 'core' or 'recommended' tasks (however as only a small number of participants took part in the analysis session, this should be treated as a hypothesis to explore rather than a firm assertion).

The issues identified above support our assertion that identifying usability issues from rich user think-aloud data is a subjective process. However, we were encouraged that most of the usability issues that the lead developer had identified from the video clip were identified by at least one participant (even if the participant described the issue using different terms).

Findings from the IB usability focus group session

We now detail our findings from the IB usability method focus group. Although participants were generally positive about the usefulness, usability, learnability and the likelihood of them using the method in future, comments from this session appear to be more mixed than those from the IB functionality method focus group session.

Findings relating to the usefulness of the IB usability method

User Experience Consultant P1 highlighted from the outset that apart from setting *information behaviour-focused* as opposed to other types of think-aloud tasks, the IB usability method did not differ much from other usability methods that they currently used. Indeed, this relates to the fact that we only tested the *analysis* part of the method, and not the task setting part of the method (which might prove to be challenging in its own right, as suggested by participant P1):

> *P1: It's always a problem to come up with the tasks. What are you going to test? Because you can't test everything.*

This issue was stated succinctly by User Experience Consultant P10 and this suggests the need to test the *task setting* part of the IB usability method in the future:

> *P10: I think a lot of the hard work is what is not visible to us on this paper, because we're not devising the scenarios that the users are performing. […] Really, that's the only place where you're using the method really. This is just using the results. I think the method **is** only devising the tasks.*

There were, however, many positive comments about the usability method. User Experience Consultant P1 highlighted that other usability evaluation approaches undertaken by the firm *"very often will be more system-centred,"* and suggested that using the IB usability method may make his team *"think more broadly in terms of user behaviour."*

User Experience Consultant P10, who commented that he did not have as much prior experience of usability evaluation as participant P1, was also positive regarding the usefulness of the method, referring to many of the user-centred features of the method:

> *P10: I don't have P1's experience of usability analysis because I'm working in a different field, but I think I would find it very useful for specific task analysis. I do a lot of work on user assistance and I'm always trying to find out what the problems are that users are trying to solve and this is quite a thorough way of going about answering that question, because you ask users to solve a prescribed task – you watch it, you analyse it, and at the end the conclusions are quite obvious about the problems that they face. You don't address those problems, you just try to find out what they are. It allows you to think about 'What is it that they are trying to do?' 'What are the pitfalls that may be consistent between different users?' and address them and help the user. […] I think the greatest benefit is that you are looking at user behaviour and addressing user needs rather than taking a more system-centred approach, which is good. That's the whole aim of usability.*

Participant P10 also commented that using the IB usability method might be a useful way of 'getting attention' for usability related problems (and therefore funding to address them) from senior management. He suggested that having 'a large body of users' thinking-aloud when using an electronic resource might help in this regard:

> *P10: What I find anyway with usability is that having a single usability problem isn't enough to get everyone's attention. If you can illustrate that you've got a number of issues in the same kind of area, then that becomes all of a sudden 'now that **is** a serious problem.' So if you had low priority issues previously, having a number of them could make it a high-severity issue. So I think having this framework, you know having a large body of users that you're doing research with, it can help highlight and get these issues addressed. Because usability falls dead in the pecking order. Businesses are always pushing for extra features, extra function and usability sometimes gets forgotten. But if you demonstrate that you're having serious problems in a particular area or areas, then it can help with your argument to get funding to address the problem.*

User Experience Consultant participant P10 also suggested that the method might be more useful in persuading management on the need to address particular usability issues if explicit reference was made to the behaviour-focused tasks when presenting results from the evaluation to others:

> *P10: You **could** make it explicit by bringing it back round in a circle. So you could say 'well I started with these behaviours based on these tasks' and you could then use that to turn that back around to say 'well this is why it's a severe issue and why we need to do something about it.'*

In relation to the usefulness of the method, participant P10 also commented on the potential subjective nature of the IB usability method:

> *P10: The level of frustration will vary greatly between different people, because they will be in different circumstances, trying to solve different problems. So I think it's subjective in terms of severity and maybe some of the other things. […] If someone's trying to do something repetitively and failing, it can appear to be very serious. […]*

The issue of subjectivity was also raised by Product Manager P9, who found that she would wonder if she had 'missed something' whilst filling out the IB usability form. Whilst she commented that the IB usability evaluation *was* useful, she stated that it could *not* 'give her objectivity':

> *P9: I must admit that I got lost when I was doing it, and I was thinking 'what was I meant to be doing?' I mean I knew roughly what I was meant to be doing, and I took lots of notes, but I was thinking 'did I miss something?' And I think it was something similar to what P1 was saying, that what we discussed this afternoon wasn't adding to what I was writing down. I think it's useful getting there, because I have a better understanding of how to decide what the tasks should be and understanding why they're important. But when I was actually sitting there doing it, I kept looking for something in my head. Particularly as I'm not very experienced in doing that kind of analysis. And so it did end up, as you were saying, quite subjective. It was useful, but if I was looking for something that could give me objectivity, I didn't find that.*

This highlights another delicate balance – between making the IB usability method prescriptive and providing with evaluators with flexibility in how to conduct the evaluation. We have always recognised that the IB usability method can be subjective in many ways – different evaluators might use different terms to describe the same usability issues or a particular 'part of the resource' or screen. Similarly we hypothesise that different evaluators will provide different severity ratings for the same problems (this remains a hypothesis as we did not have enough participants to make any firm conclusions). Interestingly, although participant P9 raised the issue of subjectivity, her IB usability form was filled out thoroughly and completely and did not show any signs of apparent misconceptions. We suggest that future tutorial sessions should include a short form-filling exercise as part of a mini-practice session, to give participants the required confidence to follow their convictions when conducting an IB usability evaluation. We do not believe that making the IB usability form more prescriptive will, necessarily, lead to an improved method.

Findings relating to the usability and learnability of the IB usability method

Several issues were highlighted relating to the usability and learnability of the IB usability method. Product Manager P9, who had already commented about the potential subjectivity of the method, mentioned that she would mix up the 'user actions/comments' column with the 'usability issues' column. However, she attributed this to it being the first time she had filled out the form. She also suggested that filling out the IB usability form was likely to get easier with practice:

> *P9: I found, and again this is just me, I mean I can go back through it and edit it all, but I was having problems with the first column [user actions/comments] because I was actually mixing that up with the*

third column, the issues. And it was the first time that I tried filling in the form and I was trying to keep up with this [points to think-aloud video] as well, but that's something I think with practice you can get down very well.

User Experience Consultant P10 commented that it would be 'too much' to fill in the reflective columns of the IB usability form whilst reviewing the video clip think-aloud data:

> *P10: I don't think there's any way you can fill in these middle columns here during the session. It's too much.*

This supports the findings of our developer pilot, which led to a change so that the columns could be filled out either *whilst* or *after* reviewing the data.

User Experience Consultant P1 suggested that the usability of the IB usability form might be improved by moving the 'screen/pages' column of the form "*more over to the left where you write the time.*" When Product Manager P9 commented that she did not fill out the 'approximate time' column, P1 suggested that perhaps the form should be 'more mechanical' by putting the 'screen/pages,' 'approximate time' and 'user comments/actions/own observations' columns together. We soon realised the potential benefit of incorporating this change into future versions of the method, particularly as when asked by the facilitator whether in their opinion would make the method 'more objective,' several participants answered affirmatively. The current usability form reflects this change (see appendix 8).

P1 also suggested that he 'might also change some of the severity codes' in order to standardise them with those already used by the firm:

> *P1: I'd have to look, but we've probably used other scales in the past and we should try to be consistent with that. We usually use a four or five point scale, 0-4. […] I wouldn't use that N/Q/V scale, but the principle of it is right – you want to break down the data. So I would definitely use that basic approach.*

When asked by the facilitator whether the scale might need more fidelity, P1 answered 'probably not,' explaining that he has previously used scales with less fidelity and therefore is *"not hung up about the fidelity just because we may use something different."*

Findings relating to the likelihood that the IB usability method will be used in future

Participant comments were generally positive about the likelihood of using the IB usability method in future, in particular with regard to the behaviour-focused aspects of the method (i.e. the setting of behaviour-focused tasks):

> *P1: I think the task creation I may use in my team. We're often charged with doing that – coming up with user tasks. Using the behaviours as a way of distributing the tasks so you're getting a full picture of the usability. I think that's good.*

User Experience Consultant P1 also suggests that the behavioural approach to task setting would also provide justification for setting particular tasks (i.e. in order to ensure that particular behaviours are covered during the usability evaluation):

> *P1: It also helps discussions too. So 'why do you have that question in there?' 'Well we want to do **this**, we wanna cover this behaviour,' you know. It gives reasons for things being done and that's often a challenge – getting agreement and things like that.*

Like comments relating to the IB functionality method, User Experience Consultant P1 suggested that the IB usability method might also benefit from being provided in electronic rather than paper form, again as an Excel spreadsheet or even in the form of a database so that findings can be compared across users and dynamically sorted:

> *P1: I don't know if he intended it to be electronic, but I would definitely use Excel. I've even used databases in the past. Not personally, but at other companies in order to compare tasks across users, so you can say 'show me how everybody did for a particular task.'* **Facilitator: So you can re-organise them?** *P1: Yeah, you can re-organise them, like you can say 'show me all the high-priority ones' and things like that. You can do that to some extent in Excel and I've done that. Where you have participant number 1, task 1, and then you can look down or across. So I'd definitely do this electronically. […] I mean we have this goal to have a usability database, with all of our data from all of our usability tests in a single format, and even categorised against a single taxonomy of problems. So we can then say, you know, 'with search query formulation across products, pick a search query formulation topic and see all of the usability tests which have issues with that.' That was our vision. So we have inklings of more standardised things that we would probably do on top of this.*

Again, we regard the suggestion to create an electronic version of both methods to be an interesting one that, if given the opportunity, we would like to follow up on in the future.

Findings from the short questionnaire on the IB usability method

The summary findings relating to the IB usability method questionnaire help to place the mixed (but overall positive) comments from the focus group into perspective. The questionnaire findings were highly similar to those relating to the functionality method, with participants providing a positive picture on the usefulness, usability, and learnability of the method and the likelihood of them using the method in future (see figure 17).
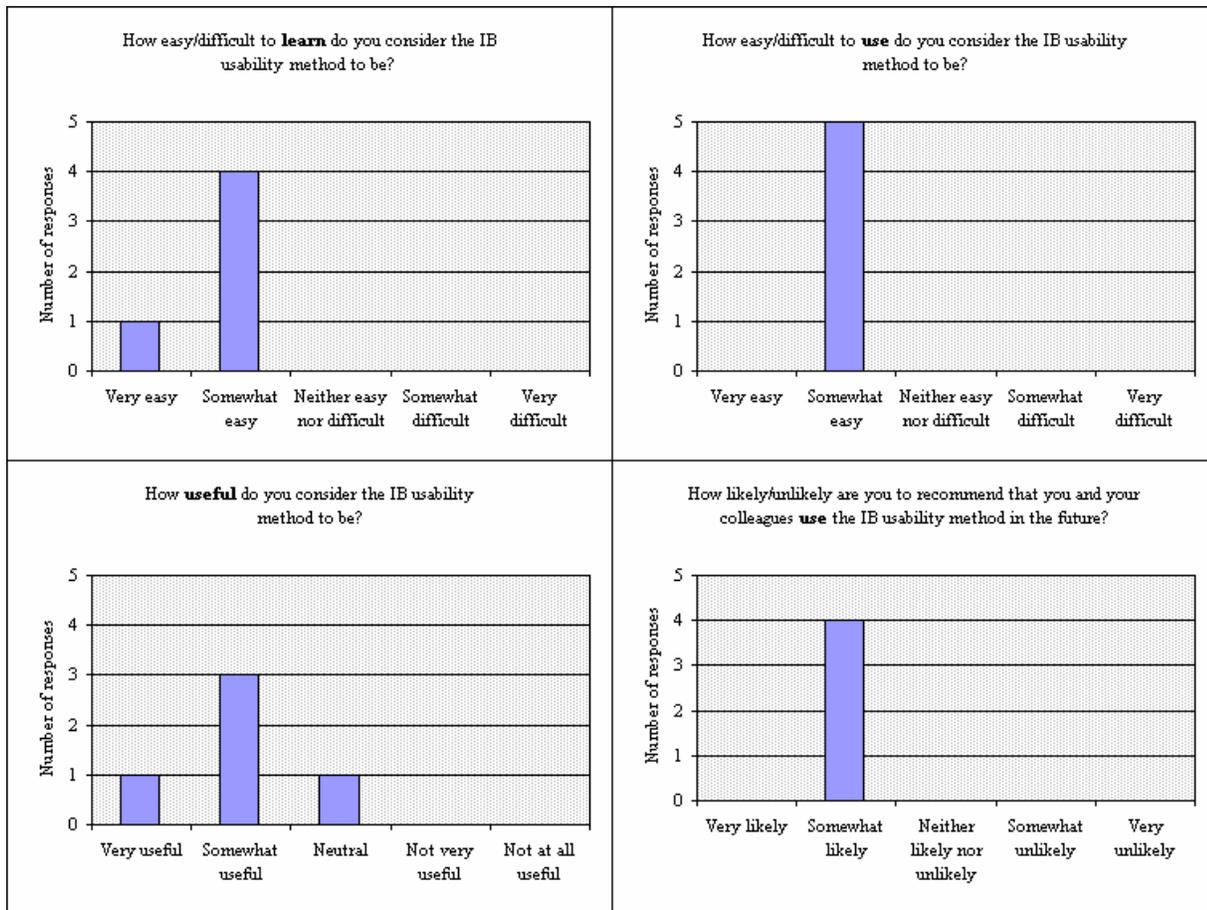
Figure 17: Summary of the questionnaire responses received about the usefulness, usability, learnability and likelihood of future use of the IB usability method. For the question on future use of the method, the remaining participant answered 'N/A.'

We believe the slight but notable discrepancy between the mixed focus group comments and the overall positive summary questionnaire data might be explained by the some of the initial comments made during the focus group. Consider the following comment about the IB usability method made by User Experience Consultant P1:

> *P1: I think the framework helps quite a bit. But the rest of the method, for me personally didn't really give anything other than a standard usability test.*

Whilst this can be interpreted in a negative light (i.e. that the participant believes that the analysis process of the IB usability method does not provide anything new compared with other usability evaluation methods), we believe the participant was simply stating that he believed the novelty of the method lay in the *setting* as opposed to the analysis of behaviour-focused tasks (a point made explicitly by User Experience Consultant P10 later in the focus group session and one which we do not disagree with). Perhaps the potentially negative spin on some of the initial comments made during the focus group session may have influenced the overall balance of the session. In any case, this provides further justification for the need to provide participants with the opportunity to set as well as analyse behaviour-focused tasks in future tutorial (and evaluation) sessions.

7.6   Summary and reflection

In this chapter, we have formatively evaluated the IB functionality and IB usability methods with a small group of stakeholders working for LexisNexis Butterworths, a large electronic legal resource development firm. The IB methods were taught during a one-day tutorial, which included a mix of theory, video-based examples and mini-practice sessions where tutorial attendees were given the opportunity to practice using the methods. Attendees were also given the opportunity to practice applying each of the methods in two hour-long practice sessions. The first session, aimed at giving them practice using the IB functionality method, involved applying the method in a group to a prototype version of an electronic legal resource that some attendees were currently involved in developing. The second session was aimed at giving tutorial attendees practice using the IB usability method. This involved the attendees, either individually or in pairs, analysing one of two video clips of lawyers performing behaviour-focused tasks using one of the firm's electronic legal resources. Most of the practice sessions were recorded. Attendees were also issued with a short questionnaire and participated in two separate focus groups, all with questions aimed at determining how useful, usable and easy to learn they consider the methods to be and how likely they are to use the methods in future.

Our findings relating to both methods were positive overall. In the focus groups, attendees highlighted the benefits of a user-centred and behaviour-focused approach for both the IB functionality and usability methods. An attendee also mentioned that such an approach might aid his team in 'getting attention' from senior management regarding usability issues and securing funding to address these issues. In the questionnaires, the vast majority of attendees deemed both of the methods to be 'somewhat' easy to use, easy to learn and useful and said they would be 'somewhat' likely to recommend that they or their colleagues use the methods in future.

The data collected also suggested useful ways that the methods might be improved. This included ensuring that there were more video-based examples in the IB functionality method to help people learning the method understand some of the more rarely-displayed levels of particular behaviours and replacing some of the behavioural theory with concrete examples in order to help people learning the method 'get to grips' with it faster (and without the need to refer to supporting examples) and to avoid confusion related to the boundaries of similar behaviours. Potential improvements to the IB functionality method also included asking summary questions about the necessity for supporting particular behaviours/levels at the *end* of an IB functionality evaluation (as opposed to throughout the evaluation for each behaviour/level). It was also suggested that IB functionality evaluations might be conducted by individuals as well as groups (although the results could be discussed with team members). It was suggested that the IB usability method might be improved by changing the order of some of the columns and standardising the 'severity' and 'easy of addressing' scales so that they match similar scales used by the firm when conducting usability evaluations. One attendee also suggested that future tutorials should include a practice session aimed at *devising* as well as analysing behavioural-focused think-aloud tasks. Finally, it was suggested that an electronic version of both the IB functionality and IB usability methods could be developed that would reduce paperwork, whilst making dynamic sorting and comparison of usability data easier. Many of these suggestions have been incorporated into the versions of the IB methods presented in this thesis.

Overall, we are encouraged by the tutorial attendees' generally positive outlook on both the IB functionality and IB usability methods. The tutorial session was well received and the hands-on examples and practice sessions led to a high level of engagement amongst participants with the IB methods. The data obtained from the formative evaluation session, particularly the focus group data, was not only useful in helping us to determine how useful, usable, learnable and likely to be used the IB methods were considered to be amongst LexisNexis Butterworths stakeholders, but also provided useful suggestions for how the method and future tutorial/evaluation sessions could be improved. In particular, the need to evaluate the task-setting part of the IB usability method was emphasised. This is something that was not considered as a priority for this formative evaluation session (partly due to the time constraints involved in teaching and evaluating the IB methods as part of a one-day tutorial). However, in hindsight, we believe that gaining feedback on the task-setting part of an IB

usability evaluation would be particularly useful for the future development of the IB usability method. This is discussed as a potential direction for future research in section 8.3.3.

We have demonstrated that the IB functionality and usability methods have the potential to make an important impact in practice. The IB methods contribute to the theory base of user evaluation methods, providing two specialist methods for evaluating electronic resources (as opposed to other types of interactive system). The IB methods have also been demonstrated to be potentially useful for practitioners, who are responsible for developing electronic resources.

We suggest that one of the main benefits of using information behaviours as a framework for evaluating electronic resources is the *structure* that they provide (an assertion echoed by some of our evaluation participants). Indeed, we hypothesise that the structure provided by the IB usability method is likely to result in the identification of a wider range of usability issues as compared with conventional user testing (or, at the very least, should lead to the identification of different types of issues). Similarly, we hypothesise that the structure provided by the IB functionality method is likely to result in the consideration of a wider range of resource functionality as compared with an unstructured functionality inspection. These are hypotheses that might be tested as part of further evaluations of the IB methods. In our final chapter, we discuss several options for further evaluation of the methods, along with other directions for future work connected to this thesis.

Chapter 8: Conclusion

This chapter at a glance…

In this chapter we:

- Summarise this thesis and its contributions to research.
- Discuss the implications of our work for shaping future studies of information behaviour and for informing the design and evaluation of electronic resources.
- Discuss the potential for future work in the areas of legal information behaviour, behavioural research evaluation and information behaviour in general.

8.1   Summary of thesis and research contributions

In this final chapter, we summarise this thesis and, in doing so, revisit each of its research contributions. This is followed, in section 8.2, by a short discussion of the implications of our thesis work for shaping future studies of information behaviour and for informing the design and evaluation of electronic resources. In section 8.3, we discuss areas related to this thesis where there is potential for future work and, finally, in section 8.4 we conclude by returning to the ethos that motivated our work in the first place.

The first part of our thesis work involved conducting an empirical study into academic and practicing lawyers' information behaviour. This study was novel for two main reasons. Firstly, there had not been many previous studies of lawyers' work with an aim of improving the interactive systems that they use (and even fewer studies aimed at improving the design of electronic legal resources). Secondly, the study employed an observational data collection and analysis approach that was useful in identifying a number of information behaviours that lawyers perform. Whilst many of these behaviours had been previously identified in other disciplines, the observational approach that we followed had not (previous studies had followed an interview approach). Our empirical study also made a number of theoretical contributions, validating Ellis's model in the new academic and practical domains of law, broadening the scope of the model by covering information *use* as well as information-seeking behaviours, enhancing the potential analytical detail of the model and extending the model to include behaviours pertinent to legal information-seeking.

The second part of our thesis work involved developing and formatively evaluating two novel methods for assessing the functionality and usability of electronic legal resources. This work is novel as, to the best of our knowledge, no other evaluation methods exist with roots grounded in Information Science theory. Nor do any other methods exist that were developed specifically with electronic resource evaluation in mind (although some have been adapted to the area, as discussed earlier). The formative evaluation of the IB methods involved teaching both methods to a small group of stakeholders working for an electronic resource development firm, allowing them to practice using the methods and recording some of the participants during the practice sessions. The evaluation also involved asking focus group and questionnaire questions on how usable, useful and learnable participants considered the methods to be and how likely they were to use them in future.

8.2   Implications of our work

Our thesis work, and in particular the research contributions it makes, have important implications for shaping future studies of information behaviour and for informing the design and evaluation of electronic resources. In this section, we discuss these implications in relation to the work we have carried out.

The first implication of our work is that **taking a broad, naturalistic approach to understanding information behaviour can be useful** when aiming to obtain data can be used to inform the design or improvement of interactive systems. The data obtained from our empirical study (presented in chapter 5) was rich and potentially useful for directly informing the design of the electronic legal resources (even though we chose to inform the design of a pair of evaluation methods instead). The broad and naturalistic nature of the study resulted in data that has the potential to be analysed under many different theoretical lenses, all with an aim of informing the design or improvement of electronic legal resources.

There are also other important implications of our work related to the methodology employed in our empirical study of lawyers (described in chapter 4). Firstly, **observing a wide range of lawyers, using a variety of different electronic resources can be useful** when trying to identify a wide range of information behaviours that can be generalised across resources. The decision to observe a vertical slice of both academic and practicing lawyers (where many of the academic students and staff had different specialisms and the practicing lawyers worked in two contrasting departments) resulted in rich, generalisable data. This was aided by the broad and flexible nature of the information task given during the observation, which did not restrict lawyers to using particular electronic resources. Secondly, **when observing information behaviour to inform the design or improvement of interactive systems, it may be necessary to adopt and adapt existing methodologies**, deviating from the classic use of particular methods. As discussed in chapter 4, we found it necessary to stop short of generating a full Grounded Theory as our aim was not to generate theory per se, but to inform the improvement of electronic resources. Similarly we asked participants to think aloud whilst using electronic resources to find information required for their work (which is not normally part of a Contextual Inquiry, but something we deemed necessary in order to generate data that could inform resource improvement). We also asked short probing questions whilst participants were thinking aloud (again not part of a traditional verbal protocol analysis, but also necessary to ensure richer data that could inform the improvement of electronic resources).

The analysis of our empirical data also suggests an important implication of our work. As discussed theoretically in chapter 3 (by using a variety of models to analyse an example information-seeking episode) and highlighted by the emergence of information behaviours related to those found by Ellis and his colleagues in other disciplines (see chapter 5), **some information-seeking models are likely to be more useful than others when used as lenses for the analysis of information behaviour**, particularly when aiming to inform the design or evaluation of electronic resources. As we discussed in chapter 3, we believe Ellis's model was particularly suitable for this purpose as it is based on observable behaviours rather than cognitive or physical processes that demand a further level of filtration to the underlying behaviour level in order to inform design.

Similarly, and related to the behaviours identified in our study, two more important implications of our work are that **information behaviours can provide a useful framework for making design suggestions based on supporting each behaviour**, and **can provide a useful framework for evaluating the functionality and usability of electronic legal resources**. Our thesis has focused on the second of these implications. Chapters 6 and 7 detail the development and evaluation of the Information Behaviour methods – two novel methods for evaluating the functionality and usability of electronic legal resources. Both of these methods are theoretically underpinned by the information behaviours identified in our empirical study and, essentially, use the behaviours as an evaluation framework. We also highlight that **it is potentially useful to document the development of evaluation methods** as very few published accounts exist of how particular methods were developed and how they evolved or changed since their initial development. The development and early testing of the IB methods is presented in chapter 6. Finally, and also related to the IB methods, another important implication of our work is

that **there is potential for evaluation methods that are theoretically underpinned by empirically observed information behaviours to be regarded as useful, usable, easy to learn and likely to be used in future**. A similar case can be made for methods aimed specifically at evaluating electronic resources (as opposed to other types of interactive system). Chapter 7 discusses the evaluation of both IB methods, which resulted in positive feedback related to each of the above 'success factors' (and a number of useful suggestions for improving the methods, which were incorporated into the versions presented in this thesis). However, a sizeable and long-term research challenge still exists in ensuring methods such as ours successfully transfer from theory to practice.

8.3   Potential for future work

Now that we have explored some of the implications of our thesis work, we briefly re-cap on the work carried out with the aim of highlighting the potential for future work. This results in the identification of three broad categories of potential future work, which we discuss in turn.

The first part of our thesis work examined the information behaviour of lawyers with the aim of improving the electronic resources that they use. However, there is also scope for **studying information behaviour in non-legal domains**. In addition, in order to identify information behaviour, we followed an observational data collection and analysis approach (where previous studies had followed an interview approach). This suggests the possibility of **studying the information behaviour of people working in other domains by employing a similar methodological approach**. As some of the identified information behaviours were unique to information-seeking in the legal domain, this suggests the possibility of **identifying behaviours unique to other domains** by conducting similar studies in future. It is, however, not necessary to focus outside the domain of law. There is also a wide range of possibilities for further work **examining various aspects of lawyers' information behaviour**.

The second part of our thesis work fed our findings related to lawyers' information behaviour into the development of the Information Behaviour evaluation methods. As the methods were developed using the empirically grounded data, this illustrates the potential for **studies to tailor the IB methods to other domains**, aimed at providing useful data for people who wish to use the IB methods outside of the legal domain. Also, although the IB methods were formatively evaluated with a group of stakeholders working for LexisNexis Butterworths, there is plenty of scope for **further developing the methods** based on participants' comments and for **conducting further evaluations** of both methods. These evaluations would not only serve to evaluate future versions of the methods but also to evaluate the methods in new ways.

In summary, the future work described above can be divided into three broad categories:

1.   Studying information behaviour in non-legal domains.
2.   Further examining lawyers' information behaviour.
3.   Conducting further development and evaluation of the IB methods.

In the remainder of this section, we discuss examples of potential future work in each of the above categories.

8.3.1   Potential for studying information behaviour in non-legal domains

There is plenty of potential to study information behaviour in domains outside law, particularly in domains where information seeking and use is important. These domains vary in nature. In the domain of medicine, doctors and other medical staff rely on information about diseases and how they can be treated. Information is particularly important in this domain in order to ensure that diagnoses are error-free and treatments are appropriate and, most of all, safe to the patient. In the humanities domain, historians rely on primary and secondary sources of information to piece together events that have happened in the past. Information is particularly important for historians so that they can ascertain the authority and reliability of these sources (see

Newman, Rimmer and Warwick, forthcoming). In a similar way to historians, journalists rely on information to piece together a story. Information is particularly important for journalists to help them form a viewpoint on particular events, hopefully leading to a story with an interesting 'angle,' or even to an exclusive story.

In general, there is not much work investigating information behaviour with the aim of *informing the design or improvement of electronic tools to support information work* (i.e. where the output includes design recommendations, working tools/prototypes or, as in our thesis work, techniques aimed at improving the design of electronic tools). However, the above examples come from domains where exceptions exist. Adams, Blandford and Attfield (2005) present an overview of a number of information studies they conducted across the UK National Health Service and discuss issues surrounding implementing electronic resources aimed at satisfying clinicians' and patients' varying information needs. Both Rimmer et al. (2008) and Newman et al. (Forthcoming) interview humanities researchers (including historians) and discuss the role of information in humanities research. Similarly Butterworth (2006) discusses personal historical research as a leisure activity. In addition, Attfield et al. (2008) observed journalists working in a large UK newsroom and fed their findings into the design of NewsHarvester, an electronic tool designed to support information-seeking in the context of writing. In this thesis, we have also discussed studies by David Ellis and his colleagues, mostly in the science domain, which have resulted in design recommendations for electronic resources.

These, however, are only a small number of studies in domains where there are still numerous possibilities for further research. In medicine, it would be interesting to examine how the information needs of medical students (who may have a great deal of medical information, but lack experience of applying this information) change as they progress through medical school into clinical practice and how new electronic tools might support (or existing tools might better support) the *use* of information in practice. Similar work might also be conducted into patients' changing information needs (for example, as they become more familiar with a particular illness they have been diagnosed with). It is also possible to design to support the *development* of information skills in the medical or other domains (which is something we discuss, but do not directly focus on in this thesis).

In the humanities domain, it would be interesting to examine information-seeking in the wider context of writing (just as Attfield et al. did with journalists). We believe it would be fruitful to observe literature academics and/or historians finding information to form or support a viewpoint. For historians, this may be a viewpoint on events – what happened and why, or a viewpoint on the reliability of a particular historical source (which might itself be information that has been sought during information-seeking). For both historians and literature academics, this may be a viewpoint on a debate. Insights gained could then be used to design electronic tools to support the process of using information to form viewpoints on issues, sources or debates. In the media domain, it would be interesting to examine how journalists decide between potential 'angles' on a story and whether there is potential to design electronic tools to support this process.

There are also a number of other domains where there is potential for future information research (with an HCI design focus). Kuhlthau (1997, 1999) observed the information work of a Securities Analyst over a five year period and whilst she does not make any explicit design recommendations, her work can be used to ask questions such as 'how can we design to support the development of expertise in this area?' Like journalists, Financial Investment Managers also use information to obtain an 'angle' to use in order to justify their opinions on a company's investment potential (Kuhlthau, 1999). It would therefore be interesting to examine how electronic tools might support the process of obtaining such an 'angle.' It would also be interesting to examine the concept of obtaining an 'angle' in order to pitch advice in other financial professions, such as Chartered Management Accountancy (where accountants provide management with advice on the potential profitability of different business strategies) and in non-financial professions. This includes any profession where information is sought and advice is given to guide a business to achieve its objectives (e.g. Management Consultancy, Marketing). Whilst these examples have focused on information behaviour in practice, there is also scope to study academics in these and other fields. Indeed, this might include fields where information seeking and use is important in academia, but not so important in practice (for example, imagine a Civil Engineer writing an academic essay comparing different methods of building a bridge or a Social Worker writing an essay on how to approach a sensitive family issue). Whilst studying these professions might yield little to support these

professionals' work in practice, useful insights could still be gained for how to improve existing electronic resources designed to support these professions. Such an approach might also provide insights into information behaviour amongst professions that might otherwise have been neglected (e.g. because information seeking and use is not of great importance to them).

In many of the domains and professions discussed above, however, there is an important need for information, and for electronic research. In these professions, it is also possible to take a similar behavioural approach to that described in this thesis (i.e. to identify the information behaviours performed in a particular profession). Findings from these observations can either be used to suggest design improvements for existing electronic resources or to highlight gaps in support for particular behaviours. Findings can also be fed into domain-tailored versions of the Information Behaviour methods, where the behaviours used to evaluate the functionality and usability of electronic resources are replaced or supplemented by those behaviours identified.

8.3.2   Potential for further examining lawyers' information behaviour

There are a number of potential sets of empirical studies that could be conducted that relate to the legal information behaviour work conducted in this thesis (and we believe would advance the current state of research in this field). All of these studies share the common theme of gaining a deep understanding of lawyers and their work in order to feed this understanding into the design of interactive systems to support their work. The first set of studies involves taking an *electronic resource focus*. It is possible to examine lawyers' use of particular electronic legal resources and use the findings to make usability-related suggestions (i.e. to suggest things that they find easy and difficult when using the resource and ways that the resource can be improved in order to make it easier to use). This is not only a possible suggestion for academic research, but could also be encouraged amongst firms that develop electronic legal resources (many of which already conduct some form of usability testing on their products). These findings can then feed into the design of tools to support difficult aspects of information-seeking (e.g. tools to support searching) or into design improvements for the electronic resource that was used. Also with a resource focus, it might be possible to examine how particular resources support a range of different types of information task. Lawyers could be asked to perform an information task of their choice (whether that be an active searching task, or a more passive monitoring task) using the resource. Insights on which types of task the resource supports well or not so well (and potentially which types of task it does not support at all) could be used to inform design improvements to the resource.

The second set of studies involves taking an *information task focus*. One option might be to examine particular legal information tasks in greater detail and feed this understanding into the design of interactive systems to support these tasks. In particular, we suggest the need to examine broader information use and re-use tasks (such as the re-use tasks examined in Blomberg et al.'s 1996 study of lawyers which led to the design of an electronic filing cabinet system). An information task focus has been adopted in UCL Interaction Centre's Making Sense of Information (MaSI) project, which has examined tasks such as the creation of a research note and monitoring/alert-related tasks (by observing practicing lawyers from the same London law firm as those in our study). There may be scope for observing tasks that the MaSI project does not focus on (for example writing internal knowledge documents as opposed to research notes and preparing client briefing documents or court documents). The flipside of taking an information task focus might involve the exploration of the potential utility, usability and usefulness of future legal technologies to support particular tasks (such as bundling information for presentation in court).

The third set of studies involves taking an *information behaviour* focus. This might involve re-examining legal information work with a particular focus on certain information behaviours (e.g. updating and history tracking, which we found to be particularly important for lawyers). For example, it is possible to observe lawyers as they perform tasks related to these behaviours and to feed findings into the design of new or existing electronic tools focused on supporting these behaviours. If examining updating and history tracking behaviours, for example, it might be useful to feed findings into the improved design of Current Legal Information (an electronic resource citator that focuses on supporting these behaviours). Although it is not possible to 'restrict' users to perform

particular behaviours, setting tasks focused on particular behaviours has the potential to encourage the display of those behaviours (as we saw with many of the tasks in our user pilot studies for the IB usability method). In a similar way to taking a task focus (and almost identical to conducting an IB usability evaluation), it is also possible to observe lawyers using a particular electronic resource to perform a range of behaviour-focused tasks and to make design improvement suggestions based on the findings (which may or may not be recorded and analysed using an IB usability evaluation structure).

The fourth set of studies involves taking an *information skills development* focus and involves gaining a detailed understanding of the development of lawyers' information skills as they progress through academia and into practice. This understanding can then be fed into the design of electronic resources or other tools that can help lawyers develop their skills, or resources that are adaptable so that functionality can be reduced (aimed at simplifying the resource at the cost of a reduction in 'power') or increased (aimed at providing increased 'power' at the cost of a potential increase in complexity). One example of an adaptive interface is Carroll's Training Wheels (see Carroll and Carrithers, 1984), where functionality is slowly introduced to the user (aimed at helping them to get to grips with the software).

The fifth and final set of studies involves *taking two different foci in order to compare different contexts in which legal work takes place*. For example, it is possible to examine the information behaviour of lawyers who work in different specialist legal departments. In the empirical study of lawyers' information behaviour described in this thesis, we observed practicing lawyers from one mainly contentious and one mainly non-contentious department of a large London law firm. It is also possible to examine commonalities and differences between the information behaviour of lawyers across a wider range of legal specialisms. Findings on how lawyers' information behaviour differs across specialisms could be fed into the design or tailoring of electronic tools to support a particular specialism (and any unique features of that specialism that have been identified). In a similar vein, it would be interesting to examine the information behaviour of practicing lawyers within a smaller UK firm than the one featured in our study, or academic lawyers from a different academic institution (perhaps a non-red-brick university such as a former polytechnic). It might also be interesting to examine differences across countries, perhaps by observing lawyers working in a country with a different legal system than the UK, or even in a country with a different legal culture. This with the view of noting any differences in the information behaviour displayed by these lawyers as compared with the lawyers in the empirical study described in this thesis and making design suggestions or designing tools that are sympathetic to these differences. These comparative studies can also be combined with ideas already discussed in this section (i.e. with a resource, task, or behavioural focus). For example, it is possible to examine how different legal specialisms use a particular electronic resource, how particular types of legal information work vary across UK firms of different sizes or the behavioural differences associated lawyers working with different legal systems (e.g. whether and how lawyers working with different legal systems perform 'updating' and 'history tracking' behaviours and what differences exist).

8.3.3 Potential for conducting further development and evaluation of the IB methods

Now that we have discussed the potential for studying information behaviour in both legal and non-legal domains, we turn to discuss the potential for conducting further development and evaluation of the IB methods. Firstly, there is the potential for investigating the applicability of the IB methods to non-legal domains. This would involve evaluating electronic resources from other domains (such as those we have discussed in this chapter) using the IB functionality and usability methods and determining how well they fare in evaluating non-legal resources. One option would be to use the methods to evaluate resources from professional domains similar to law (e.g. finance) and from domains where electronic resources differ considerably to legal ones (e.g. medicine or humanities). It is also possible to examine whether the principles of observing information behaviours apply to non-information-based domains (and therefore whether similar methods to the IB functionality and usability methods might, in future, be developed to evaluate *non* information-based systems). This suggestion comes in light of our evaluation with stakeholders working from LexisNexis Butterworths, as the two tutorial participants who did not work with information-based products on a regular basis suggested that

a behavioural approach might still be useful to them (although they would need to empirically identify behaviours using their database systems in order to be sure).

There is also potential to evaluate the IB methods that were developed as part of this thesis in additional, alternative ways. These include:

- Identifying what types of usability insights the IB usability method provides as compared to other methods in order to provide empirical data to support our hypothesis that the IB method fills a niche in the current 'market' for user evaluation methods. Whilst it is impossible to conduct a 'clean' comparison between usability evaluation methods, it is still possible to ask a number of different evaluators to use a variety of methods to evaluate a particular interactive system (see Blandford, Hyde et al.'s 2008 study comparing the output provided by a range of different evaluation methods). Ideally, the evaluators should have similar levels of familiarity with the resource under evaluation and similar amounts of experience using each method (although these factors are notoriously difficult to control). Rather than 'count' usability issues identified this study would aim to ascertain what different *types* of usability issue are uncovered by each method and whether the IB usability method does, as we predict, highlight certain issues that other methods do not. This suggestion is aimed at evaluating the *validity, reliability* and *scope* of the method. These three possible requirements for evaluation methods, along with the others mentioned in this section are discussed in chapter 6.
- Determining whether any differences exist in the *types* of usability issues identified when using each of the different versions of the IB usability method (i.e. the core, recommended and custom versions). This is similar to the suggestion above, and also requires caution when making comparisons between versions of the method as it is difficult to control evaluator differences (Hertzum and Jacobsen, 2001). This suggestion is aimed at evaluating the *scope* of the methods and the *insights derived* from their use.
- Examining the factors that can influence a functionality and/or usability IB evaluation. It would be interesting to examine whether the guidelines provided in this thesis make for successful functionality and usability evaluations. As this is a potentially mammoth undertaking, it will be necessary to isolate elements of the guidelines to 'test.' For example, it is possible to teach the IB functionality method to two groups of evaluators, enforcing one of the groups to use the IB guidance materials (which list illustrative ways of supporting each behaviour/level that they will be evaluating electronic resources in relation to) and allowing the other group to perform the evaluation without the examples (in a similar way to our formative evaluation session, where the examples were not used at all by participants, even though their use was encouraged). The difficulties the evaluators face when conducting the evaluation, and the output of the evaluations themselves could then be examined across both groups. It would also be interesting to examine what issues exist when conducting IB evaluations within particular organisational settings and cultures and how the method can be improved in light of any issues identified. This is another mammoth undertaking and would probably involve months of exploratory work – applying the IB methods in a variety of different organisations and drawing together some insights. This suggestion is aimed at evaluating the *reliability* of the methods.
- Evaluating the task *setting* as opposed to analysis aspects of the IB usability method. This was suggested by participants in our one-day tutorial and might involve conducting a follow-up tutorial with the participants (or conducting a tutorial with another firm centred on the IB usability method rather than both methods). Participants might be asked to set different behaviour-focused tasks for different evaluation purposes. These tasks can then either be reviewed by intended or actual users of the resource to ensure clarity, or performed by users (and the participants who set the tasks asked to analyse them, as they would do when actually conducting an IB usability evaluation for real). This suggestion is aimed at gaining additional insights into the *usability* (and potentially the *learnability)* of the method.
- Examining broader issues surrounding how evaluation methods can transfer successfully from theory to practice. For example, it is possible to examine evaluators' development of expertise using the IB methods by conducting a longitudinal study tracking the use of both methods over an extended time period. This suggestion is aimed at evaluating the long-term *utility* of the method and findings might be used both to improve the methods and to enrich our understanding of method adaptation and adaptation in practice.

All of these suggestions have the potential of improving the IB methods and contributing to the field of HCI method development research – a field with few documented method evaluation studies.

8.4   Conclusion

To conclude this thesis, we return to the ethos that motivated our work in the first place – the ethos that in order to design interactive systems that truly support users and their information work, it is necessary to understand these users and the work that these systems might be designed to support. The need to understand users in order to inform the design and evaluation of interactive systems has driven HCI research for decades. There has also been increased recognition of the need to understand users' work (i.e. the context in which these systems are used) amongst the wider digital library community (see Blandford and Gow – Eds., 2006). Our work can be regarded as a successful implementation of this user and work-centred ethos, bringing us full circle from understanding the use of interactive systems in the context of users' work, to informing the development of the Information Behaviour methods, to evaluating those methods to ensure that they truly meet the needs of the people they were developed to support. Indeed, the stakeholders working for LexisNexis Butterworths (which we regarded as representative potential users of the methods) regarded the methods to be useful, usable and easy to learn. They also suggested they would be likely to use the methods in future.

We must not forget, however, that the real benefit of ensuring that evaluation methods are useful, usable, learnable and used is that they might result in the design of new systems, or improvement of existing systems to support lawyers and their work. This is the real acid test for the IB methods. Whilst this thesis goes some way towards achieving this goal, there is still more work required to ensure that the outputs from this thesis (i.e. both our empirical findings and the IB methods) filter into the design of electronic legal resources and result in systems that are better at supporting lawyers' work.

Indeed, whilst we may have travelled 'full circle' in our thesis work, this is not a journey with a definite ending point. We have not travelled far round the iterative loop of evaluation and refinement whilst developing and evaluating our IB methods and therefore there is plenty of scope for future work in this area. In addition, the task of understanding users and their work is never complete: Just as the introduction of new technology can lead to changes in users' work tasks, changes in the nature of users' work over time can result in the need for new technology to support this work. Both can lead to new information behaviour displayed by users when performing these tasks, and to new ways of supporting these behaviours when designing (and evaluating) systems aimed at supporting this behaviour.

We embrace the ongoing challenge of understanding users and their work and using this understanding to help design and evaluate interactive systems that support this work. Moreover, we believe it is important for the disciplines of HCI, Information Science and Digital Libraries to face this challenge together, by encouraging further research that builds bridges between these disciplines and by never losing sight of the importance of ensuring that interactive systems are designed to truly support users and their work.

References

Adams, A. & Blandford, A. (2005). Digital Libraries' Support for the User's 'Information Journey.' In proceedings of the ACM/IEEE Joint Conference on Digital Libraries, 2005, pp. 160-169. ACM Press, New York, USA.

Adams, A., Blandford, A. & Attfield S. (2005). Implementing Digital Resources for Clinicians' and Patients' Varying Needs. In proceedings of BCS Healthcare Computing 2005. pp. 226-233. Also in Medical Informatics and the Internet in Medicine. 30(2), pp. 107-122.

Andrews, C. (1993). User Perceptions of CALR. Unpublished MSc thesis, City University, London, UK.

Attfield, S., Blandford, A., Dowell, J. & Cairns, P. (2008). Uncertainty-Tolerant Design: Evaluating Task Performance and Drag-Link Information Gathering for a News-Writing Task. International Journal of Human-Computer Studies 66(6), pp. 410-424.

Bainbridge, L. & Sanderson, P. (1995). Verbal Protocol Analysis. In Wilson, J. and Corlett, E. Evaluation of Human Work. Erlbaum Associates, London, UK, pp. 169-201.

Bates, M. (1989). The Design of Browsing and Berrypicking Techniques for the Online Search Interface. Online Review 13(5), pp. 407-24.

Bates, M. (2002). The Cascade of Interactions in the Digital Library Interface. Information Processing and Management 38, pp, 381-400.

Belkin, N. (1980). Anomalous States of Knowledge for Information Retrieval. Canadian Journal of Information Science 5, pp. 133-143.

Bell, B. (1992). Using Programming Walkthroughs to Design a Visual Language. Unpublished PhD Thesis. University of Colorado at Boulder. Also available as University of Colorado at Boulder working paper CU-CS-581-92.

Bellotti, V., Buckingham Shum, S., MacLean, A. & Hammond, N. (1995). Multidisciplinary Modelling in HCI Design… In theory and in Practice. In proceedings of CHI '95 ACM Conference on Human Factors in Computing. pp. 429-436. Boulder, Colorado, USA. ACM Press.

Beyer, H. & Holtzblatt, K. (1998). Contextual Design: Defining Customer-centred Systems. Morgan Kauffman. London, UK.

Blandford, A., Adams, A., Attfield, S., Buchanan, G., Gow, J., Makri, S., Rimmer, J. & Warwick, C. (2008). The PRET A Rapporter Framework: Evaluating Digital Libraries from the Perspective of Information Work. Information Processing and Management 44(1), pp. 4-21.

Blandford, A., Fields, B. & Keith, S. (2003). Usability Evaluation of Digital Libraries. Tutorial Presented at the ACM/IEEE Joint Conference on Digital Libraries. May 27-31. Houston, Texas, USA. Published as Middlesex University Technical Report IDC-TR-2003-001. Available online: www.cs.mdx.ac.uk/research/idc/papers/IDC-TR-2003-001.pdf [Accessed 12/07/07].

Blandford, A., Gow, J., Buchanan, G., Rimmer, J. & Warwick, C. (2007). Creators, Composers and Consumers: Experiences of Designing a Digital Library. Baranauskas, C. et al. (Eds.). INTERACT 2007, Lecture Notes in Computer Science, Part I, pp. 239-242.

Blandford, A. & Green, T. (2008). Methodological Development. In Cairns, P., & Cox, A. (Eds.). Research Methods for Human-Computer Interaction (1st edition). Cambridge University Press. Cambridge, UK.

Blandford, A., Green, T., Furniss, D. & Makri, S. (2008). Evaluating System Utility and Conceptual Fit Using CASSM. International Journal of Human-Computer Studies 66(6), pp. 393-409.

Blandford A. & Gow, J. (2006). (Eds.). Proceedings of the 1st International Workshop on Digital Libraries in the Context of Users' Broader Activities (DL-CUBA), 15th June 2006. JCDL 2006, Chapel Hill, NC, USA.

Blandford, A., Hyde, J., Green, T. & Connell, I. (2008). Scoping Usability Evaluation Methods: A Case Study. Human-Computer Interaction Journal 23(3), pp. TBC.

Blandford A., Keith, S., Connell, I. & Edwards, H. (2004). Analytical Usability Evaluation for Digital Libraries: A Case Study. In proceedings of the ACM/IEEE Joint Conference on Digital Libraries, pp. 27-36. June 7-11th 2004. Tucson, Arizona, USA.

Blandford, A., Keith, S. & Fields, B. (2006). Claims Analysis "In the Wild:" A Case Study on Digital Library Development. International Journal of Human-computer Interaction 21(2), pp. 197-218.

Blandford, A., Keith, S., Fields, B. & Furniss, D. (2007). Disrupting Digital Library Development with Scenario Informed Design. Interacting with Computers 19(1), pp. 70-82.

Blomberg, J., Suchman, L. & Trigg, R. (1996). Reflections on a Work-oriented Design Project. Human-computer Interaction, 11, pp. 237-265.

Boren, M. & Ramey, J. (2000). Thinking Aloud: Reconciling Theory and Practice. IEEE Transactions on Professional Communication 43(3), pp. 261-278.

Buckingham Shum, S. & Hammond, N. (1994). Transferring HCI Modelling and Design Techniques to Practitioners: A Framework and Empirical Work. In G. Cockton, S. Draper and G. Weir, Eds. People and Computers IX: Proceedings of British Computer Society HCI'94, pp. 21-36. Cambridge, UK. Cambridge University Press.

Butterworth, R. (2006). Information Seeking and Retrieval as a Leisure Activity. In Blandford A. and Gow, J. (Eds.). Proceedings of the 1st International Workshop on Digital Libraries in the Context of Users' Broader Activities (DL-CUBA), 15th June 2006, pp.29-32. JCDL 2006, Chapel Hill, NC, USA.

Carroll, J. & Carrithers, C. (1984). Training Wheels in a User Interface. In Communications of the ACM. 27(8), pp. 800-806.

Cheatle, E. (1992). Information Needs of Solicitors. Unpublished MSc thesis, City University, London, UK.

Colbert, M., Peltason, C., Fricke, R. & Sanderson, M. (1997). The Application of Process Models of Information-seeking During Conceptual Design: The Case of an Intranet Resource for the Re-use of Multimedia Training Material in the Motor Industry. In proceedings of the conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques, pp. 73-81. August 18-20, Amsterdam, the Netherlands.

Cole, C. & Kuhlthau, C. (2000). Information and Information-seeking of Novice Versus Expert Lawyers: How Experts Add Value. The New Review of Information Behaviour Research 1, pp. 103-115.

Cunningham, S., Knowles, C. & Reeves, N. (2001). An Ethnographic Study of Technical Support Workers: Why We Didn't Build a Tech Support Digital Library. In Proceedings of the 1st ACM/IEE Joint Conference on Digital Libraries, pp. 189-198. Roanoake, Virginia, USA.

Dempsey, B., Vreeland, R., Sumner Jr, R. & Yang, K. (2000). Design and Empirical Evaluation of Search Software for Legal Professionals on the WWW. Information Processing and Management 36, pp. 253-273.

Dervin, B. (1983). An Overview of Sense-making Research: Concepts, Methods, and Results to Date. Paper presented at the meeting of the International Communication Association, May 1999, Dallas, TX, USA.

Dervin, B. (1992). From the Mind's Eye of the User: The Sense-Making Qualitative-Quantitative Methodology. In Glazier, J. & Powell, R. (Eds.). Qualitative Research in Information Management. Englewood, CO, USA. Libraries Unlimited.

Dumas, J. & Redish, J. (1999). A Practical Guide to Usability Testing. Intellect Books. Exeter, UK.

Elliott, M. & Kling, R. (1996). Organizational Usability of Digital Libraries in the Courts. In proceedings of the 29th Annual Hawaii International Conference on System Sciences, pp. 62-71, Hawaii, USA.

Elliott, M. & Kling, R. (1997). Organizational Usability of Digital Libraries: Case Study of Legal Research in Civil and Criminal Courts. Journal of the American Society for Information Science 48(11), pp. 1023-1035.

Ellis, D. (1987). The Derivation of a Behavioural Model for Information Retrieval System Design. Unpublished doctoral dissertation, University of Sheffield, UK.

Ellis, D. (1989). A Behavioural Approach to Information Retrieval System Design. Journal of Documentation, 45(3), pp. 171-212.

Ellis, D. (1993). Modeling the Information-seeking Patterns of Academic Researchers: A Grounded Theory Approach. Library Quarterly 63(4), pp. 469-486.

Ellis, D. (2006). Ellis's Model of Information-seeking Behaviour. In Fisher, F., Erdelez, S. and McKechnie (Eds.). Theories of Information Behavior. Asisandt Monograph Series. Information Today. Melford, New Jersey, USA.

Ellis, D., Cox, D. & Hall, K. (1993). A Comparison of the Information-seeking Patterns of Researchers in the Physical and Social Sciences. Journal of Documentation 49(4), pp. 356-369.

Ellis, D. & Haugan, M. (1997). Modelling the Information-seeking Patterns of Engineers and Research Scientists in an Industrial Environment. Journal of Documentation 53(4), pp. 384-403.

Ericsson, K. & Simon, H. (1984). Protocol Analysis: Verbal Reports as Data. MIT Press. London, UK.

Feliciano, M. (1984). Legal Information Sources, Services and Needs of Lawyers. Journal of Philippine Librarianship 8(12), pp.71-92.

Flanagan, J. (1954). The Critical Incident Technique. Psychological Bulletin, 51(4), pp. 327-358.

Ford, N., Wilson, T., Foster, A. & Ellis, D. (2002). Information-seeking and Mediated Searching. Part 4. Cognitive Styles in Information-seeking. Journal of the American Society for Information Science and Technology 53(9), pp. 728-735.

Garzotto, F. & Perrone, V. (2007). Industrial Acceptability of Web Design Methods: An Empirical Study. Journal of Web Engineering 6(1), pp. 73-96.

Glaser, B. & Strauss, A. (1967). The Discovery of Grounded Theory: Strategies for Qualitative Research. Chicago, USA. Aldine Publishers.

Hainsworth, M. (1992). Information-seeking Behaviour of Judges. Unpublished PhD Thesis, Florida State University, USA.

Haruna, I. & Mabawonku, I. (2001). Information Needs and Seeking Behaviour of Legal Practitioners and the Challenges to Law Libraries in Lagos, Nigeria. The International Information and Library Review, 33(1), pp. 69-87.

Hertzum, M. & Jacobsen, N. (2001). The Evaluator Effect: A Chilling Fact about Usability Evaluation Methods. International Journal of Human-Computer Interaction 13(4), pp. 421-443.

Howland, J. & Lewis, N. (1990). The Effectiveness of Law School Legal Research Training Programs. Journal of Legal Education, 40, pp. 381-391.

Ingwersen, P. & Järvelin (2005). The Turn: Integration of Information-seeking and Retrieval in Context. Springer, the Netherlands.

John, B. & Packer, H. (1995). Learning and Using the Cognitive Walkthrough Method: A Case Study Approach. In proceedings of the SIGCHI conference on Human Factors in Computing Systems, pp. 429-436. Denver, Colorado, USA.

Jones, Y. (2006). "Just the Facts Ma'am?" A Contextual Approach to the Legal Information Use Environment. In proceedings of the 6[th] ACM Conference on Designing Interactive Systems, pp. 357-359. University Park, PA, USA.

Kerins, G., Madden, R. & Fulton, C. (2004). Information-Seeking and Students Studying for Professional Careers: The Cases of Engineering and Law Students in Ireland. Information Research, 10(1), paper 208. Available online: http://InformationR.net/ir/10-1/paper208.html. [Accessed: 07/01/06].

Kidd, D. (1978). Legal Information of Solicitors in Private Practice. Unpublished MSc Thesis, University of Edinburgh, UK.

Komlodi, A. & Soergel, D. (2002). Attorneys Interacting with Legal Information Systems: Tools for Mental Model Building and Task Integration. In proceedings of the 65th Annual Meeting of American Society for Information Science and Technology, pp. 152-163. Philadelphia, USA. ACM Press.

Kuhlthau, C. (1988). Developing a Model of the Library Search Process: Cognitive and Affective Aspects. Reference Quarterly 28(2), pp. 232-242.

Kuhlthau, C. (1991). Inside the Search Process: Information-seeking from the User's Perspective. Journal of the American Society for Information Science 42(5), pp. 361-371.

Kuhlthau, C. (1997). The Influence of Uncertainty on the Information-seeking Behavior of a Securities Analyst. In proceedings of Information-seeking in Context, University of Tampere, Finland, August 1996, Taylor-Graham, pp. 268-274.

Kuhlthau, C. (1999). The Role of Experience in the Information Search Process of an Early Career Information Worker: Perceptions of Uncertainty, Complexity, Construction and Sources. Journal of the American Society for Information Science 50(5), pp. 399-412.

Kuhlthau, C. and Tama, S. (2001). Information Search Process of Lawyers: A Call for 'Just for Me' Information Services. Journal of Documentation, 57(1), pp. 25-43.

Landauer, T. (1995). The Trouble with Computers: Usefulness, Usability and Productivity. MIT Press. London, UK.

Leckie, G., Pettigrew, K. & Sylvain, C. (1996). Modelling the Information-seeking of Professionals: A General Model Derived from Research on Engineers, Health Care Professionals, and Lawyers. Library Quarterly, 66(2), pp. 161-193.

Polson, P., Lewis, C., Rieman, J. & Wharton, C. (1992). Cognitive Walkthroughs: A Method for Theory-Based Evaluation of User Interfaces. International Journal of Man-Machine Studies 36(5), pp. 741-773.

Lynch, C. & Garcia-Molina, H. (1996). Interoperability, Scaling and the Digital Libraries Research Agenda. Microcomputers for Information Management 13(2), pp. 85-132.

Mack, R. & Nielsen, J. (1994). Executive Summary. In Nielson, J. and Mack, R. (Eds.). Usability Evaluation methods. Wiley, New York, USA.

MacLean, A., Young, R., Bellotti, V. & Moran, T. (1991). Questions, Options, and Criteria: Elements of Design Space Analysis. Human-Computer Interaction 6(3 & 4), pp. 201-250.

Marchionini, G. (1995). Information-seeking in Electronic Environments. Cambridge University Press, Cambridge, UK.

Marchionini, G. (2007). Find What You Need, Understand What You Find. International Journal of Human-Computer Interaction 23(3), pp. 205-237.

Marshall, C., Price, M, Golovchinsky, G. & Schilit, B. (2001). Designing e-Books for Legal Research. In proceedings of the 1st ACM/IEEE-CS Joint Conference on Digital Libraries, Roanoke, Virginia, USA. 2001, pp. 41-48.

Meho, L. & Tibbo, H. (2003). Modeling the Information-seeking Behavior of Social Scientists: Ellis's Study Revisited. Journal of the American Society for Information Science and Technology 54(6), pp. 570-587.

Newman, W., Rimmer, J. & Warwick, C. (Forthcoming). The Problem of Follow-ups in Scholarly Research: Using Story Elicitation to Establish the Needs of Humanities Researchers. To appear in the Journal of Documentation.

Nielsen, J. (1993). Usability Engineering. Academic Press. New York, USA.

Nielsen, J. & Landauer, T. (1993). A Mathematical Model of the Finding of Usability Problems. In proceedings of ACM INTERCHI'93. Amsterdam, The Netherlands, 24-29 April 1993. pp. 206-213.

Norman, D. (1986). Cognitive Engineering. In Norman, D. and Draper, S. (Eds). User Centred System Design, pp. 31-61. Lawrence Earlbaum Publishers, New Jersey, USA.

Norman, P. (2004). The Big Match - Lexis v Westlaw. Legal Information Management 4, pp. 90-97.

O'Brien, M. & Buckley, J. (2005). Modelling the Information-Seeking Behaviour of Programmers – An Empirical Approach. In proceedings of the 13th International Workshop on Program Comprehension (IWPC'05), St. Louis, Missouri, USA. pp. 125-134.

Otike, J. (1999). The Information Needs and Seeking Habits of Lawyers in England: A Pilot Study. International Information and Library Review 31, pp. 19-39.

Oulanov, A. & Pajarillo, E. (2003). Academic Librarians' Perception of LexisNexis. The Electronic Library, 21(2), pp. 123-129.

Pirolli, P. & Card, S. (1999). Information Foraging. Psychological Review 106(4), pp. 643-675.

Rimmer, J., Warwick, C., Blandford, A., Gow, J. & Buchanan, G. (2008). An Examination of the Physical and Digital Qualities of Humanities Research. Information Processing and Management 44(3), pp. 1374-1392.

Smith, K. (1988). An Investigation into the Information Behaviour of Academics Active in the Field of English Literature. Unpublished Masters dissertation, University of Sheffield, UK.

Stelmaszewska, H., Blandford, A. and Buchanan, G. (2005). Designing to Change Users' Information-seeking Behaviour: A Case Study. In S. Chen and G. Magoulas (Eds.) Adaptable and Adaptive Hypermedia Systems, pp. 1-18. Information Science Publishing, London.

Strauss A. & Corbin J. (1998). Basics of Qualitative Research. Sage. London, UK.

Sutcliffe, A. & Ennis, M. (1998). Towards a Cognitive Theory of Information Retrieval. Interacting with Computers, 10, pp. 321-351.

Sutcliffe, A. & Ennis, M. (2000). Designing Intelligent Assistance for End-User Information Retrieval. In proceedings of OZCHI 2000, pp. 202-210. Sydney, Australia.

Sutcliffe, A., Ennis, M. & Watkinson, S. (2000). Empirical Studies of End-User Information Searching. Journal for the American Society for Information Science 51(13), pp. 1211-1231.

Sutton, S. (1994). The Role of Attorney Mental Models of Law in case Relevance Determinations: An Exploratory Analysis. Journal of the American Society for Information Science 45(3), pp. 186-200.

Van Den Haak, M., De Jong, M. & Schellens, P. (2003). Retrospective vs. Concurrent Think-Aloud Protocols: Testing the Usability of an Online Library Catalogue. Behaviour and Information Technology 22(5), pp. 339-351.

Vollaro, A. & Hawkins, D. (1986). End-user Searching in a Large Library Network: A Case Study of Patent Attorneys. Online 10(4), pp. 67-72.

Wharton, C., Reiman, J., Lewis, C. & Polson, P. (1994). The Cognitive Walkthrough Method: A Practitioner's Guide. In Nielsen, J. and Mack, R. (Eds.). Usability Evaluation methods. Wiley, New York, USA.

Wilkinson, M. (2001). Information Sources Used by Lawyers in Problem-solving: An Empirical Exploration. Library and Information Science Research, 23(2001), pp. 257-276.

Wilson, T. (1999). Models of Information Behaviour Research. Journal of Documentation, 55(3), pp. 249-270.

Wilson, T. (2000) Human Information Behavior. Informing Science, 3(2) (Special Issue on Information Science Research), pp. 49-55.

Wixon, D. (2003). Evaluating Usability Methods: Why the Current Literature Fails the Practitioner. Interactions 10(4) (July 2003), pp. 28-34.

Yuan, W. (1997). End-user Searching Behavior in Information Retrieval: A Longitudinal Study. Journal of the American Society for Information Science and Technology. 48(3), pp. 218-234.

Appendix 1: Illustrative transcript from our study of lawyers' information behaviour

Below we present an illustrative transcript from our study of academic and practicing lawyers' information behaviour. The transcript is of an Associate Tax lawyer, working at a large London law firm (participant T6-A), who is looking for information on a rather complex Capital Gains Tax issue, which she describes below. The Associate used the LexisNexis Butterworths electronic resource, along with the firm's Knowledge Management database (referred to in the transcript as 'the firm's KM database') to complete her information task, which was to 'find information required for her work' using the electronic resource or resources of her choice. The Associate conducted this task at her desk.

Questions and comments from the researcher, aimed at eliciting further details about the Associate's interaction with the electronic resources are denoted by 'R.' User interface actions are presented in square brackets. […] denotes where some participant comments have been omitted. Words in ***bold italics*** denote that emphasis was placed on them in the audio recording.

T6-A: Often on a takeover you'll be buying shares from members of the public and you also might be buying shares from Directors of the target company. And usually both the Directors and the Shareholders will be concerned with the tax treatment that they're going to get when they sell their shares. So say, for example, that they bought shares that were worth £1 and the bidder is now offering them £5 for those shares. £4 of that will be taxable, probably, and depending how you structure it, a greater or lesser part of it will be taxable. And one thing you can do to try to mitigate, or at least defer the amount of tax that you have to pay is, instead of paying them cash – instead of paying them £5 for their shares, you can offer them something called loan notes, which is basically an IOU that says 'we'll pay you £5 but we're not going to pay it right now, we're going to pay it in the future.' And the reason that the shareholders want that is that they can then spread their gain in the shares over a number of years and the government gives us all an annual exception from Capital Gains, which at the moment is about £9,000. So every year you can make £9,000 through selling shares or various other assets and you will only be taxed on gains above that £9,000. So, in the original example, they can sell £1 of each of their shares spread over 5 years, say, and that might mean that it all falls within their annual exemption and they don't pay tax on it, which is better. But in order to do that treatment, you have to make sure that the shares that they originally have are treated as the same assets, for Capital Gains purposes, as the loan notes that you're giving them instead. So I'm looking at a question that our Corporate department have given us to see [pauses]. Normally when you're doing a takeover, you'll make an offer to Shareholders and Directors and see how many people accept it. What they want to do, before they announce the market that they're planning to take over this company, they want the directors to ***already*** have sold their shares to them because, presentationally, that looks much better to the other Shareholders. So they want to know whether they can single out the Directors and offer them a special form of loan note at the stage when they announce the deal and then when they actually make the offer, they'll offer a different type of loan note to all the Shareholders and perhaps the Directors as well. So they've asked me to tell them whether, if they ***do*** offer these loan notes at the announcement stage to the Directors, whether that will work for the legislation and they'll get the exchange treatment and the deferral of their Capital Gains.

R: So it's not a question of whether they would be ***allowed*** to do that, but whether it will give them the tax treatment that they would prefer?

T6-A: Yeah, exactly. There are no restrictions on them offering the loan notes, they are perfectly free to do that. But there would be no point if they weren't going to get the favourable tax treatment as a result of this.

R: And what have you done so far with this?

T6-A: Nothing. I've asked for a few facts that I need to know about the percentage shareholding that the Directors have individually and as a collective group in the company because I know that's relevant for the particular section that I'm looking at. I'm now going to look at the bit of legislation in question, which I've looked at a bit before. But what I intend to do is have a look at some of the commentary sources that we have online and also our internal know-how and see what that comes up with.

T6-A: I've read through this section before [opens paper version of Tolley's and points to relevant legislation] which talks about the exchange process of shares into loan notes and then what I usually do is pick up on the Simon's Direct Tax reference because Simon's is the most authoritative Tax commentary and is always in your Statute books as

they're LexisNexis [points to LexisNexis logo at the front of the paper-based version of Tolley's] and Simon's is a LexisNexis publication as well, so they link in. [Loads LexisNexis Butterworths, selects the 'Sources' tab and browses for Simon's by clicking on 'S' in the alphabetical list, then selecting the source and clicking on the 'browse' button next to it]. Sources [pauses]. Simon's Direct Tax.

R: How do you know which section of Simon's to go into?

T6-A: Because I'm looking at section 135 and then 'commentary Simon's C2.726' [points to section reference in the legislation in Tolley's]. But sometimes I can't remember whether section 'C' is in 'binder 4' or 'binder 5' or whatever on the list [points to initial browsable list of Simon's headings which do not have section numbers and browses through a couple of headings]. So it turns out it's in binder 4. C2 [drills down on section C2], point 7 [drills down further to section C2.7], 26 [drills down yet further]. So that gives me the Simon's commentary and it's 726 and 727 [points to references in Tolley's] so it's both, so you just end up printing both of them [clicks on tick box next to both sections of Simon's and selects 'print' and then 'selected documents'].

R: Does that print both of them?

T6-A: Yes. It totally wastes paper this thing if you ask me, the way that LexisNexis does it. I'll print them out so that you can have a look. That prints them both because I've ticked them on the tree and done it that way.

T6-A: Then I guess, although I probably won't do this now, I'll read through this all and follow any of the footnotes to cases and if I thought that any of them were interesting or useful I would do into them and print them too. I very rarely read things straight from my screen, I don't really like doing that. So I might print it out if I thought it was a useful case.

T6-A: So that's where I would start. Then the next thing, which is always quite useful for Tax, is the Inland Revenue Manuals. The Revenue Manuals are only the Inland Revenue's interpretation of the legislation, but still they do give you a steer on how the Revenue is going to treat your case, basically, and there is seldom any point of saying 'we think the correct interpretation of this section of the legislation is x and we know that the Revenue thinks it's y, but never mind, we'll just do with x anyway.' That's a bit of a stupid position to adopt unless you really want to fight it. So you do have to look at the manuals.

T6-A: So usually what I would do is go back to 'Sources' and then select 'I' and 'Inland Revenue Manuals' and 'browse' and it tells you again here [points to reference to Inland Revenue manuals in Tolley's] which manuals you're looking at. So 'Capital Gains manual' CG. So I usually just open it up in the tree and then decide if I want the whole thing. So we've got C45550 [points to reference in Tolley's]. I'll probably start on the ones that are on C5. Again, it doesn't actually tell you the numbers here, but I'm guessing that it's going to be under 'Shares and Securities.' [Browses within 'Shares and Securities' section and eventually locates the reference being sought]. This is C52521 which was mentioned in Tolley's and I guess what I would do again is, if I thought it was going to be useful, is tick the whole of the section and print it out [does so]. I'll just go and get the first printout from Lexis to show you [walks to printer]. So this is the Simon's Direct Tax bit that we were looking at [flicks through printout] which always gives you an annoying cover page at the front. We printed one section after the other, so you'll see that it gives me first C2.726 which was the first one we looked at, and then [flicks through pages] C2.727. So I'd go through as much of these Revenue manuals that I thought was useful.

R: Is it typical that you would print a lot of these things out and come back to them later? Or would you read them as you got them?

T6-A: I'd probably print everything out that I thought was useful and then read it in a one-er, I think rather than print, read a bit, go on to the next bit, but I guess that's just the way I am. [Skim reads through the text of the Inland Revenue Manual references]. This strikes me as being useful on takeovers. [Reads document more thoroughly then browses through subsequent headings in the tree of Inland Revenue Manual headings].

R: What are you doing now?

T6-A: I'm just looking in case any of these individual sections really jump out at me as being potentially helpful for the question that I'm considering, which none of them especially so. So what I'd do is [pauses]. [Clicks on all sections within the main 'Shares and Securities' section of the Inland Revenue Manual], maybe do a search just to check [selects 'search source' from the drop-down menu on the results list].

R: What are you trying to search for?

T6-A: I'm trying to search just within wherever we were [pauses]. I'm never particularly sure how to do it. What I want to do is [pauses]. It's says 'search source' [pauses].

R: What does that 'search source' option do?

T6-A: At the moment it's only letting me search within the Inland Revenue manuals, which is way too big. I don't want to search all the manuals for my search term. That would be absolutely useless. I just want to search one very small bit of the manuals. I'll have to say that I actually don't know how to do that. So at this point I'd probably give up on that and go to [the firm's KM database] and have a look at what our in-house learning is on this [loads KM database].

T6-A: I've got my 'Tax' preference on, so I kind of know what I'm gonna get back. And we're looking at section 135 of the legislation and want it particularly in the context of loan notes and probably, as I said we're considering this specifically for the directors at the moment, I know from past experience from the way that usually things are written up in our Tax news that finds its way onto the Info bank, people probably refer to that as 'management.' [Searches for "I35" AND "loan note" AND management].

R: Why did you search for 'loan note' rather than 'loan notes?'

T6-A: Because then I get 'note' singular and plural.

R: When you said that you had your 'Tax' preference on and therefore you knew what that would bring back, what did you mean by that?

T6-A: Well in [the firm's KM database], you can create 'preferences' which are sort of refined searches. So for Tax, that's just all of the Tax documents, but I have lots of other preferences [loads preference list]. I've got a 'Corporate,' 'Corporate intra-group,' 'Tax with no news' and if I know that I'm looking for a point that I know is most likely to come up in the Corporate department, because it's their type of work, then I'll take my 'Tax' preference off and just search within 'Corporate.'

R: So these are preferences that you've set up previously?

T6-A: Yes, but more often than not it's just 'Tax' documents that we're looking at.

T6-A: Then you can have this set different ways [points to result sorting options]. I usually have mine prioritised by 'effective date,' so you get the most recent results first. Just because, I figure, if there's something relevant then I'd rather know what we were taking about recently because that is most likely to point to the current issues between our interpretation and maybe points that the Revenue that have come up with or tricky points that people have faced on recent deals. Particularly with takeovers, you often find that there's a sort of continual market evolution going on. So if you are doing a takeover in 2006, say, someone will pioneer a new way of doing the takeover. Then other people get to hear about that and they start saying 'I want to do our takeover the way the Vodafone takeover was done,' or something. And so it kind of develops from there and people start identifying the benefits and problems of doing it that new way and then usually they find ways of overcoming the problems and it all gets refined and goes along quite nicely until someone else comes up with an even brighter idea and says 'let's do it this way.' So that's why I usually go for the newest stuff first because, more than likely, the structure of the thing that I have been asked to look at is something that other people in the market might have looked at and there might be some learning on it.

R: And what does [the KM database] usually order results by, by default?

T6-A: I can't remember if, by default [pauses]. I don't think it's 'effective date,' I think it's maybe by 'level' which is [the KM database's] own view of how relevant the description that is given here [points to document summary in result list] is to the search terms that you've entered. But I guess, well, I'm just not so convinced about that. And you kind of think 'well if I haven't got my search terms exactly right and the person who's written the [KM database] just has a different view than me about how to enter search terms,' I'd usually just have mine on 'effective date.'

R: I guess the one other way of sorting results is by percentage? [Points to '%' option at the top of results list].

T6-A: Yeah. I'm not really sure what this does. [Hovers cursor over '%' option at the top of the results list and reads out tool-tip]. 'Sort descending.' I'm not sure, I've never used that before.

R: So sorting by date is more useful in general?

T6-A: For *me*, I prefer using date, yeah.

T6-A: Once you've got these results you can go to 'relevant law,' for example, which is not going to be particularly relevant in this case and actually, we always sort this fairly badly [points to text that indicates that most of the current results have been classified as 'Worldwide Law' on the system rather than 'UK law']. Most of these documents should be under 'UK law' because this section of legislation that we're looking at is UK statute, so I'm not sure why it says 'UK law' 1 result and 'Worldwide law' 10 results out of 11, because as far as I'm concerned it should probably be 'UK law' 11 results out of 11. But if I was doing something were I thought it was more general, it wasn't a specific piece of UK legislation, then I might say 'there's no point looking at the 'French law' results,' for example, 'I'll get rid of those.'

R: So these are further ways of narrowing down the results?

T6-A: Yeah, exactly. And you can also, for example, narrow down by 'resource type.' If I knew that someone had written a memo on this issue before – as often happens one of the Partners says to you 'oh I remember when we did that deal 5 years ago and James wrote a note on such and such, you'll probably be able to find it.' So if I knew that it was a memo that I was looking for, then I might just limit my search to the memos and hopefully that would come up with it. But I'm not really looking for anything that specific at the moment.

T6-A: So I would basically just start with the most recent result [clicks on document in results list and loads full-page summary] and then kind of find where, quite usefully, it highlights where your search terms are. I'd just have a look at those and see if what I was looking at was relevant to my question. And then, similarly, if I did, which say this one seems to be, I would [pauses]. I hate printing things from [the KM database] because I think it's done really badly. So I would, instead, take the number off the bottom of this [copies document number listed at the bottom of the summary into the computer's clipboard] and I would instead search for it on the firm's document database.

R: What's bad about printing from [the KM database]?

T6-A: Just because I never know what the printing options mean [reads out print options for full-page summary of document]. I mean 'everything' should theoretically mean everything, but 'business and legal commentary,' 'profile,' 'see also.' I guess they correspond to parts of the text on the screen, but I often find that if I print the 'commentary,' it just doesn't print what I'm expecting it to. I'll show you what happens. [Selects the print 'business and legal commentary' option]. I do want to print this document, but just not really in the [KM database] format. I think it's printed out to my default printer, so I'll just go and get it [walks to printer]. You see I pressed print 'commentary' and then it printed, but I would've thought that means it prints the page of text, but all I get is this [points to printout] which is absolutely useless.

R: So that's a cover page without any sort of summary?

T6-A: Yeah. So I don't know. If you print 'everything,' let's see what that does instead [prints 'everything'].

R: I guess another option is to copy and paste the text into Word?

T6-A: I *could* do that, but as I know these numbers at the bottom refers to the document number that it was written on originally before it was transferred to [the firm's KM database], that's why I just went in and copied it over. And I would just open it read-only probably and print it.

T6-A: I'll see what the [KM database] version printed [walks to printer]. See it's still not printed the text. It's just given me some kind of profile classification thing [points to hardcopy of meta-data relating to the document] and URLS. It hasn't actually printed the text.

T6-A: Then I'd start going back through the other [KM database] documents. [Clicks on each document in the results list in turn]. This one, for example, I know because it's referring to 'Earnings' [points to the word 'earnings' in one of the document headings] which is not a situation that I'm looking at, so I wouldn't bother with that. [Continues to look through document summaries]. So I would now go through the rest of my 11 search results and see if I thought any of them were useful and if any of them were, then I'd probably print them out and then once I'd done all of Simon's, the Revenue and all of our internal correspondence, then I'd probably sit down at that point and that would be my starting bank of information. I'd have a read through that and then either specific points came up whilst I was reading that or absolutely nothing was useful at all and I thought that I would have to widen my search to different sources as well, then I would maybe go to our paper-based sources such as books.

Appendix 2: Guidance for conducting IB evaluations

In this document we describe the IB functionality methods and present guidance for those who seek to carry out either usability, functionality or both types of evaluations. The broad process of conducting a functionality or IB usability evaluation involves:

1. Defining the **purpose and boundaries** of the evaluation.
2. Deciding on the **practicalities** of the evaluation.
3. Considering the **ethical issues** surrounding the evaluation.
4. **Conducting the evaluation itself** and **recording the output**.
5. **Communicating the findings** from the evaluation. This stage is not discussed in detail in this guidance document as the ways in which users of the IB methods are likely to choose to communicate their findings are likely to vary widely (ranging from using them as the basis of formal reports to using them as a basis of informal presentations or discussions).

There are a number of resources that you can use to support you in conducting an IB evaluation. The first is the behaviour definition and examples document provided in appendix 3. This documentation lists the definition of each of the information behaviours at the core of the methods, along with illustrative ways that electronic resources might support these behaviours at each applicable level. The document also contains screenshots illustrating some of the examples using Justis, an electronic legal resource.

The second supporting resource that you can use are the forms to be used to record the output of the evaluation and the user think-aloud information sheet and list of behaviour-focused tasks that evaluators can choose to present to users in order to obtain think-aloud data (which they will, in turn, identify usability issues from). The forms for recording the output from an IB functionality evaluation are presented in appendices 6 and 7. The form for recording the output from an IB usability evaluation is presented in appendix 8. The think-aloud information sheet and task list to be used to structure user think-aloud sessions as part of an IB usability evaluation are presented in appendix 5.

The third (and in our opinion most important) resource for conducting an IB evaluation is this guidance document, which explains in detail how to conduct functionality and usability evaluations using the behaviour definition and examples document in appendix 3 and the forms and think-aloud task list presented in appendices 5-8. It is important to note, however, that although some of the recommendations in this section for conducting an IB evaluation are based on the Usability Evaluation Methods literature and on the results of our pilot studies, others are based on beliefs of 'good practice' as held by the development team (and not on observed evidence or direct literature involvement). As testing the appropriateness of this guidance falls beyond the scope of this thesis, these 'good practice' recommendations are not intended to be accepted as gospel. Furthermore we do not claim that our recommendations are the 'only' or 'correct' way to carry out aspects of an IB evaluation. We suggest the recommendations simply serve to illustrate how *we believe* a successful IB evaluation might be carried out. In light of this discussion, it should be assumed that unless evidence from the literature or pilot studies is cited, any recommendations highlighted in this section are 'good practice' guidelines and have not been tested empirically.

We now begin our explanation of how to conduct an IB evaluation, starting with guidance how you might decide which electronic resource to evaluate and how to decide whether to conduct a functionality evaluation, a usability evaluation or both. We then provide guidance for the remainder of the process: deciding on the practicalities of the evaluation, defining the boundaries, considering the ethical issues surrounding the evaluation and conducting the evaluation itself and recording the output. We provide this guidance first for conducting a functionality evaluation, then for conducting a usability evaluation.

a)   Defining the purpose and boundaries of the IB evaluation

The first task when deciding to conduct an Information Behaviour evaluation is for you to decide on its purpose and boundaries. More specifically, this involves:

1. Deciding which electronic resource to evaluate. It is possible to evaluate an '*own resource*' or a '*competitor resource.*' An 'own resource' is one in which you have a direct stake (e.g. it is developed or sold by your firm). A 'competitor resource' is developed by another firm and will usually share similar aims or design features to one or more of the resources your firm develops or sells.
2. Deciding what type of evaluation to carry out (it is possible to conduct an IB functionality evaluation, an IB usability evaluation, or both).
3. Deciding which parts of the resource to evaluate.
4. Deciding which sets of behaviours to evaluate the resource in relation to.

For usability evaluations only, it is also necessary to decide whether to conduct a *core*, *recommended* or *custom* IB usability evaluation.

Deciding which electronic resource to evaluate and what type of evaluation to carry out

We believe value can be gained from evaluating competitors' resources as well as own resources. In addition, value can be gained from carrying out both a functionality and usability evaluation of the same resource as different insights can be gained from each type of evaluation.   Therefore deciding whether to conduct a functionality evaluation, a usability evaluation, or both and on which electronic resource to evaluate depends on the *purpose* of the IB evaluation – i.e. why you are turning to the method in the first place:

- Is it to identify **ways that a competitor's product currently supports certain functionality** in order to somehow compare this functionality to that provided by one of your own products? If so a functionality evaluation of the competitor's resource should be considered.
- Is it to identify **ways that you might increase or even reduce the functionality provided by your own product**? If so a functionality evaluation of your own resource should be considered.
- Is it to identify **issues that users have difficulty with when using one of your products** so that these issues can be addressed and the product improved? If so a usability evaluation of your own resource should be considered.
- Is it to identify **issues that users have difficulty with when using a competitor's product** in order to determine what usability advantage you have over the competitor or they have over you? If so, a usability evaluation of the competitor's resource should be considered.
- Is it **an assortment** of these? If so, more than one type of IB evaluation should be considered.

Assessing the *functionality* provided by an electronic resource can lead to identifying gaps in functionality that can be filled by adding additional functionality support to an electronic resource. Assessing resource functionality (and in particular considering whether it is still necessary to support all of the behaviours/levels that the resource currently supports) can also lead to identifying ways in which functionality might be *reduced*. Although reducing functionality is not normally considered by those with a stakehold in a resource, it is particularly pertinent to consider for electronic legal resources as many of these resources have a variety of, sometimes complicated, system features that might make the resource more difficult to use than if the features were not there. Assessing the *usability* of an electronic resource can identify usability issues that can be addressed in future design interventions. In addition, we believe that analysing think-aloud data obtained from observing actual or intended users of an electronic resource can be more useful than a developer or other stakeholder evaluating the usability of the resource themselves (i.e. by conducting other forms of usability evaluation that do not require user involvement).

Deciding which parts of the resource to evaluate

Usually you will want to conduct a functionality or usability IB evaluation on an entire electronic resource and it will rarely be necessary to exclude various screens or system features from the evaluation. However evaluators *may* be interested in evaluating only a subset of the resource if their remit is restricted to designing and/or evaluating particular parts of the resource. This is particularly likely to be the case for large legal resources, which can provide dedicated screens and system features for accessing a range of legal documents, from a range of topical areas and jurisdictions. For these large resources, it may be impractical to focus evaluation efforts on the entire electronic resource and, instead, the functionality or usability evaluation should be focused on certain parts of the resource (for example those parts that are commonly used, or enabled as default for the majority of users or particular groups of users). Therefore evaluators conducting an IB evaluation should aim to evaluate the entire electronic resource unless there are specific reasons to exclude certain parts.

Deciding which sets of behaviours to evaluate the resource in relation to

As with deciding which parts of an electronic resource to evaluate, usually you will want to evaluate the resource using all three sets of information behaviours - the core and legal-specific information-seeking behaviours and the wider information use behaviours. However, there may be some exceptions. Only *legal* resources should only be evaluated in relation to the legal-specific behaviours. In addition, it may not be within the scope of the electronic resource under evaluation to support wider information use behaviours such as analysing, synthesising, recording, collating, editing and distributing information. If this is the case, these behaviours should be excluded from the evaluation. As a general rule, we would encourage evaluating an electronic resource in relation to *all three sets* of behaviours unless any of the above exceptions apply.

Deciding whether to conduct a core, recommended or custom evaluation (for IB usability evaluations only)

If you have decided to conduct an IB usability evaluation, it is also necessary to decide which type of evaluation to carry out: a *core*, *recommended* or *custom* evaluation. These three types of usability evaluations allow users of the usability method to choose an evaluation that best suits the focus of their evaluation and fits in with any time or cost constraints they might have.

The 'core' IB usability evaluation involves asking intended or actual users of the resource to think aloud whilst performing three core information-seeking tasks:

1. Gain access to the electronic resource.
2. Find out which parts or sources within the resource you have access to.
3. Think of some information that you currently need or have recently needed to find for your work and demonstrate, using the electronic resource, how you might go about finding it.

A core IB usability evaluation is recommended as a 'quick and dirty' way of acquiring user think-aloud data that can then be analysed to identify usability issues. The core evaluation is highly naturalistic, but is only somewhat behaviour-focused as the tasks users are asked to perform are very broad. As illustrated by our first and second user pilot studies, the broad nature of the tasks means that the core method is only likely to allow users to demonstrate a potentially limited range of user behaviours which, in turn, may limit the range of related usability issues that can be identified from the think-aloud data. The core evaluation does, however, require very little user involvement (around 20-30 minutes per user).

The 'recommended' usability evaluation involves asking users to think aloud whilst performing the core tasks as above, plus a broader range of behaviour-focused tasks (presented in appendix 5). The recommended evaluation can be used to acquire rich and behaviour-focused think-aloud data (as illustrated by our second user pilot study). This type of evaluation is naturalistic (although not as naturalistic as the 'core' evaluation, as the

recommended evaluation includes tasks that are more prescriptive) and requires around 1 hour of user involvement (more involvement than required under a core evaluation, but less than under a custom evaluation).

Finally the 'custom' usability evaluation involves asking users to think aloud whilst performing a custom set of tasks, chosen from the bank of tasks in appendix 5. This allows users of the method the most flexibility in deciding what tasks to set. Users of the method can choose tasks related to particular behaviours and/or levels of interest, which is likely to result in rich and highly behaviour-focused think-aloud data. Although the custom evaluation is only somewhat naturalistic due to the highly prescriptive nature of the custom tasks, we hypothesise it has the potential to result in users demonstrating tasks that they might not have demonstrated in a recommended or core evaluation and therefore highlighting usability issues that might not otherwise have been highlighted. Whilst the amount of user involvement in a custom evaluation varies tremendously depending on which and how many tasks are set for users to attempt, we suggest trying to make sessions last no longer than an hour and a half. The trade-offs between the core, recommended and custom versions of the IB usability evaluation are summarised in table a.

Type of IB usability evaluation	When is it appropriate to use this type of evaluation?	How naturalistic is this type of evaluation?	How behaviour-focused is this type of evaluation?	What range of behaviours is this type of evaluation likely to encourage users to demonstrate?	How much user involvement is required?
Core	Recommended as a 'quick and dirty' way of acquiring user think-aloud data that highlights usability issues.	Highly naturalistic.	Somewhat behaviour-focused.	Aims to encourage demonstration of a (potentially limited) range of user behaviours and to highlight related usability issues.	Around 20-30 minutes per user.
Recommended	Recommended as a way of acquiring rich and behaviour-focused user-think-aloud data that highlights usability issues.	Naturalistic.	Behaviour-focused.	Aims to encourage demonstration of a broad range of user behaviours at the most commonly observed levels and to highlight related usability issues.	Around 1 hour per user.
Custom	Recommended as a way of acquiring rich user-think-aloud data that highlights usability issues focused on particular behaviours and/or particular levels at which these behaviours can operate.	Somewhat naturalistic.	Highly behaviour-focused.	Aims to encourage demonstration of particular behaviours and/or particular levels of interest to the usability evaluator (i.e. not necessarily the most commonly observed levels) and to highlight related usability issues.	Around 1½ hours per user.

Table a: A comparison of the 'core,' 'recommended' and 'custom' types of IB usability evaluation.

b)   Conducting an IB functionality evaluation

After deciding which electronic resource to evaluate and which type of evaluation to conduct, the broad process for conducting a functionality and IB usability evaluation is similar and involves deciding on the practicalities of the evaluation, defining the boundaries, considering the ethical issues surrounding the evaluation, conducting the evaluation itself and recording the output and reporting the findings. We now turn to discuss this process for conducting a *functionality evaluation*, beginning with deciding on the practicalities. The process is discussed in relation to conducting a *usability evaluation* in section c.

Deciding on the practicalities of the functionality evaluation

Deciding on the practicalities of an IB functionality evaluation involves considering a number of questions.

When in the design and evaluation cycle should the functionality of the resource be evaluated? As with a Cognitive Walkthrough, the IB functionality method can be applied either with a detailed specification of an interface, or an actual interface. Therefore there is no set place in the design and evaluation cycle to evaluate the functionality of a resource. However we hypothesise that, like a Cognitive Walkthrough, the functionality method might be most useful when applied *"after an intermediate milestone such as prototype creation"* (Wharton et al. 1994, p. 109) and the results used to inform the next revision of the electronic resource. Indeed, we recommend that you use both the functionality and usability methods as part of an iterative design and evaluation process, where results of evaluations are used to inform the design of subsequent versions of the resource.

Who should participate in the functionality evaluation? We suggest IB functionality evaluations are conducted in small groups as we believe that the value of the method lies primarily in *encouraging and supporting discussion* about functionality issues amongst people with a direct stake in the electronic resource under evaluation. We recommend this group comprises usability professionals, who are less likely to become 'attached' to particular functionality than members of the development team. However, an IB functionality evaluation *can* include members of the development team if it is felt that they will remain 'politically unbiased.'

The group should also include someone with the assigned role of note-taker who has the task of filling out the functionality forms (provided in appendices 6 and 7). The note-taker is encouraged to participate in the functionality discussions, but should be aware that their primary role is to fill out the functionality forms clearly. The note-taker is reminded that the forms provide a section allowing them to make general notes related to each behaviour and applicable level at which it can be performed. Therefore the job of the note-taker involves filling out the functionality forms *and* making general notes in the space provided on the form in appendix 8 – as long as those notes relate to discussions surrounding functionality support for the current behaviour/level. Finally, it is recommended that the group should include an independent facilitator, such as a Human-Computer Interaction expert. The role of the facilitator involves ensuring that the functionality evaluation runs at an appropriate pace, ensuring that all group members are given the opportunity to voice their opinions and ensuring that no individual group members dominate discussions at the expense of others. Regarding recommended group size, we suggest that an IB functionality evaluation is carried out in small groups of no more than around 10 people. The reasons for this are twofold: firstly, diminishing returns are likely to be gained from increasing the number of group members conducting the evaluation and secondly, it is more difficult to facilitate larger groups. There is no lower recommended limit for group size, although we hypothesise that functionality evaluations with three or more participants are likely to be more useful than with just two participants.

When evaluating the functionality of an *own resource*, it is also important to include people in the group who are already familiar with the functionality of the resource. The group need not include individuals with the responsibility to make ultimate decisions on what functionality to support in a particular resource as the aim of the functionality evaluation is not to make binding decisions on what functionality to support. Indeed this would not be practical as such decisions are likely to be made by several individuals, perhaps based in different physical locations. In addition, final decisions regarding what functionality to support are also likely to be informed by other data, such as usage statistics or user training data.

When evaluating the functionality of a *competitor resource*, our recommendations for selecting who to participate in the evaluation are less prescriptive. As the evaluation involves *evaluating* the resource and as it is beyond the scope of the evaluation to discuss anything other than the ways the resource *currently supports* behaviours at particular levels, no prior knowledge of the functionality provided by the competitor product is needed by participants. However, when evaluating a competitor resource, one member of the group should be assigned the role of exploring the resource with direction from the rest of the group (referred to in this chapter as

a 'demonstrator'). The demonstrator need not think-aloud whilst using the resource, but instead should participate in the discussions surrounding the evaluation.

Comments from the evaluation of the IB functionality method suggested that it might not always be possible to conduct the evaluation in a small group. If this is found to be the case (for example, because there is not enough time to involve more than one member of the team), it is also possible to conduct the evaluation *in pairs* or, if necessary, *individually*. We maintain, however, that the likely value of the method is in the discussion it encourages, therefore even if the evaluation is carried out by a single individual, it is important to feed back findings to other team members with the purpose of opening a discussion on how functionality might be increased and/or reduced for the resource under evaluation.

How much time should be devoted to evaluating the functionality of the resource with regard to each behaviour or set of behaviours and how much time should be devoted to the evaluation overall? The answer to this question will vary depending on the size and complexity of the electronic resource to be evaluated and the range of behaviours that will be used to evaluate it. Based on the results of our three user pilots and one developer pilot, we do not suggest setting firm time limits for the evaluation. However we suggest that it is reasonable to put aside around two or three hours to conduct a functionality evaluation of a large competitor or own resource with many features or around an hour for smaller resources. Based on experience in our user pilot studies, we suggest that more important than setting an overall time limit for the evaluation is to keep a rough eye on the time spent evaluating a competitor resource with regard to or discussing functionality issues relating to each behaviour or group of behaviours. We suggest that is it important not to spend too long examining a particular behaviour/level, nor to rush through the behaviours and levels as this might discourage participants from making useful functionality-related comments.

How should the functionality evaluation be recorded? IB functionality evaluations should be recorded by the designated note-taker on the functionality forms provided. Although it is possible to audio-record the evaluation in order to capture the functionality discussions in their entirety (rather than in note form), we do not suggest that this is necessary in most cases. This is because the time and cost associated with transcribing the audio is likely to outweigh the benefits of having a word-for-word account of discussions. This recommendation is supported by our developer pilot, where an audio recording of the brief functionality discussion did not prove to provide much added value above the notes made during the discussion.

Should the functionality method be conducted on its own, integrated into current methods of assessing functionality or conducted alongside current methods of assessing functionality? It is certainly possible to conduct an IB functionality evaluation alongside other ways of evaluating the functionality of resources. It is also possible to integrate the method into current practices for assessing functionality. In the latter case, however, particular care must be taken to avoid the method being adopted as a simple functionality checklist, where the concern is to support particular behaviours/levels rather than use the method as a means for *encouraging and supporting discussion surrounding* supporting particular behaviours/levels. Therefore it is important that the spirit and ethos of the method is preserved if any changes are made to it so that it can be integrated into everyday working practices. This was an issue partly raised by our developer pilot session, as the developer suggested that the 'biggest drawback' of using the concepts of information behaviours and levels to frame functionality discussions was the potential to misunderstand and therefore misapply them.

Considering ethical issues surrounding the functionality evaluation

Before conducting the functionality evaluation itself, it is necessary to consider a number of potential ethical issues surrounding the evaluation. Perhaps the most important ethical issue to consider when conducting a functionality evaluation surrounds ensuring that all participants in the functionality evaluation feel that they are able to contribute to the evaluation and to voice their true opinions surrounding the functionality of a product, even if they or another member of the team has been involved in or had advocated the development of the product and/or the functionality in question. Although this is by no means easy to achieve, some positive steps

can be taken to ensure that the functionality evaluation is viewed in a positive light by all members of the group and that individuals do not feel as though the evaluation is aimed at undermining the value of their design work. For example, the independent facilitator should *emphasise the exploratory and non-binding nature of the functionality evaluation.* If members of the group are aware that the evaluation aims to encourage and support discussion surrounding support for certain functionality, not to help them identify superfluous functionality that they must remove, we believe they will be more receptive to the method.

Conducting the functionality evaluation itself and recording the output

In an IB functionality evaluation of a *competitor's* electronic resource, evaluators explore the resource to find out whether and how the competitor resource currently supports a range of different information behaviours. Similarly, in a functionality evaluation of an *own resource,* evaluators ask themselves whether and how their own resource currently supports certain behaviours. If the answer to this question is 'yes,' evaluators then ask whether there are any additional ways that the resource might support the behaviour. If the answer to this question is 'no,' evaluators ask how the resource might support the behaviour. When conducting an IB functionality evaluation of an own resource, evaluators also ask themselves whether there are any behaviours/levels (or ways of supporting them) that it may be no longer necessary to support. If they identify any behaviours, levels, or ways of supporting them that they are considering ceasing support for, an IB functionality evaluation also involves considering the potential arguments for and against ceasing support.

Although it is possible for IB functionality evaluations to be conducted in pairs (or even by individuals), we evaluations are conducted in a small group (see earlier discussion). If conducted in a group, we recommend that the group includes an independent facilitator (such as an HCI expert or someone who has no professional or personal ties with the individuals conducting the functionality evaluation). This is primarily to ensure that all group members are given the opportunity of fair and equal participation. The facilitator's role is to inform the group which behaviour and level they are currently evaluating, to try to ensure the evaluation runs at an appropriate pace, to try to ensure that all group members are given the opportunity to voice their opinions and to try to ensure that that no individual group members dominate discussions at the expense of others. In practice, this should involve minimal intervention. For large groups, the facilitator might decide to ask group members to raise their hands when they have something to contribute in the discussion (thereby helping to ensure a balanced discussion). The facilitator should also aim to ensure that functionality discussions do not turn into discussions about *how exactly* a resource might support a particular behaviour/level at the interface level. This involves steering the discussion back on course if necessary.

If conducted in a small group, this group should also include someone with the role of note-taker, who should fill out the functionality forms in appendices 6 and 7 and make general notes related to the discussion surrounding functionality support for the current behaviour/level in the space provided on the form in appendix 7. When evaluating a *competitor resource*, the group should also include someone with the role of demonstrator, who should explore the competitor resource with direction from the rest of the group (and the facilitator). More specifically, the demonstrator will explore the functionality provided by the competitor resource related to each behaviour and at each applicable level. This should be carried out on a screen that all members of the group can see clearly (such as a projector screen).

Before beginning the functionality evaluation, the note-taker (or the individual or pair conducting the evaluation) should keep a record of the details pertaining to the evaluation on the summary functionality form in appendix 6. This involves recording the name/version of the electronic resource under evaluation, the date of the evaluation and the names of the group members that were present during the evaluation. This also involves listing the parts of the resource that the functionality evaluation has been restricted to (only if applicable) and indicating which set or sets of behaviours the evaluation has been restricted to (again, only if applicable). Users of the method can use their own shared terminology to refer to 'parts' of the resource. This may be based on the names of screens/pages within the resource, names of different pieces of functionality etc. An IB functionality evaluation might be restricted to only the core information-seeking behaviours (when evaluating non-legal resources that do

not aim to support information use behaviours), to the core and law-specific behaviours (when evaluating legal resources that do not aim to support information use behaviours), or the core and information use behaviours (when evaluating non-legal resources that aim to support information use behaviours). When examining legal resources that aim to support information use behaviours, there is no need to restrict the evaluation to a particular set or sets of behaviours and the evaluation should involve assessing the functionality of the electronic legal resource in relation to all fourteen of the information behaviours listed in the table on the summary functionality form in appendix 6. This table is re-printed below as table b.

	Resource level	Source level	Document level	Content level	Search query/result level
Accessing	Currently supported? (Y/N)	Currently supported? (Y/N)	Currently supported? (Y/N)		
Surveying		Currently supported? (Y/N)	Currently supported? (Y/N)		
Monitoring		Currently supported? (Y/N)	Currently supported? (Y/N)		
Searching		Currently supported? (Y/N)	Currently supported? (Y/N)	Currently supported? (Y/N)	
Browsing and extracting		Currently supported? (Y/N)	Currently supported? (Y/N)	Currently supported? (Y/N)	
Chaining		Currently supported? (Y/N)	Currently supported? (Y/N)		
Selecting, distinguishing and filtering		Currently supported? (Y/N)	Currently supported? (Y/N)		
Updating			Currently supported? (Y/N)		
History tracking			Currently supported? (Y/N)		
Analysing				Currently supported? (Y/N)	
Recording		Currently supported? (Y/N)	Currently supported? (Y/N)	Currently supported? (Y/N)	Currently supported? (Y/N)
Collating			Currently supported? (Y/N)	Currently supported? (Y/N)	
Editing			Currently supported? (Y/N)		
Distributing			Currently supported? (Y/N)	Currently supported? (Y/N)	Currently supported? (Y/N)

Key: Core information-seeking behaviours ☐   Law-specific behaviours ▨   Information use behaviours ▨

Table b: Table of information behaviours and applicable levels to be considered in an IB functionality evaluation.

In order to evaluate a **competitor resource**, the demonstrator (or the individual or pair of evaluators) should explore the resource in relation to each of the behaviours and levels listed in table b (apart from any behaviours that were excluded from the evaluation). Definitions of each behaviour and examples of how electronic resources might support the behaviours at each applicable level are provided in the supporting documentation in appendix 3. **It is not advisable** to conduct a functionality evaluation without using the supporting documentation during the evaluation as a reference guide, as the potential for wrongly interpreting the intended meaning of behaviours and levels is high (as illustrated by our developer pilot study). The examples in the supporting documentation greatly reduce the scope for error during an evaluation. The remaining process of conducting an IB functionality evaluation of a competitor resource is as follows (and should be carried out in sequence for each behaviour and applicable level, one behaviour/level at the time):

- **Explore the competitor resource to find out whether the behaviours, at each applicable level, are currently supported by the resource.** Using the table of behaviours and applicable levels on the summary functionality form to help structure the evaluation, the group should first attempt to ascertain whether each behaviour is supported by the competitor resource, *at each applicable level* (i.e. those levels in the table on the functionality forms). This involves directing the demonstrator to perform tasks related to different ways that the behaviour might be supported at this level. The supporting documentation in appendix 3 can be used to help the group decide on ways to explore the resource in

relation to each behaviour and level (i.e. which tasks to ask the demonstrator to perform). Note, however, that the emphasis of a competitor IB functionality evaluation is on *exploring* the competitor resource and it is not necessary to use the supporting documentation to provide a rigid structure to the exploration – the document should be used to ensure the group understands what supporting each behaviour and level entails and to assist the group if they are unsure how to explore the resource in relation to a particular behaviour/level. The note-taker should indicate on the summary functionality form in appendix 6 whether each behaviour/level is currently supported. Disagreements about whether a behaviour/level is currently supported should, wherever possible, be resolved by referring to the supporting documentation in appendix 3.

- **Determine in which way(s) the resource currently supports the behaviour at each applicable level.** For each behaviour and applicable level that is currently supported by the competitor resource, determine in which way(s) the resource currently supports the behaviour at this level. List these under (a) on the detailed functionality form, filling out a new one-paged detailed form for each level of each behaviour that has been analysed. The supporting documentation in appendix 3 should be used so that the group can familiarise themselves with ways that an electronic resource might support the behaviour and level under consideration.

In order to evaluate an **own resource**, it is not necessary to have a demonstrator exploring the resource as it is a pre-requisite that group members participating in the evaluation will have a sound knowledge of the functionality provided by their own resource. However, it is useful to have a shared screen (such as a projector) set up to display the resource in order to support the functionality discussion (e.g. to remind group members whether or in which ways a particular behaviour/level is supported or to settle any disagreements about whether a behaviour/level is currently supported). It is expected, however, that at least some of the evaluation will be carried out without direct reference to a running version of the resource. As with a competitor functionality evaluation, an own resource evaluation is also supported by the documentation in appendix 3 and, once again, we do not believe it would be fruitful to conduct the evaluation without using the supporting documentation during the evaluation as a reference guide. The remaining process of conducting an IB functionality evaluation of an own resource is similar to the process of conducting a functionality evaluation of a competitor's resource, but also involves considering additional ways that the own resource might support behaviours at a particular level (if the resource currently supports the behaviour/level) or ways in which the behaviour/level might be supported (if the resource does not currently support the behaviour/level). A functionality evaluation of an own resource also involves considering whether it is still necessary to support all of the behaviours/levels that the resource currently supports. This process is as follows (and, as with evaluating a competitor resource, should be carried out in sequence for each behaviour and applicable level, one behaviour/level at the time):

- **Discuss whether the behaviours, at each applicable level, are currently supported by the own resource.** Using the table of behaviours and applicable levels on the summary functionality form to help structure the evaluation, the group should first attempt to ascertain whether each behaviour is supported by the own resource, *at each applicable level*. Unlike with a functionality evaluation of a *competitor* resource, this primarily involves *discussion* amongst the group as opposed to exploring the resource. However, this exercise can still be supported by a running or prototype version of the resource under evaluation. Again, the supporting documentation in appendix 3 should be used to ensure the group understands what supporting each behaviour and level entails and, in the same way as when conducting an evaluation of a competitor resource, the note-taker should indicate on the summary functionality form in appendix 6 whether each behaviour/level is currently supported. As with conducting an IB functionality evaluation of a competitor resource, disagreements about whether a behaviour/level is currently supported should, wherever possible, be resolved by referring to the supporting documentation in appendix 3.
- **For levels of a behaviour that the resource *currently supports*, determine in which way(s) the resource currently supports and in which additional ways it might support the behaviour at this level.** For each behaviour and applicable level that is currently supported by the competitor resource, determine in which way(s) the resource currently supports the behaviour at this level. List these under (a) on the detailed functionality form. Also for each behaviour and applicable level that is currently

supported by the competitor resource, determine in which additional way(s) the resource might support the behaviour at this level. This is a somewhat creative task, as it may be possible to think of innovative ways to support a certain behaviour at a particular level. List these ways under (b) on the detailed functionality form in appendix 7. The supporting documentation in appendix 3 should also be used so that the group can familiarise themselves with ways that an electronic resource might support the behaviour and level under consideration. As when conducting a functionality evaluation of a competitor resource, it is necessary to fill out a new one-paged detailed form for each level of each behaviour that has been analysed.

- **For levels of a behaviour that the resource *does not currently support*, determine in which way(s) the resource might support the behaviour at this level.** For each behaviour and applicable level that is *not* currently supported by the competitor resource, determine in which way(s) the resource might support the behaviour at this level. List these under (c) on the detailed functionality form, once again using the supporting documentation in appendix 3 to provide guidance.

Finally, after repeating this process for each of the behaviours and levels listed on the summary IB functionality form in appendix 6, an IB functionality evaluation of an own resource involves asking the following general questions (also on the summary form):

- **Are there any *behaviours/levels* that it may no longer be necessary to support?** For any behaviours/levels which you are considering ceasing support for, **what are the potential arguments for and against support?**
- **Are there any *ways that you currently support any of the behaviours/levels* that may no longer be necessary?** For ways of supporting a particular behaviour/level which you are considering ceasing support for, **what are the potential arguments for and against support?**

These questions aim to address the often assumption that increasing functionality will always lead to an electronic resource that is easier to use. This might indeed be the case, however increasing functionality might also serve to make the resource more complicated and hence confuse users, making the resource more difficult to use overall. Therefore we suggest it may not always be useful to support or continue to support functionality related to a particular behaviour/level. It is important to emphasise that decisions made when answering these (or any of the other IB functionality) questions are non-binding. By this, we mean that we do not expect that conducting an IB functionality evaluation will automatically lead to increased or reduced functionality support in the electronic resource under evaluation. These questions are intended to *facilitate discussion* surrounding *whether the functionality evaluators believe that particular behaviours/levels should be supported* by the resource if they are not already (and, if so, *in which ways they might be supported*) and *whether it is necessary to continue supporting* particular behaviours/levels. The questions are not intended to 'force' the evaluator(s) to introduce or cease support for particular behaviours and/or levels.

Potential arguments *for* supporting a particular behaviour/level include the fact that usage statistics or other marketing material suggest that the functionality related to the behaviour/level is frequently used or highly valued by users, that the functionality related to the behaviour/level is not supported (or not supported well) by competitor products and therefore helps to create a competitive advantage, or that the functionality reflects an important aspect of user behaviour (e.g. lawyers' work) even if it is not frequently used. Potential arguments *against* supporting a particular behaviour/level include the fact that usage statistics or other marketing material suggest that the functionality related to the behaviour/level is infrequently used or poorly valued by users or that the functionality does not reflect an important aspect of user behaviour. Another important argument against supporting a particular behaviour/level, highlighted from our evaluation of the method, is if there are any technical or resource constraints which might prevent support.

c) Conducting an IB usability evaluation

Deciding on the practicalities of the usability evaluation

Deciding on the practicalities of an IB usability evaluation involves considering similar questions to those in a functionality evaluation. However, the guidance we provide in this section differs considerably to that provided for conducting a functionality evaluation due to the fact that the IB functionality and usability methods have different purposes and are conducted in very different ways. The questions to be considered when deciding on the practicalities of a usability evaluation are:

When in the design and evaluation cycle should the usability of the resource be evaluated? Although it is possible to ask users to think aloud whilst using or low-fidelity (e.g. paper or wire-frame) prototype of an electronic resource, we hypothesise that most value can be gained from evaluating either a high-fidelity prototype or a fully-implemented electronic resource. As we have mentioned earlier, we recommend that the evaluation be conducted as part of an iterative design and evaluation process.

Who should participate in the usability evaluation? IB usability evaluations involve intended or actual users of an electronic resource thinking-aloud whilst performing a number of behaviour-focused tasks. Although it would be theoretically possible to substitute these with other stakeholders in the resource, we do not believe that it is possible for those with intimate knowledge of how to use a particular resource to behave like a 'regular user' in order to generate think-aloud data. Therefore we do not suggest conducting an IB usability evaluation unless it is possible to recruit at least one intended or actual user of the resource to think aloud. By 'intended' user, we mean a user that a particular resource has been designed to support. For example, intended users of a new version of an electronic legal resource might include users of the previous version of the resource and users of competitor resources. There is also potential value to be gained from asking actual users, with a variety of previous experience using the resource, to perform the tasks and think aloud. Alternatively (or even in addition), it is possible to recruit different user groups (e.g. for electronic legal resources, taught students, research students, academic staff, practicing lawyers with various legal specialisms etc.). There is potential for these different user groups to provide varied, but equally useful and rich think-aloud data, as illustrated by our earlier empirical study of the information behaviour displayed by different groups of lawyers.

The number of users recruited depends on time and cost constraints. It is possible to use the IB usability method to generate think-aloud data that can highlight usability issues from a single user (as illustrated by our user pilot studies – the usability issues identified from each user's think-aloud data are presented in appendix 4). However, generalising from this data will be problematic and there is no guarantee that other users will encounter similar usability problems (indeed a variety of usability issues were identified across our three user pilot studies). However, the process of analysing IB think-aloud data is relatively time-consuming and, therefore, it would be impractical to recruit and analyse think-aloud data from scores of users. We suggest aiming to recruit around 5 users to begin with and, resources permitting, recruiting a small number of further users if feasible. This is based on a de facto guideline that diminishing returns are provided as the number of users recruited to participate in a usability evaluation increase (see Nielsen and Landauer, 1993). Nielsen and Landauer recommend testing with 5 users first, then feeding the findings of the evaluation into design intervention and conducting further user tests afterwards. This is also a possibility for conducting an IB usability evaluation, as we also advocate conducting evaluations as part of an iterative design and evaluation process.

The user think-aloud data generated from an IB usability evaluation is intended to be analysed by individuals or pairs. Preferably, these individuals or pairs should have a grounding in the field of Human-Computer Interaction or some experience in conducting usability evaluations in general. However they do not need to be HCI experts. The think-aloud data can be analysed (and usability issues extracted) by usability experts or other stakeholders in the resource, for example, provided they have some knowledge of usability evaluation in general. One of these individuals can also act as facilitator for the user think-aloud sessions (with a main role of ensuring that users do not over-run on tasks and that the think-aloud data has been recorded properly). Once again, we discourage

asking members of the development team to evaluate resources that they have been involved with designing, unless you can be reasonably sure that they will be impartial.

How much time should be set for users to attempt each behaviour-focused task and how much time should be devoted to the usability evaluation overall? This depends on the type of usability evaluation that will be conducted. A recommendation (in number of minutes) is provided next to each task and sub-task that might be used as part of an IB usability evaluation (see appendix 5). Also, as a general rule (as illustrated by our user pilot studies), a core usability evaluation should generate around 20-30 minutes of user think-aloud data, a recommended evaluation around an hour of data and a custom evaluation anything up to one and a half hours of data.

How should the usability evaluation be recorded? User think-aloud data should be both audio and video recorded. Software tools such as Camtasia can be used to support recording user interaction, saving all of their screen activity (and the related think-aloud audio) into a video clip. Whilst it is possible to identify usability issues by only audio recording the think-aloud data, we do not recommend this approach as it is often necessary to view the user's screen activity in order to make useful assumptions about what they are trying to achieve (as illustrated by our user pilot studies, where extraction of usability issues proved to be extremely difficult without watching the user's screen interaction as well as listening to the accompanying audio think-aloud data). It is also very useful to keep a record of users' interaction with the resource as it is difficult to remember exactly what interface-level actions the user performed from an audio recording alone. The output of the IB usability evaluation (for example the list of usability issues identified) should be recorded on the form provided in appendix 8.

Should the usability method be conducted on its own or conducted alongside other methods of assessing usability? We believe an IB usability evaluation can complement many other forms of evaluation and the output can, as we have mentioned, feed in to future design discussions. It is, for example, possible to conduct an IB evaluation alongside other usability evaluation methods such as Cognitive Walkthroughs, Heuristic Evaluations or Expert Evaluations. This can provide different perspectives on things that users might find difficult when using the electronic resource. Similarly, the usability method can also be used alongside other user evaluation methods such as traditional (non-behaviour-guided) think-alouds. It also goes to follow that by using behaviour-focused tasks as a structure, we believe that an IB usability evaluation will provide different insights to these other methods. This, however, should be considered a hypothesis that still requires testing (and is discussed further in chapter 8).

Considering ethical issues surrounding the usability evaluation

As highlighted by Blandford et al. (2008), it is important to ensure that participants are informed of the purpose of the study they are participating in (in this case the think-aloud session associated with the usability evaluation) and that they are free to view or ask for their data (in this case the audio and video recording of their think-aloud session) to be deleted at any time, for any reason and without penalty. Blandford et al. also highlight the importance of keeping user think-aloud data as anonymous as possible and respecting participants' confidentiality and privacy. We also believe it is important to inform informing participants of how their think-aloud data will be used and, if applicable, disseminated. An example informed consent form that users of the IB usability method might adapt for their purposes is presented in appendix 10.

Conducting the usability evaluation itself - obtaining user think-aloud data

The first step in order to obtain user think-aloud data is to **recruit intended or actual users** of the electronic resource under evaluation. We suggest aiming to recruit around 5 users to begin with and, resources permitting, recruiting a small number of further users if feasible. Think-aloud users might be recruited from a number of channels. Probably the most straightforward of these channels, provided users of the resource have given their

consent, is to contact subscribers of an own resource that are based in a suitable geographical area to see whether they would be willing to take part in think-aloud session. It may also be possible to place an advertisement on the Internet (e.g. on the login page of the resource website). Alternatively (and particularly suitable for conducting a usability evaluation of a *competitor* resource), it may be possible to recruit users through a market research agency. Further issues surrounding deciding on which types and how many users to recruit are discussed in the 'deciding on the practicalities of the evaluation' section of this chapter. After users have been recruited, a suitable quiet setting should be selected for them to conduct the think-aloud observation. This setting should have a computer capable of making an audio and screen recording of the session (i.e. it should be equipped with Internet access and a high-quality microphone).

The next step in obtaining user think-aloud data is to **brief think-aloud participants about the session**. The following guidelines were tested and refined during our three successive user pilot sessions, and appeared to be adequate (none of the pilot participants claimed in the de-brief interview that they found the instructions unclear, not did they raise any additional questions before commencing the think-alouds). Think aloud participants should be informed that they will be using the electronic resource to perform a number of information tasks and that they should think-aloud whilst performing the tasks. It should be explained to participants that they should not worry if they have not used the resource before, have not used this version of the resource before, or have not used it to perform similar tasks before, as prior familiarity with the tasks or with the resource is not necessary. Next, it should be explained to participants how to think aloud (i.e. that they should mention exactly what they are doing as they are doing it and that they should try to mention everything that is going through their heads) and that as the session is aimed at yielding insights that will hopefully make the resource easier to use, they should make sure to mention anything that you find easy/clear or difficult/unclear when using the resource. Think-aloud participants should be encouraged to explore the resource during each task and use any of the features within it that they think might help them perform the task. However, they should also be instructed to always maintain a task focus in order to avoid providing a 'tour' of the resource. It should be requested that think-aloud participants only use the resource under evaluation during the session and no other resources, pieces of software or websites. The facilitator should direct participants back towards the resource if they attempt to access another resource/piece of software/website.

Additionally, think aloud participants should be informed that the facilitator will not ask them any questions during the think-aloud session, although they might be prompted to explain what they are currently doing if they have not spoken aloud for a while. Participants should also be informed that the facilitator is not permitted to assist them in any way or answer questions about a task after they have commenced it. Participants should be instructed that in order to make the most of their session, they might be asked to move on to the next task if time is running short. Finally, it should be explained to think-aloud participants that the session is not a test of their performance or how they use the resource and that there is no 'right' or 'wrong' way of performing the tasks. Participants should be encouraged to perform each task in a way that is natural to them (i.e. a way that they might normally attempt this or a similar task when using an electronic resource). The instructions and guidance in these two paragraphs is available in bulleted list form in appendix 5, for use as an instruction sheet to give to think-aloud participants at the start of the session. After these instructions have been given to participants (preferably verbally as well as in writing), participants should be directed to ask any questions and to read and sign the informed consent form (see appendix 10 for an example form).

The third step in obtaining user think-aloud data is to **set tasks to participants**. The tasks that think-aloud participants will be asked to perform will depend on whether it has been decided to conduct a core, recommended or custom IB usability evaluation (see section a for a comparison of these three types of usability evaluations). In a 'core' IB usability evaluation, you should ask participants to perform the following three information-seeking tasks in table c below:

> **'Core' tasks in an IB usability evaluation:**
>
> 4. Gain access to the electronic resource
> 5. Find out which parts or sources within the resource you have access to
> 6. Think of some information that you currently need or have recently needed to find for your work and demonstrate, using the electronic resource, how you might go about finding it.

Table c: The three information-seeking tasks that think-aloud participants are asked to perform as part of a 'core' IB usability evaluation.

In a 'recommended' evaluation, you should ask participants to perform the three core tasks listed above, *plus any of the tasks in table d below that are currently supported by the electronic resource*:

> **'Recommended' tasks in an IB usability evaluation:**
>
> Gain an overview of an area by:
> - Trying to gain a basic understanding of the law relating to a particular legal area (e.g. Breaches of contract).
> - Trying to gain an appreciation of the importance of a certain legal journal author's role in a particular legal area.
> - Trying to locate a legal journal article written by an author who has published many articles or many highly cited articles.
>
> Gain a current or historical understanding of the importance of a document by:
> - Trying to find out whether a particular case is still good law.
> - Trying to find out what amendments have been made to a particular piece of legislation over a certain time period.
> - Trying to find out whether a particular piece of legislation is currently in force.
> - Trying to locate a historical version of a particular piece of legislation (i.e. a previous version that has since been amended).
>
> Maintain awareness of developments in an area by:
> - Trying to find out whether there have been any recent developments in a particular legal area (e.g. Discrimination law).
> - Trying to set up an alert so that you can be informed every time new documents are added to the system that match particular search terms (e.g. when new documents that match the term 'discrimination' are added).
> - Trying to set up an alert so that you can be informed every time there are new developments in a particular legal area.
>
> Return to any one of the tasks where you found useful documents and:
> - Determine which sections of a document that you have found are important to you.
> - Keep a softcopy (downloaded or saved) record of a document that you have found.
> - Keep a hardcopy (printed) record of a document that you have found.
> - Download two documents into a single file (i.e. you should end up with one file saved on the computer that includes the text of two separate documents, e.g. two different legal journal articles or two different sections of a particular piece of legislation).
> - Keep a softcopy or hardcopy record of *part* of a document that is important to you (e.g. print or download only certain parts of a case report).
> - Distribute a document that you have found, by e-mail, to a fictitious colleague from within the electronic resource.
> - Store a document on the server of the resource (i.e. save a copy of the document to a personalised area on the electronic resource itself, so you can access it again quicker in future).

Table d: Tasks that think-aloud participants are asked to perform as part of a 'recommended' IB usability evaluation. In addition to these tasks, participants are also asked to perform the three core tasks in table c.

In a 'custom' IB usability evaluation, you should present the tasks in tables c and d as default. However, the final choice of tasks depends on the focus of the usability evaluation. If the focus is on particular information behaviours, custom tasks related to these behaviours should be used. A list of custom tasks is presented in appendix 5 and example custom tasks related to the behaviour of chaining are presented in table e below:

Custom tasks related to 'chaining' behaviour:

Try to follow a hyperlink or other form of connection from:
- A legal case to a previous case or a particular piece of legislation mentioned in the case report.
- A piece of legislation to other pieces of legislation mentioned in the text of the Act.
- A legal journal or commentary article to a case or piece of legislation mentioned in the article.

Try to follow a hyperlink or other form of connection from a document to other documents that have been written *since* this document and have subsequently mentioned it. Specifically, try to:
- Find a particular case, then find out which *more recent* cases have mentioned it (if any).
- Find a particular piece of legislation, then find out which more recent pieces of legislation have mentioned it (if any).
- Find a particular legal journal article, then find out which articles that were written *after* this article have mentioned it (if any).

Table e: Custom think-aloud tasks related to chaining behaviour.

Both the recommended and custom tasks are based on the illustrative ways of supporting each behaviour/level presented in the supporting documentation in appendix 3 (which was itself devised by evaluating the functionality supported by a range of electronic legal resources and matching this to the data from our earlier empirical study of lawyers' information behaviour). These tasks are not exhaustive, however, at it may be possible to devise additional or alternative tasks related to particular behaviours/levels. Note that if any of the recommended or custom tasks are to be used to evaluate a *non* legal resource, it will be necessary to remove references to 'legal areas,' 'legal journal articles' etc. in the wording of the tasks and tailor the tasks to the new domain. This includes providing alternative examples of topical areas related to the domain. Also, setting tasks that are not supported in any way by the electronic resource under evaluation (i.e. not possible) should be avoided.

If think-aloud participants encounter difficulties logging in to the resource under evaluation, they should be allowed a few minutes to try alternative ways of logging in and then the facilitator should offer assistance. No other assistance should be offered by the facilitator, apart from verbally re-phrase the wording of a task if it is clear that a think-aloud participant has misinterpreted its meaning. If this is the case, the task may need to be re-worded in future sessions.

The facilitator is also responsible for **ensuring that the audio and user's screen is captured**. As highlighted by our user pilot sessions, it is suggested that the facilitator tests the recording facilities before the user commences their tasks. This includes setting the microphone volume, ensuring the computer has enough space on the hard drive to store a large video clip (several hundred megabytes) and creating a short test clip of the screen and playing back the resultant video clip to ensure that the screen and audio have been recorded properly. After the recording, the facilitator should also ensure that a meaningful filename is given to the video clip (perhaps including an anonymous code to identify the participant and the date and time of the evaluation) and that appropriate backups are made. If using regular video camera equipment to film the screen and audio, the facilitator should perform a similar 'dry run' and ensure the video media is appropriately labelled.

Conducting the usability evaluation itself – identifying usability issues from the think-aloud data

Usability issues should be identified by someone who has (or a pair who have) a basic knowledge of Human-Computer Interaction in general or other usability evaluation methods. Asking the developers of the resources themselves to perform the IB usability evaluation is discouraged, unless you can be fairly confident that their participation will not bias the output of the evaluation (i.e. the number, nature or subjective ratings of the usability issues identified).

Issues can be identified from the think-aloud data either *on-the-fly*, during the user think-aloud session or *after the session*, by playing back the video clip. Even if it is decided to identify issues during the session, an audio and screen recording should still be made as it may be difficult to keep up with the users' interactions with the resource or filling out the form that records the output of an IB usability evaluation. If identifying usability issues on the fly, the evaluator present at the think-aloud session may also act as facilitator. If identifying issues after the session, by playing back the audio and screen recordings, it may be necessary to pause and re-wind parts of the think-aloud session in order to get an adequate grasp of the user's interaction with the resource and to allow enough time to fill out the usability form (as illustrated by the need for the developer in our developer pilot to pause and re-wind the video clip often). Although it may be possible for the facilitator to edit the video clip to remove parts where similar user interactions are repeated multiple times or the same problems are encountered more than once, we discourage editing the video clips at all as it is a time-consuming process and may affect the perceived credibility of the think-aloud session amongst other members of the group.

Whether identifying usability issues during or after the user think-aloud session, a **usability form should be filled in to record the output of the evaluation** (see appendix 8 for the form). As with the summary functionality form, the usability form allows usability evaluators to record a number of details about the evaluation: the name and version number of the electronic resource being evaluated, whether this is a core, recommended or custom evaluation, the date of the evaluation, the video filename that accompanies the evaluation, an anonymous identification code for the participant that took part in the think-aloud session, the name of the facilitator present and the name of the evaluator (i.e. the person conducting the analysis of the think-aloud data). As with the summary functionality form, there is space on the form to list the parts of the resource that the usability evaluation has been restricted to and to indicate which set or sets of behaviours the evaluation has been restricted to (if applicable). There is also space to (optionally) record details about how experienced the user considers themselves to be with using (a) general electronic resources in their domain and (b) the electronic resource currently under evaluation. This should be recorded on a scale of 'not at all experienced/somewhat experienced/very experienced.' This information is likely to be useful if seeking to recruit users with a variety of experience using the resource under evaluation (or if seeking to recruit users that regard themselves as particularly experienced or inexperienced in using the resource).

The remainder of the usability form should be used to record the usability issues identified from the think-aloud data and other relevant details. Specifically, users of the IB usability method should use the usability form in appendix 8 to record:

- The task number (from the list of tasks in appendix 10) that the user was attempting when the usability issue was identified. Alternatively this column can be filled with an abbreviation for the task rather than the number (e.g. 'subsequent case' for 'find a particular case, then find out which more recent cases have mentioned it, if any').
- The user actions, user comments, or evaluator observations that might suggest a usability issue.
- The approximate time in the video clip the actions/comments/observations occurred. This need only be an approximate pointer to part of the video clip in order to aid reviewing the video footage later. It does not need to be exact.

The above details should be recorded whilst watching the user think-aloud session (regardless of whether the evaluator is watching a live or recorded version of the session). Other details to be recorded on the form include:

- The usability issues identified from the actions/comments/observations.
- The screen(s)/page(s)/part(s) of the resource that the actions/comments/observations relate to.
- The evaluator's perception of the severity of the usability issue (i.e. not severe/does not need immediate attention, quite severe/needs attention in the future or very severe/needs immediate attention).
- The evaluator's perception of the amount of effort required to address the usability issue (i.e. small, medium or large).

These details can be recorded either whilst watching the think-aloud session or after watching the session. This is because some degree of reflection may be required in answering these questions. In particular, the evaluator's perceptions of the severity of an issue and the amount of effort required to address it might change as the session progresses. Therefore users of the method should also be encouraged to review any of the above details that were recorded whilst watching the session at the end of the session and amend the details as they see fit.

Finally, the form can also be used to record separate evaluator reflections on the usability issues identified (such as broad themes arising from issues identified, issues for future discussion etc.). As these are notes based on the evaluators reflections on the issues, these will usually (although not always) be noted *after* watching the session.

As illustrated by our developer pilot, sometimes it may be difficult to decide whether to note particular user actions or comments (or indeed personal observations) on the form. We suggest making a note of all actions, comments or observations that you believe might highlight a usability issue, even if you have doubts. Often user actions or comments will become clearer as the interaction (and the think-aloud session) progresses. In particular, we encourage a broad interpretation of the term 'usability issue.' By this, we mean that instances where the task was not directly supported by the resource and instances where the user did not know how to go about performing a particular task, even though the task was possible should be noted. In addition, we also recommend noting any usability issues that have been inferred from user actions, comments or apparent misconceptions or misunderstandings surrounding the resource. Similarly, we recommend noting issues that are unlikely to be addressable (for whatever reason) and small-scale usability issues, which may be minor but have caused the user problems nonetheless.

Appendix 3: Illustrative ways that electronic legal resources might support information behaviours at each applicable level

Listed below are several illustrative ways that an electronic resource might support each behaviour, at each applicable level. The majority of these examples were obtained by inspecting six electronic legal resources that were used by some of the academic and practicing lawyers in our study. These resources were:

- LexisNexis Butterworths (abbreviated below as LNB). Available online: http://web.LexisNexis.com/professional [Accessed 25/05/07].
- LexisNexis Professional (abbreviated below as LNP). Available online: http://www.lexisnexis.com/uk/legal [Accessed 25/05/07].
- Westlaw (abbreviated below as W). Available online: http://www.westlaw.com [Accessed 25/05/07].
- Kluwer Arbitration (abbreviated below as K). Available online: http://www.kluwerarbitration.com [Accessed 25/05/07].
- Justis 5 (abbreviated below as J). Available online: http://www.justis.com [Accessed 25/05/07].
- HeinOnline Journals (abbreviated below as H). Available online: http://heinonline.org [Accessed 25/05/07].

A few of the examples listed below are purely illustrative of ways in which future electronic legal resources *might* support a particular behaviour or level (and hence were not identified as supported by any of the electronic resources above). These examples are denoted by a ❍ symbol and are included to illustrate that (a) some behaviours such as analysing are performed by lawyers bur are not directly supported by the electronic legal resources above and (b) it is possible to identify additional ways that a particular behaviour/level might be supported by an electronic resource.

The abbreviated name of each electronic legal resource that was found to support each way or performing a particular behaviour is listed next to each example. For example, all of the six resources 'provide a login facility that asks for credentials such as a username and password' (see below). Some of the examples are illustrated by screenshots from the Justis electronic legal resource. These screenshots have been included with permission from the management at Justis UK.

Accessing - the process of gaining access to an electronic resource, or to sources or documents and content within a resource, for example by logging in.

Illustrative ways of supporting **accessing** electronic *resources* and particular *sources* and *documents* within them include (but are not restricted to):

- Providing a login facility that asks for credentials such as a username and password. **LNB/LNP/W/K/J/H.**
- Providing support for logging in through a third-party site or resource (e.g. such as the Athens devolved login for educational users). **LNB/LNP/W/K/J/H.**

The 'personal user' sign in credentials page in Justis (which allows users to log in to their own personalised version of Justis). Also note the 'Athens user sign in' button, which allows devolved login through Athens.

- Providing automatic access without a noticeable access procedure through recognition technologies (such as IP recognition). **LNB/LNP/W/K/J/H.**

Surveying - the initial search for information to obtain an overview of sources or documents within a subject field, or to locate key people operating in this field.

Illustrative ways of supporting **surveying** *sources* in an electronic resource include (but are not restricted to):

- Providing a summary of the topical coverage of sources. Providing other overview information alongside the summary may also be useful, such as a description of the types of documents within the source and details of when the source's coverage of the document begins and, if appropriate, ends. **LNB/LNP/W/K/J/H.**

The Times Law Reports

Coverage	1990 – Present
Format	Internet & CD ROM
Data Provider	Times Newspapers Limited
Frequency	Daily (as appropriate)

A fully searchable database of all cases published in The Times newspaper since 1990.

The Times Law Reports contains the full text of every case report in *The Times* newspaper since 1990. The cases cover every branch of law and come from a wide range of jurisdictions, including the UK courts, the European Court of Human Rights, the Court of Justice of the European Communities and the International Court of Justice in The Hague.

A summary of the topical coverage of 'The Times Law Reports' source within Justis.

Illustrative ways of supporting **surveying** *documents* include (but are not restricted to):

- Providing access to secondary documents that provide an overview of a subject field by consolidating a number of primary documents. In the legal domain an example of secondary documents are commentary articles, which often provide overviews of particular legal topics. Not only might we designate a main part of the electronic resource to provide access to secondary overview documents, but we might provide hyperlinks to overview documents within the text of documents that are related to the overview topic. **LNB/W/K/J/H.**
- Providing support for searching for documents in a particular subject field to gain an appreciation of the area, or for important people in the field to gain an appreciation of their role in the area. This support may be specifically focused on surveying (e.g. tailored to bring back key overview documents in a

particular area or documents relating to key people in a particular area) or may support surveying incidentally **(e.g. LNB/LNP/W/K/J/H).**

Search In Articles and Support Material

☑ HERMES   ☐ BLISS

Clear

ℹ Title

ℹ Full Text

ℹ Subject   competition

ℹ Date

Search

A Justis search for articles by subject (in this case, aiming to find articles in the area of Competition Law).

- Providing support for browsing for documents in a particular subject field to gain an appreciation of the area, or for important people in the field to gain an appreciation of their role in the area. Once again, this support may be specifically focused on surveying (e.g. tailored to facilitate browsing of key overview documents or documents relating to key people in a particular area) or may provide incidental support for surveying **(e.g. LNB/LNP/W/K/H).**

Monitoring - maintaining awareness of developments in a field.

Illustrative ways of supporting **monitoring** *sources* in an electronic resource include (but are not restricted to):

- Providing an indication when new sources have been added to the resource (for example through the use of a 'new' icon whenever the source name is displayed or by announcing the new sources in a part of the site dedicated to monitoring sources and/or documents. **⬤**

Illustrative ways of supporting **monitoring** *documents* include (but are not restricted to):

- Providing a dedicated section of the electronic resource access to showcase recently published documents (of all types that the electronic resource contains) on a particular topic. This might allow users to monitor by the age or date of publication of documents on a particular topic, the type of document (e.g. in the legal domain, case reports/commentary/legislation/journal articles) and/or, for legal documents, the jurisdiction. This dedicated section of the resource might support monitoring by one of these aspects at a time, or allow users to combine monitoring criteria (e.g. to look at all pieces of legislation published in the last week). **LNB/W.**
- Providing support for conducting similar, regular searches for documents in a particular subject field to maintain awareness of any developments in the area. Like surveying, this support may be specifically focused on monitoring (e.g. tailored to only search for recent materials or new materials since the last search) or may support monitoring incidentally. **LNB.**
- Providing support for regularly browsing for documents in a particular subject field to maintain awareness of any developments in the area. Once again, this support may be specifically focused on monitoring (e.g. tailored to facilitate browsing of recent materials on a particular topic) or may provide incidental support for surveying. **LNB.**

Illustrative ways of supporting **monitoring both** *sources* **and** *documents* include (but are not restricted to):

- Providing e-mail alert services. These services might draw users' attention to new documents or sources related to pre-defined topics of interest **(e.g. LNB/W)** or list new documents (of a specified type) that match pre-defined search terms of interest **(e.g. LNB/W/J).** Not only might we designate a main part of the electronic resource to provide access to subject-related highlights of new materials that have been added to the resource, but we might allow any search or browse operation to be set up as an alert so that details of any new documents matching the search terms or added to the browsing category of interest are sent as an e-mail alert to the user **(e.g. LNB/W/J).**

E-mail alerting options in Justis. In this case, an e-mail will be sent when new documents with a title beginning with 'lotter' are added to the resource.

Searching - formulating a query in order to locate information.

Illustrative ways of supporting **searching** (for either *sources*, *documents* or *content*) in an electronic resource include (but are not restricted to):

- Providing one or more global or local Boolean or natural language search fields that can be used to formulate and submit a query. These fields may restrict the search to particular types of documents (e.g. for the legal domain, case reports, legislation, journal articles etc.), to particular time periods (e.g. for a piece of legislation, the enactment date) and, only for legal documents, to a particular jurisdiction (e.g. EU law). As well as searching the full-text of the document, these fields may also restrict the search for the query terms entered to a particular part of the document or meta-data. For legal cases, searches may be restricted to a particular case citation, to particular party names, judges presiding, counsel involved, keywords/headnotes, levels of court, other cases or legislation cited etc. For legislation, searches may be restricted to the title of the Statute or Statutory Instrument, particular provisions within the legislation and particular keywords relating to the legislation. For legal journal articles, searches may be restricted to a particular journal or article title, particular author names, particular keywords and particular legislation or cases cited. **LNB/LNP/W/K/J/H provide Boolean segmented field search facilities. W provides a Natural Language Search.**

The default segmented search fields provided by Justis when searching for UK legal cases.

- Providing a query wizard to guide users through a series of segmented field searches and search restrictions (see examples above). **W.**
- Providing a simple text-matching search field (particularly relevant for supporting searching within a particular document for a particular word or phrase). **LNB/LNP/W.**

Browsing and extracting – browsing involves semi-directed searching for sources, documents or content. Extracting involves systematically working though a particular resource to identify sources of interest, a particular source to identify documents of interest and/or a particular document to identify content of interest.

Illustrative ways of supporting **browsing,** and **extracting** in an electronic resource include (but are not restricted to):

For **browsing** within resources to extract *sources*:

- Providing a local or global facility for browsing by source (e.g. a facility to facilitating browsing a list of the different journal titles or databases of cases/legislation available in the current electronic resource, for the current user). This facility might also allow users to continue to browse within a particular source for documents. **LNB/LNP/W/K/J/H.**

Directory tree illustrating browsing Justis by source. In this case, the source 'The Weekly Law Reports' has been selected and the volumes of the source (starting in 1953) are listed in the expanded tree.

For **browsing** within resources or sources to extract *documents*:

- Providing a local facility for browsing by various aspects of document meta-data (i.e. an individual part of the resource designed to support browsing). **LNB/W/K/H.** This includes browsing by the type of document (ideally all types of document that the electronic resource contains), by particular time periods (e.g. the date cases were heard or the date a legal article was published) and, only for legal documents, by jurisdiction. For legal cases, it may also be possible to facilitate browsing by party name, judge presiding, counsel involved, keyword/headnote, level of court etc. For legislation it may be possible to facilitate browsing by the title of the Statute or Statutory Instrument or particular keywords relating to the legislation. For legal journal articles it may be possible to facilitate browsing by journal or article title, author name, keyword and journal volume or issue number.
- Providing a global facility for browsing by various aspects of document meta-data (accessible from most or all parts of the electronic resource). This might involve, for example, providing hyperlinks within documents to other similar documents (e.g. documents about the same topic, written by the same author etc.). For legal cases, this might include the facility to find all other cases involving one of the parties in the currently displayed case, all other cases with the same judge presiding, same counsel involved, classified by the same keyword as used to classify the current article etc. For legislation, this might include the facility to find all other pieces of legislation classified by the same keyword as used to classify the current piece. For legal journal articles, this might include the facility to find all other articles in the current volume/issue of the journal, all other articles written by the same author or all articles classified by the same keyword as used to classify the current article etc. **O**

For **browsing** within documents to extract *content*:

- Highlighting instances of the search terms in the document text. **LNB/W/K/H.**
- Providing the facility to highlight particular words or phrases in the document text. **LNB/LNP/W.**
- Providing the facility to jump between instances of the search terms or particular words or phrases in the document text. **LNB/J.**

Jumping between instances of the search term 'beer' in the current document loaded in Justis.

- Providing the facility to jump to particular section headings in the document text. **J.**

Jumping to the 'headnote' section of the currently displayed law report in Justis. The highlighted block of text is the headnote of the case.

Chaining - following chains of citations or other forms of referential connections between sources or documents.

Illustrative ways of supporting **chaining** between *sources* include (but are not restricted to):

- Providing hyperlinks that link between sources (for example between a source and another that precedes or supersedes it). These hyperlinks can be presented within documents from each source or in another part of the electronic resource that describes the content of the source. ⭘

Illustrative ways of supporting **chaining** between *documents* include (but are not restricted to):

- Providing hyperlinks that link between related documents. These hyperlinks can be presented in the document text or on a summary or overview page.
    - o We might want to support chaining from a document to previously written documents that are cited in the document text (known as backwards chaining). **LNB/W/K/J/H.**
    - o We might want to support chaining from a document to documents which have been written subsequently that have cited it (i.e. documents which the document has been cited *by*). This is known as forwards chaining and may also be supported by a dedicated citator service (which catalogues and presents the connections between documents). **LNB/W/J.**

The 'subsequent cases' option in Justis, which allows subscribers to their sister resource, JustCite, to find out which cases have subsequently cited the current case.

- o We might also want to support chaining between documents on similar topic areas, written by the same author (or for legal cases, where the same judge has presided over the case) by providing hyperlinks within a document that will find documents based on these or similar relations based on the meta-data of the document. ⭘

Selecting, **distinguishing** and **filtering** - different ways of choosing relevant information (i.e. resources, sources or documents). Selecting involves carefully choosing resources, sources or documents as being potentially useful for the information task at hand (based on own or shared perceptions). Distinguishing is similar to selecting, but involves *ranking* information sources or documents according to their relative importance (again based on own or shared perceptions). This means deciding that one or more source or documents from a group is likely to be more useful than the others. Filtering involves the use of certain criteria or mechanisms when searching or browsing for information to make the information as relevant and as precise as possible (for example restricting a search to return documents by a particular author).

Illustrative ways of supporting s**electing and distinguishing and filtering** *sources* include (but are not restricted to):

- Providing a summary of the topical and date-related coverage of sources. For example, this might include a description of the types of documents within the source, the general subjects covered by the documents and when the source's coverage of the document begins and, if appropriate, ends. **LNB/LNP/W/K/J/H.**

Illustrative ways of supporting **selecting and distinguishing and filtering** *documents* include (but are not restricted to):

- Providing useful document meta-data for each type of document to be displayed as part of the search results or when browsing for documents. For legal cases, this meta-data might include the party names involved, the date the case was heard, the jurisdiction of the case, the level of court, the name(s) of the judge(s) presiding, the names of the counsel involved, the official legal citation, For legislation, this meta-data might include the title of the Statute or Statutory Instrument and the date the legislation came into force (or whether it is not yet in force) and the date it ceased being in force (if applicable). For legal journal articles this meta-data might include the journal or article title, author names and the official citation. For all types of document, this meta-data might include a short summary or abstract, an indication of how heavily cited the document is, keywords relating to the document and/or a snippet of the search query term in the context of the document. **LNB/LNP/W/K/J/H all provide some document meta-data at the search result stage.**

The meta-data provided by Justis as a result of searching for cases with the term 'beer' in the subject field. This meta-data includes the citation of the case, the title and subject (keywords) and the year the case was heard.

- For legal cases and legislation, by providing iconic indicators of the positive, negative or neutral treatment of each case or piece of legislation when search results are displayed or when browsing for documents. For cases, icons might signify whether a case has been affirmed, superseded, followed, or overruled by later cases as an indication of whether or not it is still good law. For legislation, icons might signify whether a case is not yet in force, currently in force or no longer in force. These icons

might be placed in search results lists (and on lists that the documents can be browsed by) as well as within the documents themselves. **LNB/W.**

- Providing the option for users to comment on or submit ratings on the usefulness of a source or document. It is, for example, possible to critique a judge's decision in a case, identify gaps in coverage in a particular piece of legislation or to rate the overall comprehensiveness of a commentary source. It is also possible to comment on or rate the perceived authority of a particular source (e.g. a particular journal series). **⬤**
- Providing dynamic filtering options when search results are displayed or when browsing for documents. This dynamic filtering may be by document type, title, date, keywords etc. **J.**
- Providing sorting options when results are displayed or when browsing for documents. It may be possible to sort the document or results list by document title, document type, date etc. and also by many other possible items of meta-data. **LNB/K/J/H.**

		Reference	Title and Subject	▼
☐			⚙ Sorting...	
			~~Weigh~~ Sort by title ~~easures – Sale of goods – Sale by weight or measure~~	
☐	≡	[2006] 3 WLR 148	Crehan v Inntrepreneur Pub Co (CPC) and another (Office of Fair Trading and others intervening) [House of Lords] European Community – Free competition – Restriction or distortion of competition	
☐	≡	[2003] 1 WLR 1331	Fox v Customs and Excise Commissioners [Queen's Bench Division] Revenue – Customs and excise – Importation of dutiable goods	
☐	≡	[1966] 1 WLR 410	Hall v. Hyder [Queen's Bench Division] Licensing – "Intoxicating liquor" – Mixture	
☐	≡	[1972] 1 WLR 90	McDonald v. Howard Cook Advertising Ltd. [Queen's Bench Division] Town Planning – Advertisements, control of – Permitted display	
☐	≡	[1968] 1 WLR 729	Reg. v. Leicester Licensing Justices, Ex parte Bisson [Queen's Bench Division] Licensing – Justices – Jurisdiction	

Sorting search results in ascending order of title in Justis.

- Providing the facility to search within the current search results. **LNB/LNP.**
- Providing the facility to highlight and/or jump between instances of search terms in the document text. **LNB/LNP/W/K/J/H.**
- Providing the facility to highlight and/or jump between instances of user-defined words or phrases in the document text. **W.**

Updating - gaining a *current* understanding of the importance of a particular legal *document* (e.g. an understanding of whether a particular case is still good law or a particular piece of legislation is currently in force). **History tracking** - gaining a *historical* understanding of the importance of a particular legal *document*. (e.g. an understanding of the judicial history of a particular case that has been heard at several levels or court or an understanding of how a particular piece of legislation has been amended over time).

Illustrative ways of supporting **updating** and **history tracking** of *documents* in an electronic resource include (but are not restricted to):

- Providing incidental support for searching for and/or browsing between documents that are temporally related (e.g. for cases that have affirmed, superseded, followed, or overruled a particular case of interest). Incidental support can be provided for updating and history tracking even when the search tools within the electronic resource are not tailored for supporting these behaviours (and therefore are only aimed at facilitating searching in general). **LNB/W/J.**
- Providing explicit support for searching for and/or browsing between documents that are temporally related. This can be achieved through the provision of dedicated tools (such a citator service) to provide details about the current status and/or history of a particular document (for example, the positive, negative or neutral treatment of a particular case or secondary source, details about whether a particular piece of legislation is not yet in force, currently in force or no longer in force and a history of

amendment details for a particular piece of legislation). History details can be presented in either text or diagrammatic form. **LNB/W.**

- Providing current status and/or history details of documents within the document itself or as part of an overview/summary page. This involves providing similar details to those described above and, in effect, involves integrating a citator service into the search and/or browse functions of the electronic resource so that its operation is transparent. **LNB/W/J.**

The 'cases considered' tab in Justis which lists history details related to the currently displayed case (which, in this example, is the 1966 Hall v. Hyder Queen's Bench case). The 'cases considered' tab shows that the case has been 'mentioned' by two subsequent cases.

- Providing access to the current *and* historical versions of documents (such as pieces of legislation) or access to only the current versions of documents but include mark-up or annotations to explain the history of changes that has been applied to the document. **LNB/LNP/W/K/J do the latter (and LexisNexis Butterworths allows users to request historical versions of legislation by e-mail).**
- Providing iconic indicators of the positive, negative or neutral treatment of each case or piece of legislation when search results are displayed or when browsing for documents. For cases, icons might signify whether a case has been affirmed, superseded, followed, or overruled by later cases as an indication of whether or not it is still good law. For legislation, icons might signify whether a piece of legislation is not yet in force, currently in force or no longer in force. These icons might be placed in search results lists (and on lists that the documents can be browsed by) as well as within the documents themselves. **LNB/W.**

Analysing - examining in detail the elements or structure of the *content* found during information-seeking.
Synthesising – combining the elements of the content found during information-seeking into a coherent whole.

Illustrative ways of supporting **analysing** and **synthesising** *content* include (but are not restricted to):

- Providing the facility within the electronic resource to make notes about particular documents or add notes to document text. **O**
- Providing the facility within the electronic resource to arrange notes made about documents under different categories, topics or headings. **O**
- Providing the facility to export particular parts of the document to a word processing package. **O**

Recording - making a record of resources or sources used, of documents or content found or of the query terms used or results returned in a search.

Illustrative ways of supporting **recording** *sources* include (but are not restricted to):

- Providing the facility to save a list of sources that are frequently used by an individual user or those deemed by the user to be useful. **LNB/LNP.**

- Providing the facility to customise the interface of the resource to showcase those sources that are frequently used by an individual user or deemed by the user to be useful. **LNB/LNP/W/K. LNP and K only allow the start screen to be customised to show different types of legal material.**

Illustrative ways of supporting **recording *documents*** and ***content*** in an electronic resource include (but are not restricted to):

- Providing the facility to download entire documents or parts of documents. These parts may be user-defined parts of interest, particular section headings, particular page numbers etc. The facility may be provided to download documents in a variety of formats (e.g. Adobe Portable Document format, Microsoft Word format, HTML format etc.). **LNB/LNP/W/K/J/H allow entire documents to be downloaded. LNB/W/H allow parts of documents to be downloaded.**
- Providing the facility to print entire documents or parts of documents. **LNB/LNP/W/K/J/H allow entire documents to be printed. LNB/W/H allow parts of documents to be printed.**
- Providing the facility to e-mail entire documents or parts of documents from within the electronic resource itself (as opposed to within a separate e-mail client). **LNB/LNP/W/J allow entire documents to be e-mailed. W allows parts of documents to be e-mailed.**
- Providing the facility to store entire documents or parts of documents on the server of the electronic resource. **W/J.**

Saving a document on the Justis server. Documents that are saved can be accessed later at will from the 'my justis' tab in Justis).

- Providing the facility to note which source a particular document came from or which document particular content came from, perhaps along with the date and time that the document or content was recorded. **O**
- Providing the facility to highlight or otherwise distinguish content of interest within a document that has been downloaded, printed, e-mailed or stored on the server. **O**

Illustrative ways of supporting **recording *search query terms or results*** include (but are not restricted to):

- Saving an automatic search history trail that records the date and time of the search, the search query terms used, any restrictions the user has placed on the search, the number of hits returned and, perhaps, takes a snapshot of the results returned and provides an indication of which results were selected by the user. These searches can then be revisited to remind the user of what they have previously searched for and re-run if necessary. **LNB/W/J/H.**

The automatically-saved search history trail in Justis. Clicking on the hyperlinks re-runs the search.

- Providing the facility to manually save the current search (along with some or all of the above details such as the date and time, search query terms used etc.). As with an automatic search history, these searches can then be revisited to remind the user of what they have previously searched for and re-run if necessary. **LNB/W/J.**

The saved search list in Justis.

Collating – the physical act of drawing together documents and/or content for later use.

Illustrative ways of supporting **collating** batches of *documents* include (but are not restricted to):

- Providing the facility to download/export a batch of documents and collate them in one document. **LNB/LNP/W.**
- Providing the facility to download/export a batch of documents as an archive (e.g. ZIP) file. **J.**

Downloading a batch of three documents (the ticked search results) into a ZIP file from within Justis.

- Providing the facility to print a batch of documents and collate them in one hardcopy. **LNB/LNP/W/J.**

239

- Providing the facility to e-mail a batch of documents and collate them in one message, perhaps as attachments. **LNB/LNP/W/J.**

E-mailing a batch of documents (the three ticked search results) from within Justis.

Illustrative ways of supporting **collating** parts of documents (i.e. *content*) include (but are not restricted to):
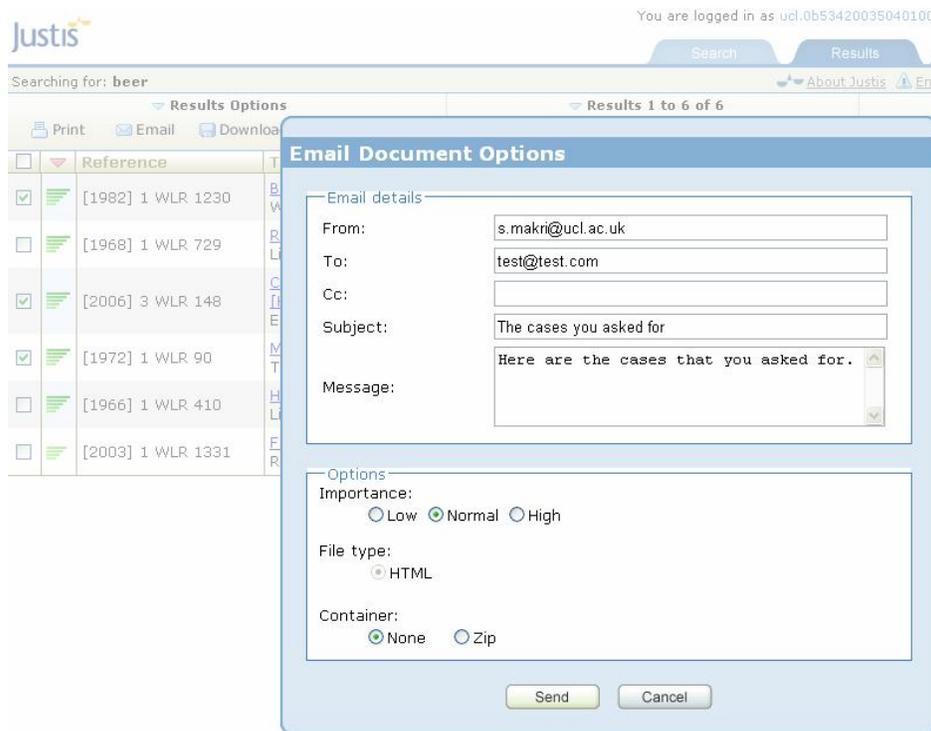
- Providing the facility to download/export particular sections of a document and collate them in one document. **LNB/W/H.**
- Providing the facility to print particular sections of a document and collate them in one hardcopy. **LNB/W/H.**
- Providing the facility to e-mail particular sections of a document and collate them in one message or attachment. **W (although only options to e-mail all pages containing search terms or first page only are given).**

Editing – preparing and arranging documents and/or content for later use by making revisions or adaptations.

Illustrative ways of supporting **editing** *documents* include (but are not restricted to):

- Providing the facility to download/export documents into an editable format (for example Microsoft Word format) for later editing. **LNB/LNP/W.**
- Providing the facility to copy parts of documents of interest to the computer clipboard for later use in other packages. This might involve adding a clipboard feature to the electronic resource itself (as opposed to using the generic computer clipboard) in order to support the storage of more than one chunk of text at a time in the clipboard's memory. **O**
- Providing the facility to edit and annotate document text from the within the electronic resource (and then to download or print a record of the edited version of the document). **O**

240

Distributing – handing or sharing out entire documents, particular content or search queries/results to others.

Illustrative ways of supporting **distributing** batches of or individual *documents* include (but are not restricted to):

- Providing the facility to e-mail a batch of documents or individual documents to others from within the electronic resource itself (as opposed to within a separate e-mail client). **LNB/LNP/W/J.**
- Providing the facility to fax a batch of documents or individual documents to others from within the electronic resource itself. **J (although this functionality has since been discontinued).**
- Providing the facility to print a batch of documents or individual documents for manual distribution from within the electronic resource itself (as opposed to from the Internet browser itself, which may not provide adequate print options). **LNB.**
- Providing the facility to store entire documents on the server of the electronic resource and *share access* to these documents or parts of documents with other users. These documents may either be identical to the versions already stored on the electronic resource or annotated/edited by the user. **O**

Illustrative ways of supporting **distributing** parts of documents (i.e. *content*) include (but are not restricted to):

- Providing the facility to e-mail particular sections of documents to others from within the electronic resource itself (as opposed to within a separate e-mail client). **O**
- Providing the facility to fax particular sections of documents to others from within the electronic resource itself. **O**
- Providing the facility to print particular sections of documents for manual distribution from within the electronic resource itself (as opposed to from the Internet browser itself, which may not provide adequate print options). **LNB/W/J.**

Selecting particular parts of the currently displayed document to print from within Justis.

- Providing the facility to store particular sections of documents on the server of the electronic resource and *share access* to these documents or parts of documents with other users. **O**

Illustrative ways of supporting **distributing** *the search history or particular saved searches to others* include (but are not restricted to):

- Providing the facility to e-mail the search queries used in selected searches and the number of hits returned and (optionally) a list of the results returned with hyperlinks to the documents listed to others from within the electronic resource itself (as opposed to within a separate e-mail client). **O**

Appendix 4: List of usability issues identified by the lead developer of the IB usability method and by the participants in our evaluation

The tables below list the usability issues identified by the lead developer of the IB usability method and the participants in our evaluation study. These issues were identified from the user pilot sessions (i.e. from the pre-collected think-aloud data provided to participants that took part in the evaluation as part of the one-day tutorial on the IB methods). Issues listed in the second column have been transcribed verbatim from the IB usability forms filled out by the participants. Issues identified by participants have been paired alongside any relevant corresponding issue identified by the lead developer of the method from the same data (provided the developer had previously identified the issue from the data). Although many participants included time references to the video clip so that we could ensure that they were referring to particular user comments/actions, this information was not *always* provided. We do not believe this had much of an effect on the issue pairing process, although the second column of the tables below should be regarded, strictly speaking, as an *interpretation* of participants' output with the aim of collating the data.

Recommended Tasks

1. Try to find out whether a particular case is still good law.

Usability issues identified by lead developer of the IB usability method	Similar usability issues identified by stakeholders working for a large electronic legal resource development firm	Additional usability issues identified by stakeholders (that were not identified by the lead developer)
Confusion over what the Table of Context bar displays during a search for cases. First, participant assumed that it listed other cases that were mentioned in the case report displayed on-screen. Then participant assumed that the TOC was a results list (as his search had returned only 1 document and therefore jumped straight to displaying the document full-text along with the TOC).	P7 (paired with P1): Unsure of what list is.	P1 (paired with P7): Too broad terms. Too narrow. Query formulation issues. P7 (paired with P1): Initially can't find info. Needed. Query formulation issues. P9: Issues formulating search. Producing either too few or too many results. System not providing sufficient support in search query formulation. P7 (paired with P1): Had to add search item before results were valid. Meaning unclear even after reference to the part of the video clip indicated. P9: Needs to search again to find a case previously viewed. Recent documents function not found by user.
Not clear how to find citator details for a particular case (i.e. to find out whether a case is still good law or not). Participant read the 'find a case in CaseSearch' option aloud and still did not associate this with the task. Also unclear how to *directly* find out whether a particular case is still good law.	P9: Can't find citing documents. Citator document not visible. P9: "Presuming this is up-to-date."	
How to use the 'Get a specific document' combo box is unclear. Participant was unaware of the necessity to type over the example text in grey.		
Participant used the 'general search' field (rather than segmented fields) for all searches and did not use any Boolean operators. Perhaps the other search possibilities were not made prominent at the interface.	P1 (paired with P7): Can't find segments (visibility). Didn't see. P7 (paired with P1): No filters on search terms (title) etc. (Can't find them). Design issues. P9: Can't find guided search form for cases. Search form tables not visible or function not obvious to user.	
Participant was not sure how to expand the Table of Contents side menu bar when this bar is presented.	P1 (paired with P7): Horizontal scrolling.	
'Back' button is sometimes non-responsive when navigating between pages.	P1 (paired with P7): Back button clicked twice. P7 (paired with P1): Back button.	

2. Try to find out whether a particular piece of legislation is currently in force.

Usability issues identified by lead developer of the IB usability method	Similar usability issues identified by stakeholders working for a large electronic legal resource development firm	Additional usability issues identified by stakeholders (that were not identified by the lead developer)
Initially unclear how to find only pieces of legislation and no other types of legal materials.		
Confusion over whether 'find source' would allow the participant to search for a particular piece of legislation.	P7 (Paired with P1): Searching on sources unclear if in source or within. Possibly clarity issue on search sources. P9: Can't find where to search legislation from. 'Find a source' search logic unclear – feature and function unclear.	
Searching for a particular Act in general search does not necessarily bring the Act with the same title as the search terms back in a high position in the results list.		
Not immediately clear that it is possible to filter a general search by result group. Participant was not aware at first that when looking for the text of the Unfair Contracts Act 1977 with search terms 'Unfair Contracts Act,' it would be necessary to filter by the group 'Unfair Contracts Act 1977' to bring back documents specifically relating to the Act.	P1 (Paired with P7): Visibility of sub-tabs. P7 (Paired with P1): Again unclear on functionality of side bar.	
Some confusion over whether a particular act is still in-force. Participant was not sure whether a lack of 'overturned' in the text of the Act meant it was still in-force, but assumed so.	P1 (Paired with P7): Can't determine if in force but assumes yes. P7 (Paired with P1): Unclear if still in force. Works on presumption. Clarity issue. P9: No clear statement as to whether legislation still in force. Lack of clarity and currency of data.	

3. Try to find out whether there have been any recent developments in a particular legal area.

Usability issues identified by lead developer of the IB usability method	Similar usability issues identified by stakeholders working for a large electronic legal resource development firm	Additional usability issues identified by stakeholders (that were not identified by the lead developer)
Confusion over whether and how the resource supports finding out whether a particular case is still good law.	P7 (paired with P1): No clear route to check amendments. Shown too late in process.	P1 (paired with P7): > 3000 results returned.

P1 (paired with P7): Didn't really complete task.

P7 (paired with P1): Left bar now showing different functionality. |
Confusion over where to click within the breadcrumb trail to return to the current search results.		
A lack of warning of misspellings in search terms and that two subsequent searches, with the same search terms and syntax but different capitalisation of search terms, had the same effects and had brought back the same results.		
Confusion over how to search within Family Law sources and what the effect of restricting a search to 'all subscribed Family sources' would be.	P7 (paired with P1): Different functionality on hyperlink and checkbox.	
Confusion over which legal practice area it would be possible to find the legal area of 'crime' under.	P9: Not obvious what sources being searched from practice area. Source selection unclear.	
Confusion over whether a document had already been clicked on in the current search session.		

4. Try to set up an alert so that you can be informed every time there are new developments in a particular legal area.

Usability issues identified by lead developer of the IB usability method	Similar usability issues identified by stakeholders working for a large electronic legal resource development firm	Additional usability issues identified by stakeholders (that were not identified by the lead developer)
Confusion over where to go in order to set up a scheduled search for particular search terms (and, related to this issue, Initial confusion over what the 'update wizard' would allow the participant to achieve). Both of these issues are symptoms of confusing setting up a 'scheduled search' with setting up an 'alert.'	P1 (Paired with P7): Conceptual clarity between alerts and saved searches.

P9: Doesn't immediately find 'alerts' (under my research). Function not visible and obvious. Not apparent how to set up alerts/legal updates. Functionality not visible and difference between legal updates and scheduled searches unclear. | P1 (Paired with P7): Query formulation.

P9: Difficulty formulating search: either too many or too few results.

P7 (Paired with P1): 'Last 10' label unclear. Meaning unclear

P7 (Paired with P1): 'Tenancy' and 'Tenant' should be treated as similar terms, but are very different. |

5. Try to download two documents into a single file.

Usability issues identified by lead developer of the IB usability method	Similar usability issues identified by stakeholders working for a large electronic legal resource development firm	Additional usability issues identified by stakeholders (that were not identified by the lead developer)
Not immediately clear how to download/save a document or if a document has actually been downloaded saved. Related to these issues, it is also not clear what 'view tagged' will allow the participant to achieve. Participant assumes it will open both the selected documents on the screen in a collated format.	P7 (Paired with P1): No apparent way to download. Icons too small.	P1 (Paired with P7): Not clear what checkboxes. Meaning unclear and no time reference point provided.
Not clear if a document has been saved already, as saving in HTML format sometimes gives documents (e.g. different parts of the same Act) the same filename.		
Not clear how to collate more than one document.	P9: Unclear how to select two documents. Tagging functionality not clear.	

6. Try to keep a softcopy or hardcopy record of *part* of a document that is important to you (e.g. print or download only certain parts of a case report).

Usability issues identified by lead developer of the IB usability method	Similar usability issues identified by stakeholders working for a large electronic legal resource development firm	Additional usability issues identified by stakeholders (that were not identified by the lead developer)
Not clear how to save only part of a document. Related to this issue, it is not immediately clear what choosing 'select items' radio button when downloading or printing will do.	P1 (Paired with P7): Unable to complete (cut and paste). P7 (Paired with P1): Task unable to be completed. P9: Cannot print part of a document. Unable to complete task as functionality missing.	
Confusion over whether 'search source' would search within a particular document or not.	P9: Search source function unclear. Functionality unclear.	

Core and Custom Tasks

1. Try to gain access to the electronic resource.

Usability issues identified by lead developer of the IB usability method	Similar usability issues identified by stakeholders working for a large electronic legal resource development firm	Additional usability issues identified by stakeholders (that were not identified by the lead developer)
Participant was not able to log in to the resource directly and found the login facility across various websites that belong to the resource development firm to be confusing.	P4 and P8 pair: Logging in appears to be problematic on account of portal design. P2: Sign on screen should be linked and highly visible from marketing site.	

2. Try to find out which parts of the electronic resource you have access to.

Usability issues identified by lead developer of the IB usability method	Similar usability issues identified by stakeholders working for a large electronic legal resource development firm	Additional usability issues identified by stakeholders (that were not identified by the lead developer)
Participant was not sure how to find out exactly what he has access to within LexisNexis Butterworths and what he does *not* have access to.	P4 and P8 pair: Unclear which sources are available for this subscription. P2: User does not know that application is subscription sensitive. Training issue.	

3. Think of some information that you currently need or have recently needed to find for your work and demonstrate, using the electronic resource, how you might go about finding it.

Usability issues identified by lead developer of the IB usability method	Similar usability issues identified by stakeholders working for a large electronic legal resource development firm	Additional usability issues identified by stakeholders (that were not identified by the lead developer)
Participant was not sure how to expand the Table of Contents side menu bar when this bar is presented.	P2: Left-hand TOC frame – unable to make wider. Visibility of frame grab.	P4 and P8 pair: Keyword search was initially too broad to provide expected results.
Participant used the 'general search' field (rather than segmented fields) for all searches and did not use any Boolean operators. Perhaps the other search possibilities were not made prominent at the interface.	P2: Search syntax formulation does not include any Boolean operators. User inexperience in constructing searches. P2: User enters search terms in free text field, thereby searching entire documents not specific segments. User inexperience in using general search fields above source type search fields.	
Participant assumed that 'old, obscure [cases] are more likely to be on Westlaw now' and suggest that 'perhaps there isn't much on here as there used to be.'		
Confusion over what the Table of Context bar displays during a search for cases. First, participant assumed that it listed other cases that were mentioned in the case report displayed on-screen. Then participant assumed that the TOC was a results list (as his search had returned only 1 document and therefore jumped straight to displaying the document full-text along with the TOC).	P2: User confuses TOC with results list and citator info. Labelling of left-hand pane in results/doc. View. Also, confusion when left-hand bar switches from results groups to TOC items.	
'Back' button is sometimes non-responsive when navigating between pages.		

4. Try to find out which sources contain information about a particular legal area.

Usability issues identified by lead developer of the IB usability method	Similar usability issues identified by stakeholders working for a large electronic legal resource development firm	Additional usability issues identified by stakeholders (that were not identified by the lead developer)
Participant uses the source search 'keyword' field incorrectly by searching for a topic keyword as opposed to a keyword referring to a particular source name.	P2: User enters search terms with no thought as to what source they were searching. Visibility of source selection component.	

5. Try to set up an alert so that you can be informed every time new documents are added to the system that match particular search terms.

Usability issues identified by lead developer of the IB usability method	Similar usability issues identified by stakeholders working for a large electronic legal resource development firm	Additional usability issues identified by stakeholders (that were not identified by the lead developer)
Participant confuses setting up a 'scheduled search' with setting up an 'alert.'	P4 and P8 pair: Slightly slow start – not immediately obvious how to create an alert.	

6. Try to conduct a more advanced search that is restricted to a particular legal area.

Usability issues identified by lead developer of the IB usability method	Similar usability issues identified by stakeholders working for a large electronic legal resource development firm	Additional usability issues identified by stakeholders (that were not identified by the lead developer)
Not immediately clear how to restrict a search to a particular legal area.	P4 and P8 pair: User struggled to focus search on specific area – tried most sub-tabs before using practice area. P2: Unable to see access point for legal taxonomy terms. Taxonomy not used – needs to be more visible.	

7a. Try to follow a hyperlink or other form of connection from a legal case to a previous case or piece of legislation mentioned in the case report and 7b. Find a particular case, then find out which more recent cases have mentioned it (if any).

Usability issues identified by lead developer of the IB usability method	Similar usability issues identified by stakeholders working for a large electronic legal resource development firm	Additional usability issues identified by stakeholders (that were not identified by the lead developer)
Confusion over what clicking on a CaseSearch result will bring back. Participant assumes she is clicking on the full-text of the case.	P2: User accessed citator document when trying to access full-text case (as they did not notice search was across all subscribed cases). Source selection needs to be more visible.	

8. Try to find a particular legal journal article, then found out which more recent articles have mentioned it (if any).

Usability issues identified by lead developer of the IB usability method	Similar usability issues identified by stakeholders working for a large electronic legal resource development firm	Additional usability issues identified by stakeholders (that were not identified by the lead developer)
Unclear how to forwards chain between legal journal articles (i.e. find out which articles have subsequently cited the current article).	P2: Looking for a function to search for articles referring to current document. User inexperience to an extent, but this could be a future enhancement.	
Unclear as to what information clicking on the 'source info' icon will provide. Participant expected to view citator details for the current article, not information about the source itself (i.e. the journal series).		
Initially unclear what 'hits' buttons would allow user to achieve (i.e. to jump between instances of her search terms in the document).		

9. Try to determine what information is provided whilst you are *browsing* that might help you decide what documents might be relevant or which documents to click on and read.

Usability issues identified by lead developer of the IB usability method	Similar usability issues identified by stakeholders working for a large electronic legal resource development firm	Additional usability issues identified by stakeholders (that were not identified by the lead developer)
Confusion over whether it is possible to browse by legal topic (participant assumes it is only possible to browse by individual source).		
Participant does not believe that the information provided when browsing is sufficient for deciding whether a document might be relevant or worth examining in further detail. Participant claims that, when browsing for articles, 'the titles aren't overly helpful.'		
Not immediately clear what different 'view' options will do when viewing document text and when viewing results list.		

Note the P4 and P8 pair did not have time to review tasks 7, 8 and 9 – he finished at the end of task 6. Participant P2 did not have time to review task 9 – they finished at the end of task 8. Participant P10 did not hand in his IB usability output form.

Appendix 5: Think-aloud instruction sheet and behaviour-focused tasks

This instruction sheet and task list was developed based on our user pilot studies. It is intended to be adapted by users of the IB usability method. Users of the method should replace the name of the electronic resource under evaluation (in this case LexisNexis Butterworths) with the name of the resource they are evaluating.

- Tasks 1, 2 and 3 should be presented to think-aloud participants as part of a **core** IB usability evaluation. Participants should spend around 20-30 minutes performing these tasks.
- Tasks 1-7 (apart from sub-tasks marked **Source Level/Search Level/Content Level**) should be presented to participants as part of a **recommended** IB usability evaluation. Participants should spend around an hour performing these tasks.
- For a **custom** IB usability evaluation, tasks 1-7 should be presented to participants as default. Depending on the focus of your usability evaluation you might, however, choose to only present some of these tasks. Also, depending on your focus, you might also choose to present some or all of the custom tasks related to particular behaviours (tasks 8-12) or some of all of the tasks related to less common levels (tasks marked with **Source Level/Search Level/Content Level**).
- It is important to note that the list of tasks presented is **non-exhaustive** as it is possible to identify additional ways that an electronic resource can support a particular behaviour/level. Therefore this task list should be regarded as **customisable** and **extendable.**
- All of the tasks are based on information behaviour performed naturalistically by one or more academic or practicing lawyers in our empirical study are therefore empirically grounded. The functionality survey in appendix 3 served to ensure that no important tasks that are currently supported by electronic legal resources were overlooked.

Key: | Task performed in 1st user pilot | ❶ |
Task performed in 2nd user pilot	❷
Task performed in 3rd user pilot	❸
Task featured in video shown in developer pilot	①
Task not performed/featured in pilots as LexisNexis Butterworths did not appear to support it.	⊗

- Times in minutes shown in italics in brackets next to each task indicate the approximate time the facilitator should allow before encouraging the think-aloud paricipant to move on to the next task.
- Assistance logging in to the electronic resource should be provided by the facilitator if task 1 over-runs by more than a couple of minutes.

LexisNexis Butterworths Think-aloud session: Participant instructions

Overview

In this session, you will be asked to use the LexisNexis Butterworths electronic legal resource to perform a number of information tasks and to think aloud whilst doing so. This is with the aim of highlighting ways that LexisNexis Butterworths can be improved to make it easier to use.

Do not worry if you have not used LexisNexis Butterworths before, have not used this version before, or have not used it to perform similar tasks before, as prior familiarity with the tasks or with LexisNexis Butterworths is not necessary.

Information tasks

1. **Gain access** to LexisNexis Butterworths (i.e. to log in to the electronic resource). *(2 mins.)* ❶❷❸①

2. Find out which parts of LexisNexis Butterworths or sources within LexisNexis Butterworths you have access to. Two example names of sources are the 'Cambridge Legal Journal' and 'Halsbury's Laws of England.' *(3 mins.)* ❷①

3. **Think of some information that you currently need or have recently needed to find for your work** and demonstrate, using LexisNexis Butterworths, how you might go about finding it. *(15-20 mins.)* ❶❷①

   If using a recent example (as opposed to a current need for information), do not try to remember exactly what you did in the past in an attempt to re-enact it during the think-aloud session. Instead, use the example as a springboard for using LexisNexis Butterworths to complete the task(s) any way that you like.

   Note there are two ways of finding information, **searching** and **browsing**. Searching involves formulating a search query, submitting it and reviewing the results, whereas browsing does not involve formulating a query (think of browsing titles alphabetically in a bookshop). If relevant to the information you are looking for, **try to demonstrate both searching and browsing**.

   Example information-seeking tasks that other participants have demonstrated include:

   - *"I want to find out more about the cause of action known as 'deceit' to see whether we can allege that the administrator of a fund has been fraudulent as well as negligent."*
   - *"I want to find out what the sentencing guidelines are for the crime of Grievous Bodily Harm in order to ascertain whether, in a particular case that I am interested in, these guidelines have been followed properly."*
   - *"I want to form an argument about whether it is against the law to keep an item that was found on somebody else's property."*

For tasks 4-7, you will need to think of more specific examples. You **are** permitted to use the same or similar legal areas in more than one example, provided your examples are still based on a **genuine** current or recent need for information. You **are also permitted** to use the same or similar legal areas to any that you used in the previous task, provided you will not be demonstrating exactly the same actions as in the previous task.

4. **Gain an overview of an area** by:
   a. Trying to gain a basic understanding of the law relating to a particular legal area (e.g. Breaches of contract). *(5 mins.)* ❶❷
   b. Trying to found out which sources contain information about a particular legal area. Two example names of sources are the 'Cambridge Legal Journal' and 'Halsbury's Laws of England.' **Source level.** *(3 mins.)* ❸
   c. Trying to gain an appreciation of the importance of a certain legal journal author's role in a particular legal area. *(3 mins.)* ⊗
   d. Trying to locate a legal journal article written by an author who has published *many articles* or *many highly cited* articles. *(3 mins.)* ⊗

5. **Gain a current or historical understanding of the importance of a document** by:
   a. Trying to find out whether a particular case is still good law. *(5 mins.)* ❶❷
   b. Trying to find out what amendments have been made to a particular piece of legislation over a certain time period. *(3 mins.)* ❷
   c. Trying to find out whether a particular piece of legislation is currently in force. *(5 mins.)* ❸

d.   Trying to locate a historical version of a particular piece of legislation (i.e. a previous version that has since been amended). *(3 mins.)* ⊗

6.   **Maintain awareness of developments in an area** by:
   a.   Trying to find out whether there have been any recent developments in a particular legal area (e.g. Discrimination law). *(5 mins.)* ❶❷
   b.   Trying to set up an alert so that you can be informed every time new documents are added to the system that match particular search terms (e.g. when new documents that match the term 'discrimination' are added). *(3 mins.)* ❷
   c.   Trying to set up an alert so that you can be informed every time there are new developments in a particular legal area. *(3 mins.)* ❷❸

   To preserve anonymity in this task, you may use the e-mail address 'test@test.com' when setting up an alert.

7.   **Return to any one of the tasks where you found useful documents** and:
   a.   Determine which sections of a document that you have found are important to you. *(3 mins.)* ❷
   b.   Use the electronic resource you are using to help you electronically 'highlight' or otherwise distinguish content of interest within a document. *(2 mins.)* ⊗
   c.   Keep a softcopy (downloaded or saved) record of a document that you have found. *(1 min.)* ❶❷
   d.   Keep a hardcopy (printed) record of a document that you have found. *(1 min.)* ❶❷
   e.   Download two documents into a single file (i.e. you should end up with one file saved on the computer that includes the text of two separate documents, e.g. two different legal journal articles or two different sections of a particular piece of legislation). *(2 mins.)* ❸
   f.   Keep a softcopy or hardcopy record of *part* of a document that is important to you (e.g. print or download only certain parts of a case report). *(2 mins.)* ❷❸
   g.   Distribute a document that you have found, by e-mail, to a fictitious colleague from within the electronic resource. *(2 mins.)* ❶❷
   h.   Store a document on the server of LexisNexis Butterworths (i.e. save a copy of the document to a personalised area on the electronic resource itself, so you can access it again quicker in future). *(3 mins.)* ⊗
   i.   Find and view a list of the sources that you have used so far in this session. Two example names of sources are the 'Cambridge Legal Journal' and 'Halsbury's Laws of England.' **Source level.** *(2 mins.)* ❸
   j.   Without using your Internet browser's 'save bookmark' or 'add favourite' command, save details of a source that you have used in this session onto LexisNexis Butterworths so you can search or browse within the source quicker in future. Two example names of sources are the 'Cambridge Legal Journal' and 'Halsbury's Laws of England.' **Source level.** *(2 mins.)* ⊗
   k.   View your search history (i.e. an automatic record that the resource keeps of the searches you have submitted). **Search level.** *(2 mins.)* ❸
   l.   Save a particular search on the server of the electronic resource (i.e. save details of the search terms you have used on the e electronic resource itself, so you can conduct the search again in future). **Search level.** *(2 mins.)* ⊗

To preserve anonymity in this task, you may use the e-mail address 'test@test.com' when sending an e-mail.

(Custom searching tasks):

8.   Try to:
    a. Conduct a more advanced search that is restricted to a particular *time period* (e.g. for a piece of legislation, the enactment date). *(2 mins.)* ❸
    b. Conduct a more advanced search that is restricted to a particular *part of the document* (e.g. for legal journal articles - the journal or article title, particular author names, particular keywords, particular legislation or cases cited etc.). *(2 mins.)* ❸
    c. Conduct a more advanced search that is restricted to a particular *legal area* (e.g. the area of Family Law). *(2 mins.)* ❸
    d. Search within a particular *source* for information on a particular legal topic (e.g. search within Halsbury's Laws of England for information on Human Rights Discrimination). *(2 mins.)* ❸
    e. Without using your Internet browser's 'find' command, search within a particular *document* to locate certain content (e.g. search for all occurrences of the term 'injunction' in the 2004 White and White case). **Content Level.** *(2 mins.)* ⊗

(Custom browsing and extracting tasks):

9.   Try to:
    a. Browse from within a particular document to another document on a similar topic (e.g. browse from a case report to another case about a similar topic or from a commentary article to piece of legislation related to the topic). *(2 mins.)* ⊗
    b. Browse from within a legal journal article to other articles written by the same author. *(2 mins.)* ⊗
    c. Browse to see whether a particular *source* is available in LexisNexis Butterworths (e.g. by browsing through a list of sources that the resource contains). Two example names of sources are the 'Cambridge Legal Journal' and 'Halsbury's Laws of England.' **Source level.** *(3 mins.)* ❸
    d. Conduct a search for particular search terms, and then find occurrences of your search terms in the text of one of the documents that was returned in your search. **Content level.** *(2 mins.)* ❸
    e. Search for a particular document, then use features provided by LexisNexis Butterworths to help you 'jump between' occurrences of your search terms in the text. Alternatively, you can try to jump between section headings in the document or between particular words or phrases in the text. **Content level.** *(2 mins.)* ❸

(Custom chaining tasks):

10.   Try to follow a hyperlink or other form of connection from:
    a. A legal case to a previous case or a particular piece of legislation mentioned in the case report. *(2 mins.)* ❸
    b. A piece of legislation to other pieces of legislation mentioned in the text of the Act. *(3 mins.)* ❸
    c. A legal journal or commentary article to a case or piece of legislation mentioned in the article. *(3 mins.)* ❸

11.   Try to follow a hyperlink or other form of connection from a document to other documents that have been written *since* this document and have subsequently mentioned it. Specifically, try to:
    a. Find a particular case, then find out which *more recent* cases have mentioned it (if any). *(3 mins.)* ❸
    b. Find a particular piece of legislation, then find out which pieces of legislation (if any) that were drafted *after* this Act have mentioned it. *(3 mins.)* ❸
    c. Find a particular legal journal article, then find out which articles that were written *after* this article have mentioned it (if any). *(3 mins.)* ❸

(Custom selecting, distinguishing and filtering tasks):

12. Try to:
    a. Determine what information is provided in your search results that might help you decide which documents might be relevant, or to decide which documents to click on and read. You are free to search for any type of legal document (e.g. a case report, piece of legislation, commentary or journal article etc.). *(3 mins.)* ❸
    b. Determine the different ways that it is possible to sort your search results (e.g. by date, title, relevance etc). You are free to search for any type of legal document (e.g. a case report, piece of legislation, commentary or journal article etc.). *(2 mins.)* ❸
    c. Conduct a search that brings back a number of different types of legal document (e.g. case reports, legislation, journal articles), then *filter* your search results by selecting or excluding particular types of documents after your search results have been returned. *(3 mins.)* ❸
    d. Search *within* your search results (i.e. conduct a search, retrieve some results and then try to conduct another search within those results). *(3 mins.)* ❸
    e. Determine what information is provided whilst you are *browsing* that might help you decide which documents might be relevant or which documents to click on and read. You are free to search for any type of legal document (e.g. a case report, piece of legislation, commentary or journal article etc.). *(3 mins.)* ❸

You will be asked to skip any parts of tasks 3 onwards that you have already demonstrated.

Further instructions to think-aloud participants

Instructions for thinking aloud whilst using LexisNexis Butterworths

- Whilst using LexisNexis Butterworths, you should **think aloud**, mentioning exactly what you are doing while you are doing it. Try to mention everything that is going through your head.
- The facilitator will not ask you any questions, although you might be prompted to explain what you are currently doing if you have not spoken aloud for a while. The facilitator is not permitted to assist you in any way or answer questions about a task after you have started it.
- As this exercise is aimed at making LexisNexis Butterworths easier to use, make sure to mention anything that you find **easy/clear** or **difficult/unclear** when using the resource.
- You are **encouraged** to explore LexisNexis Butterworths during each task and use any of the features within it that you think might help you perform the task. However, you should always maintain a task focus – the aim of this session is not to provide a 'tour' of LexisNexis Butterworths.
- To make the most of your session, you might be asked to **move on** to the next task if time is running short.

Important things to remember

- You are requested to **only use LexisNexis Butterworths** during the session and no other electronic resources/pieces of software/websites. The facilitator will direct you back towards LexisNexis Butterworths if you load another resource/piece of software/website.
- This is **not a test** of your performance or how you use LexisNexis Butterworths and there is no 'right' or 'wrong' way of performing the tasks.
- The facilitator will be **unable to answer any questions** while you are carrying out the tasks, so please ask any questions before you begin.
- You are encouraged to **perform each task in a way that is natural to you** (i.e. a way that you might normally attempt this or a similar task when using an electronic legal resource).

Please ask if you are unsure about any of these instructions or have any questions at all, then read and sign the informed consent form.

Appendix 6: Summary IB functionality evaluation form

Name/version of electronic resource:	Date of functionality evaluation:	Members of team present:

If evaluation is restricted to certain parts of the resource, list which parts:
If evaluation is restricted to a particular set or sets of behaviours, select which set(s):   Core behaviours / core + law-specific behaviours / core + information use behaviours

Step 1: Decide which of the behaviours below are supported by the electronic resource under evaluation.

	Resource level	Source level	Document level	Content level	Search query/result level
Accessing	Currently supported? (Y/N)	Currently supported? (Y/N)	Currently supported? (Y/N)		
Surveying		Currently supported? (Y/N)	Currently supported? (Y/N)		
Monitoring		Currently supported? (Y/N)	Currently supported? (Y/N)		
Searching		Currently supported? (Y/N)	Currently supported? (Y/N)	Currently supported? (Y/N)	
Browsing and extracting		Currently supported? (Y/N)	Currently supported? (Y/N)	Currently supported? (Y/N)	
Chaining		Currently supported? (Y/N)	Currently supported? (Y/N)		
Selecting, distinguishing and filtering		Currently supported? (Y/N)	Currently supported? (Y/N)		
Updating			Currently supported? (Y/N)		
History tracking			Currently supported? (Y/N)		
Analysing				Currently supported? (Y/N)	
Recording		Currently supported? (Y/N)	Currently supported? (Y/N)	Currently supported? (Y/N)	Currently supported? (Y/N)
Collating			Currently supported? (Y/N)	Currently supported? (Y/N)	
Editing			Currently supported? (Y/N)		
Distributing			Currently supported? (Y/N)	Currently supported? (Y/N)	Currently supported? (Y/N)

Key: Core information-seeking behaviours ☐   Law-specific information-seeking behaviours ▨   Information use behaviours ▨

Step 2: Fill out a detailed functionality form for each of the behaviours/levels above (i.e. for each cell in the grid where you have not excluded that particular behaviour/level).

Step 3: Answer the following two questions*:   **These questions were added as a result of the formative evaluation of the method.*

a) Are there any *behaviours/levels* that it may no longer be necessary to support? For any behaviours/levels which you are considering ceasing support for, what are the potential arguments for and against support?

b) Are there any *ways that you currently support any of the behaviours/levels* that may no longer be necessary? For ways of supporting a particular behaviour/level which you are considering ceasing support for, what are the potential arguments for and against support?

Appendix 7: Detailed IB functionality evaluation form

One of these forms should be filled out for each behaviour and level that you have chosen to evaluate the electronic resource in relation to.

Name/version of electronic resource:	Date of functionality evaluation:	Page __ of __

Step 1: Tick **one** box in the grid to indicate the behaviour and level you will be answering the questions below about:

	Resource level	Source level	Document level	Content level	Search query/ result level
Accessing					
Surveying					
Monitoring					
Searching					
Browsing and extracting					
Chaining					
Selecting, distinguishing and filtering					
Updating					
History tracking					
Analysing					
Recording					
Collating					
Editing					
Distributing					

Step 2: Answer the relevant questions below about the behaviour and level that you have ticked:

If the behaviour you have ticked **is** currently supported at this level by the resource, answer the following two questions:

1) **In which way(s)** does the resource currently support this behaviour at this level?

2) **In which additional way(s)** might the resource support the behaviour at this level? You do not need to answer this question if you **do not have a direct stake in the resource** (e.g. you have no input into the design, improvement or marketing of the resource or it is a competitor resource)

If the behaviour you have ticked **is not** currently supported at this level by the resource, in which way(s) might the resource support the behaviour at this level?

Additional notes from discussion of issues surrounding supporting the behaviour(s) at this level:

For guidance answering questions (a), (b) or (c), see illustrative ways of supporting this behaviour and level in appendix 3. Prior to the formative evaluation of the method, this form also required evaluators to consider the 'potential arguments for and against supporting the behaviour(s) at this level.' This question was replaced by the two questions on the summary functionality form.

Appendix 8: Form used to record the output of an IB usability evaluation (and detailing usability data extracted from the think-aloud transcript in appendix 9).

Name/version of resource: *LexisNexis Butterworths (Jul'07)*	Type of evaluation: Core/Recommended/Custom *(N/A)* Date of user evaluation: *23/7/07* Page 1 of *2*
Video filename(s): *Pilot1.avi*	Anonymous participant identifier: *Pilot 1* Facilitator/evaluator name: *Stephann Makri*

If evaluation is restricted to certain parts of the resource, list which parts: *Evaluation restricted to default screens available to academic users in the UK.*

If evaluation is restricted to a particular set or sets of behaviours, select which set(s):   Core behaviours / core + law-specific behaviours / core + information use behaviours   *(N/A)*

On the following scale, **how experienced** does the think-aloud user consider themselves to be with using:

	Not at all experienced	Somewhat experienced	Very experienced
General electronic resources in their domain (e.g. Law)?	☐	☑	☐
The electronic resource currently under evaluation?	☑	☐	☐

Fill in these four columns whilst watching the think-aloud session   Fill in these columns either whilst watching or after watching the session.   Fill in this column after watching the session

Task n°	*User actions or comments, or **personal observations** that might suggest a usability issue*	*Approx. time in video actions/ comments/ observations occurred*	***Screen(s)/ page(s) / part(s)** of resource that actions/ comments/ observations relate to*	*Usability issue(s) identified from actions/ comments/ observations. Do not repeat issues identified as a result of considering several sets of actions/ comments/ observations –list them once.*	*Severity of issue (Not severe - does not need immediate attention, **Q**uite severe – needs attention in the future, **V**ery severe – needs immediate attention)*	*Amount of effort required to address issue (**S**mall/ **M**edium/ **L**arge)*	*Reflections on usability issues identified (e.g. broad themes arising from issues identified, issues for future discussion etc.)*
N/A	a, g) User about to type/types search query into incorrect segmented field	1 min 0s, 11min 10s	'Cases' search page	Not immediately clear which segmented field to enter data in	N	L	As a novice user, the trainee solicitor was able to perform basic searches, but it was not always clear how to achieve particular goals using the resource, or what effect performing certain interface actions might have.

The help and tutorial screens frustrated the user rather than aiding the updating task. |
| | b) User initially enters incorrect search syntax | 1 min 10s | 'Cases' search tab and all other search pages | Required syntax not immediately clear | N | S | |
| | c) User clicks on 'next steps' combo box hoping to check whether case has been updated | 4 min 50s | All full-text document pages | Actions facilitated by 'next steps' combo not transparent or clear to user | Q | M | |

Task n°	User actions or comments, or personal observations that might suggest a usability issue	Approx. time in video actions/ comments/ observations occurred	Screen(s)/ page(s) / part(s) of resource that actions/ comments/ observations relate to	Usability issue(s) identified from actions/ comments/ observations. Do not repeat issues identified as a result of considering several sets of actions/ comments/ observations –list them once.	Severity of issue (Not severe - does not need immediate attention, Quite severe – needs attention in the future, Very severe – needs immediate attention)	Amount of effort required to address issue (Small/ Medium/ Large)	Reflections on usability issues identified (e.g. broad themes arising from issues identified, issues for future discussion etc.)
N/A	d) Help pages 'seem to be frustrating' user	9 min 0s	TotalHelp pages	Help pages did not provide the required help	Q	L	
	e) 'Slow and cumbersome' tutorial	10 min 20s	Tutorial pages	Tutorial pages did not help	Q	M	
	f) User unsure which 'submit' button to use	10 min 40s	'Home' search page	Relation of submit buttons to search fields potentially unclear	V	M	
	h) User unaware that search was processing even though 'go' button turned grey upon submit	10 min 40s	'Home' search page	Unclear feedback provided relating to search progress	Q	M	
	i) User unsure of what happened when selecting case name from 'view' combo box	12 min 50s	All full-text document pages	Effect of selecting current case from 'view' combo potentially unclear	N	M	

Prior to the formative evaluation of the method, the 4th column ('screen(s)/pages(s)/parts of the resource…') was presented on the right-hand-side of the 'usability issue(s) identified…' column.
The order was changed based on feedback from the evaluation.

Appendix 9: Think-aloud transcript of a Trainee Solicitor using LexisNexis Butterworths to 'find out whether a particular case is still good law'

The transcript below is a verbatim account of the comments and actions made by a Trainee Solicitor who was asked to use the LexisNexis Butterworths electronic legal resource to 'find whether a particular case is still good law.' The Trainee Solicitor was asked to think aloud whilst performing the task, verbalising his thoughts, actions and feelings.

The Trainee's actions are presented in square brackets [ ]. The highlighted parts of the transcript denote sections that are referred to, by the letter alongside them, on the form presented in appendix 8.

a b	There's one case called White and White, which is used in ancillary relief cases and I want to find out about that. I want to find out if that case is still valid as good law, you know, or whether any amendments have been made to the case over a period of time. So what I'm gonna do is I'm back at the LexisNexis home screen. I'm gonna click on 'cases' and under 'enter search terms' I'm going to put [pauses]. No, I'm going to put it in the 'case name.' so White and [Participant types in 'White v W,' then pauses]. [Reads caption underneath 'case name' field]. 'To find Smith v Jones, enter Smith *and* Jones.' So I'm going to use the word 'and' rather than 'v.' I don't have a citation unfortunately and I'm going to click 'search' and see what that brings back. 1700 cases, which may make it a bit difficult to find. [Clicks on 'cases' results group]. So I've clicked on 'cases' to expand it [scrolls hrough list of cases]. It's going to be quite difficult to find. I'm just going to look on the left-hand side at where the case has been reported [scrolls through side-bar detailing result groups that comprise mostly report series and levels of court]. This is a family case, so I'm just going to have a quick browse. If not, I'll have to go back and narrow the search down slightly. Aah, 'Family Court Reports' [reads report series aloud and clicks on hyperlink]. I'm hoping that that *may* bring up the case. Right, White and White 2000 and it's a House of Lords case. I think that may be it. So I'll click on that. Right, [reads case name aloud] White and White. [Begins to read keywords aloud]. 'divorce,' 'financial provision, ancillary relief,' 'available assets exceeding parties' financial needs for housing and income.' Aah, there we go, this is the case. 'Principles to be applied.' [Reads first paragraph of full-text of case]. 'The parties' marriage broke down after 30 years.' 'The overall net worth of their assets was approximately £4.6 million.' It's a big money case. Ok, this is the case. [Scrolls through text]. Now I just want to see if it's been updated in any way. Hmm. I've just scrolled to the bottom. That's the law lords' dictum. [Scrolls back up through text]. Ok. I'm looking at the cases referred to in the opinions. I can't see anything there. I'm just wondering how I can check to see if the case has been updated or not. [Clicks on 'view' combo box]. I'm clicking on 'view,' [reads options aloud]. 'Expanded list,' 'list.' These are presumably lists. Next steps [clicks on 'next steps' combo box and reads some of the options aloud]. 'Edit search.' Here we go. 'Save search,' 'create alert,' 'search the source,' 'find
c	related cases.' Hmm. I'll click on 'find related cases' and see what that brings up. [Scrolls through results]. And this has brought up 48 further cases. I'm just going to click on the top one, which is dated the 18th June 2007 [clicks on Wood vs. Rost case] which is called Wood and Rost [reads keywords of case aloud]. None of that's got anything to do with the split. [Scrolls through text of case]. Aah. Here we go. Paragraph 7. [Reads out text in paragraph 7 of the case report]. 'Mr Mostyn, fortified by the decision of the House of Lords in White v White.' And it's got the references. And it's got one of the references in red. I'm going to click on that in a minute. [Continues to read paragraph 7]. Let's click on the red link. Presumably that will take me back to what I was looking at before. Yeah, it has done.
d	[Scrolls through text of White and White case again]. As to see whether or not it is still good law, I think I'm a bit stuck. So I think what I'm going to have to do is press 'help.' [Reads title in TotalHelp screen]. I'll click on 'home' because that doesn't really help me. [Hovers cursor over headings, then clicks on 'tutorials' hyperlink]. Tutorials. Well I know how to search. Home. [Clicks on 'home' hyperlink]. It's taken me back to where I was. I'll click on 'search for information.' Nope. [Clicks on 'locate sources' heading]. Nope. [Clicks on 'review search results' heading and reads out some of the sub-headings]. 'How do I work with my search results?' I'll click on that as it might help me in some way. [Reads out sub-headings]. No, that doesn't help me. Let's search. [Searches within TotalHelp for 'current cases' and clicks on a result entitled 'results page' then soon after closes the help page]. I've closed that down because it just seems to be frustrating me.

e f/g h i	[Participant returns to White and White case and scrolls through text again]. What I ***am*** going to do is go back to the Lexis homepage as there are often tutorials on the right-hand-side. [Reads headings in 'related links' side box]. 'View tutorials,' 'overview,' 'selecting sources,' 'working with results.' [Reads other headings aloud, then scrolls down and back up the LexisNexis Butterworths homepage]. No. [Clicks on 'overview' under 'view tutorials' heading]. I'll click on that. Maybe it will bring up something different. It has linked to a tutorial. [Click on 'working with results' tutorial option]. Working with results. I'm waiting for it to start. [Clicks mouse two or three times]. Ok, the tutorial seems very slow and cumbersome, so [closes tutorial window], I'm going to try searching again. [Enters 'white and white' into the 'get a specific document' part of the LexisNexis Butterworths homepage with 'find a case' selected in the combo box. Participant then presses the grey 'go' button next to the field he has filled out, but after a few seconds presses the main red 'search' button. [An information box is displayed, which says 'Please enter a search term']. [Participant then clicks on the 'cases' tab on the LexisNexis Butterworths homepage]. [Participant types 'White' in 'citation' field]. I'm not looking for a citation. [Participant then enters 'White and White' in the case name field]. I know it's a House of Lords case [ticks 'House of Lords' in the 'Court' selection box] so that will narrow everything down a bit [conducts search]. Cases [clicks on 'cases' results group]. Family Court Reports [clicks on Family Court Reports sub-group and selects the top case from the list, White and White]. Actually, I'm just going to go back [clicks browser 'back' button and ticks box beside the White and White case that participant had just previously clicked on]. [Clicks on 'sort' combo box and 'view' combo box, then clicks on 'view tagged' button. View tagged. I'm assuming that will be the same as if I just clicked on it. Yeah. I don't think I can [pauses]. [Clicks on 'next steps' combo box and reads some of the options aloud again]. 'Find related cases,' 'find related commentary' [reads other options]. View [clicks on 'view' combo box again and reads out options]. 'List,' 'expanded list,' 'White v White 2000.' [Selects White v White 2000]. Oops, I clicked on something. [Scrolls through text of White and White 2000 case]. Unfortunately I think I'm stuck and can't go any further on this one.

Appendix 10: Example informed consent form (used during our user pilot studies)

Dear participant,

I am conducting a study looking at how easy to use electronic legal resources are. The aim of this study is to improve the design of electronic resources.

By signing below, you signify that you:

- Agree to take part in this study.
- Understand what the study involves (please ask if you have any remaining questions).
- Give permission for the session to be audio and screen recorded.
- Are aware that any details that could be used to identify you arising from this study will be anonymised.
- Are aware that the audio and screen recordings will be disseminated and stored in accordance with the Data Protection Act 1998. In particular, you are aware that the anonymised audio and screen recordings will be shown to third-parties strictly for the purposes of use in academic presentations and tutorials. These third-parties will use the recordings to try to improve the usability of the electronic resource that you have been using.
- Are aware that you will be able to review, amend or delete any data you provide for the study at any time, and without penalty. This means that you can withdraw from the study at any time and without any risk of penalty.

Signed:

Print:

Date:

Please make a note of the following contact details for if you have any queries or questions after the study or if you would like to withdraw from the study:

Stephann Makri
University College London Interaction Centre, 31-32 Alfred Place, London WC1E 7DP.
Phone: +44 (0)20 7679 5242
E-Mail: s.makri@ucl.ac.uk

Appendix 11: Focus group questions examining the usefulness, usability, learnability and likelihood of future use of the IB <u>functionality</u> method

Usefulness-related questions

- How useful did you find the functionality method to be for helping you to determine what functionality your own resource should have? Why?
- What previous approaches have you used to evaluate the functionality of your own resources and how does this method compare to these approaches?
- In your opinion, how useful is the functionality method for helping you identify opportunities for *increasing* the functionality provided by your own resources? Why?
- In your opinion, how useful is the functionality method for helping you identify opportunities for *reducing* the functionality provided by your own resources? Why?

Usability-related questions

- How easy was it, overall, to use the functionality method to evaluate the functionality of *your own* resource? Why?

Of the following usability-related questions, only ask those that have not already been addressed by participants:

- How easy was it to decide whether *behaviours/levels* are currently *supported* by your own resource ? Why?
- How easy was it to identify ways that your own resource *currently supports* a particular behaviour at a certain level? Why?
- How easy was it to identify *ways or additional ways* that your own resource *might support* a particular behaviour at a certain level? Why?
- How easy was it to identify arguments for and against supporting behaviours at a particular level. Why?

Learnability-related questions

- How easy was it to learn the functionality method? Why?
- What were the positive aspects of the *learning experience* with regard to learning the functionality method?
- What could have improved the *learning experience* with regard to learning the functionality method?

Future use-related questions

- How likely is it that you will use the functionality method in the future? Why?
- If you decide to use the functionality method in the future, are you likely to use it to evaluate the functionality of own resources/products, competitor resources/products or both? Why?
- If you decide to use the functionality method in the future, are you likely to make any changes to it? If yes, what changes are you likely to make and why?

General wrap-up questions

- What do you believe are the greatest *benefits and drawbacks* of using the functionality method and why?
- What improvements, apart from those we have already discussed, would you make to the functionality method?

Appendix 12: Focus group questions examining the usefulness, usability, learnability and likelihood of future use of the IB <u>usability</u> method

Usefulness-related questions

- How useful is the usability method in helping you to highlight usability issues that might inform design/re-design? Why?
- What previous approaches have you used to evaluate the usability of your products and how does this method compare to these approaches?
- How balanced is the usability method in highlighting severe vs. not so severe usability issues and in method in highlighting usability issues that are easy to address vs. those that are not so easy to address? Why do you think this is the case?
- Did you find filling in all of the columns of the usability form to be useful? If no, which columns were not useful and why? Would you want to record any additional information on the form?

Usability-related questions

- How easy was it, overall, to use the usability method to conduct an analysis of the video data? What was particularly easy/difficult about it?

Of the following usability-related questions, only ask those that have not already been addressed by participants:

- How easy was it to identify user actions or comments that might suggest a usability issue from the video data? What was particularly easy/difficult about it?
- How easy was it to make own observations from the video data that might suggest a usability issue? What was particularly easy/difficult about it?
- How easy was it to identify usability issues from the video data? What was particularly easy/difficult?
- How easy was it to associate user actions/comments/own observations with an underlying usability issue? What was particularly easy/difficult about it?
- How easy was it to identify the screen(s)/page(s)/parts of the resource that a usability issue relates to? What was particularly easy/difficult about it?
- How easy was it to determine the severity of usability issues? What was particularly easy/difficult about it?
- How easy was it to determine the amount of effort required to address usability issues? What was particularly easy/difficult about it?
- How easy was it to make reflections on the usability issues identified after you had finished watching the think-aloud session? What was particularly easy/difficult about it?

Learnability-related questions

- How easy was it to learn the usability method? Why?
- What were the positive aspects of the *learning experience* with regard to learning the usability method?
- What could have improved the *learning experience* with regard to learning the usability method?

Future use-related questions

- How likely is it that you will use the usability method in the future? Why?
- If you decide to use the IB usability method in the future, are you likely to make any changes to it (i.e. apart from customising it)? If yes, what changes are you likely to make and why?

General wrap-up questions

- What do you believe are the greatest *benefits and drawbacks* of using the usability method and why?
- What improvements, apart from those we have already discussed, would you make to the usability method?