



City Research Online

City, University of London Institutional Repository

Citation: Feduzi, A., Faulkner, P., Runde, J., Cabantous, L. & Loch, C. (2022). Heuristic methods for updating small world representations in strategic situations of Knightian uncertainty. *Academy of Management Review*, 47(3), pp. 402-424. doi: 10.5465/amr.2018.0235

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/25201/>

Link to published version: <https://doi.org/10.5465/amr.2018.0235>

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

City Research Online:

<http://openaccess.city.ac.uk/>

publications@city.ac.uk

**HEURISTIC METHODS FOR UPDATING SMALL WORLD REPRESENTATIONS IN
STRATEGIC SITUATIONS OF KNIGHTIAN UNCERTAINTY**

Alberto Feduzi
University of Cambridge
a.feduzi@jbs.cam.ac.uk

Philip Faulkner
University of Cambridge
pbf1000@cam.ac.uk

Jochen Runde
University of Cambridge
j.runde@jbs.cam.ac.uk

Laure Cabantous
City, University of London
laure.cabantous.1@city.ac.uk

Christoph Loch
University of Cambridge
c.loch@jbs.cam.ac.uk

This paper has been accepted for publication in the *Academy of Management Review*.
<https://doi.org/10.5465/amr.2018.0235>

Acknowledgements

We are grateful to our action editor Sharon Alvarez for her guidance and two anonymous referees for their valuable comments. For discussion and comments on earlier versions of this paper, we thank Joep Cornelissen, Saïd Elbanna, Florian Ellsäßer, Ronald Klingebiel, Russ McBride, Olga Petricevic, Joe Porac, Stratos Ramoglou, Claus Rerup, Mark Zbaracki, Charlene Zietsma, and the participants of the 2018 conference on *The Implications of Uncertainty for Management and Organization Theories* at the University of Pittsburgh.

Abstract

Recent studies on the construction and use of “small world representations” in strategic decision-making under Knightian uncertainty say little about how such representations should be updated over the implementation phase. This paper draws on the psychology of reasoning to take a step towards answering this question. We begin by theorizing small world representations and how the scenario spaces they contain are constructed and may be updated over time. We then introduce two well-known heuristic methods of inquiry, disconfirmation and counterfactual reasoning, translate them into practical procedures for updating scenario spaces, and compare the relative performance of these procedures in strategic situations of Knightian uncertainty. Our principal findings are that the procedure based on counterfactual reasoning is superior to the one based on disconfirmation with respect to (1) counteracting the confirmation bias, (2) promoting the exploration of the set of imaginable scenarios, and (3) facilitating action to mitigate or exploit the consequences of what would otherwise have been Black Swans. We close with some broader implications for the study of strategic decision-making under Knightian uncertainty.

It is widely recognized that decision-makers operating in situations of *Knightian uncertainty* (Knight, 1921; Alvarez & Barney, 2005, 2007; McMullen, & Shepherd, 2006; Reymen et al., 2015; Packard & Clark, 2020) employ simplified cognitive representations, sometimes called “small world representations” (SWRs), to anticipate the longer-term consequences of their strategic decisions (Gavetti, Greve, Levinthal, & Ocasio, 2012; Grégoire, Barr, & Shepherd, 2010; Levinthal, 2011, 2018; Maitland & Sammartino, 2015a,b; Packard, Clark, & Klein, 2017; Porac, Thomas, & Baden-Fuller, 1989). There has been a good deal of research focusing on the initial *construction* of such representations. Much of this work investigates the factors that affect decision-makers’ ability to do so effectively, such as their experience, expertise and cognitive processes (Gavetti & Levinthal, 2000; Gary, Wood, & Pillinger, 2012; Grégoire et al., 2010; Lovallo, Clarke, & Camerer, 2012; Maitland & Sammartino, 2015a; Simon, 1947) and the quality and quantity of the information available to them (Gavetti & Rivkin, 2007; Maitland & Sammartino, 2015b).

The question of how SWRs should be *updated* during the implementation phase of strategic decisions has however received considerably less attention, perhaps because adaptive responses to emerging information are often seen as extempore and *ad hoc* rather than the result of systematic analysis (Klingebiel & De Meyer, 2013; Langley, Mintzberg, Pitcher, Posada, & Saintmacary, 1995). The question is nonetheless a pressing one in situations of Knightian uncertainty, where the scope for constructing adequate SWRs *ab initio* is often highly limited, and where decision-makers relying on heuristics to guide their actions (Bingham & Eisenhardt, 2011; Busenitz & Barney, 1997; Maitland & Sammartino, 2015a; Shepherd, Haynie & McMullen, 2012)¹ often find it difficult to adapt their cognitive representations in a timely way (Tripsas & Gavetti, 2000).

¹ Heuristic methods are understood differently in different disciplines (Hertwig & Pachur, 2015). In the present paper, we take a pluralistic, Simonian, view of heuristics as any simple but useful

The purpose of this paper is to make a start on answering this question. To this end, we draw on the psychology of reasoning (e.g., Evans, 2016; Manktelow, 2012; Poletiek, 2001) to theorize how two well-known heuristic methods of inquiry—disconfirmation and counterfactual reasoning—might inform procedures for updating SWRs. We show that, in strategic situations of Knightian uncertainty, the procedure based on counterfactual reasoning is likely to be more effective in supporting the process of updating SWRs than the one based on disconfirmation with respect to (1) counteracting the confirmation bias, (2) promoting the exploration of the set of imaginable scenarios, and (3) facilitating action to mitigate or exploit the consequences of what would otherwise have been Black Swans (Taleb, 2007).

Our argument begins with a detailed exposition of SWRs, focusing particularly on how the associated scenario spaces are constructed and may change as they are updated over time. We then provide a brief account of disconfirmation and counterfactual reasoning as they appear in the psychology literature, translate them into practical procedures for updating SWRs, and compare the relative performance of these procedures in strategic situations of Knightian uncertainty. We conclude with a discussion of the wider implications of our work.

SMALL WORLD REPRESENTATIONS AND KNIGHTIAN UNCERTAINTY

The concept of small world representations can be traced to the “small world” idea introduced by Leonard Savage in his landmark *Foundations of Statistics* (Savage, 1954). While the underlying technicalities are considerable and continue to be debated to this day (Shafer, 1986; Binmore, 2008; Bradley, 2017), a small world is most easily understood as a decision situation that has been

method people might deploy, consciously or unconsciously, in the acquisition and processing of information, problem solving and decision making, in situations in which optimal methods are unavailable or impractical.

distilled to a choice of one from a set of courses of action or “acts”, each of which lead to one of a set of “consequences” depending on which of a set of contingent and exogenously determined “states of the world” obtains. Consequences are regarded as “anything at all about which the person could possibly be concerned” (Savage, 1954: 14) and states as descriptions of the world that leave “no relevant aspect undescribed” (Savage, 1954: 9). The decision-maker is assumed to know the sets of acts, states and consequences, and the only uncertainty concerns which of the set of states will turn out to be the true one, what is sometimes called *state uncertainty* (Bradley & Drechsler, 2014).

It has often been pointed out, even by Savage (1954: 16) himself, that many if not most practical decision situations are too complicated to be reduced to small worlds by the decision-maker (Binmore, 2008, 2009; Bradley & Drechsler, 2014; Gigerenzer & Gaissmaier, 2011). That is to say, in many real-world situations it is simply beyond human ability to arrive at exhaustive lists of the relevant acts, states and consequences, never mind describe the members of these lists with any high degree of completeness or precision (Gilboa & Schmeidler, 1995, 2001; Levinthal, 2011). Some authors have accordingly written about the problem of decision-making in “Large” (Binmore, 2008; Drechsler, 2012) or “Grand” (Bradley, 2017) worlds, by which they mean situations in which small world formulations are not feasible. Since strategic decisions typically involve Knightian uncertainty, they will invariably be of this kind (Levinthal, 2011).

However, none of this should be allowed to obscure that the broad acts/states/consequences framework provides a natural and quite common way to structure decision-making problems, and in particular that strategic decision makers often do seem to employ highly simplified cognitive representations that *approximate* small worlds to guide their actions (Gavetti et al., 2012; Csaszar & Levinthal, 2016). We will follow recent contributions to the strategy and entrepreneurship literature in calling such approximations small world representations (SWRs) (Levinthal, 2011,

2018; Maitland & Sammartino, 2015a; Packard et al., 2017), which we interpret as rough analogues of Savage small worlds insofar as they also share the same broad acts/states/consequences structure.

For our purposes, there are two key differences between small worlds and SWRs. The first is that the descriptions of SWR analogues of acts, states and consequences tend to be partial and vague by comparison to what is assumed in small worlds, and where, importantly, the membership of these three categories is typically incomplete (Alvarez & Barney, 2005; Gavetti & Menon, 2016; Levinthal, 2011; Maitland & Sammartino, 2015a; Reymen et al., 2015; Packard et al., 2017). That is to say, whereas the uncertainty facing the decision maker in a small world is limited to state uncertainty, SWRs typically reflect the many additional uncertainties characteristic of many real-world decisions, including uncertainties about the number and nature of possible acts, states and consequences, as well as about the probabilities of states. We will henceforth use the term *Knightian uncertainty* to refer to situations in which decision-makers face these additional types of uncertainties to significant degrees.²

² This definition of Knightian uncertainty encapsulates the various ways in which Knight's (1921: 20) original notion of "unmeasurable uncertainty" has typically been used in decision theory, economics and business studies, where it refers, for instance, to the absence of knowledge of sharp numerical probabilities that might be assigned to each of a list of exhaustive and mutually exclusive states of the world (Ellsberg, 1961; Bewley, 1986; Runde, 1998) or, in some cases, to the impossibility of determining both the possible outcomes of a decision and their associated probabilities (Alvarez & Barney, 2007; Townsend et al., 2018). However, it goes further in extending the concept to include the uncertainty stemming from possible omissions from the *state space* (Bradley & Drechsler, 2014, Dekel, Lipman, & Rustichini, 1998; Mukerji, 1997) as distinct

The second key difference between small worlds and SWRs concerns the treatment of time, and thus the scope for learning within each formulation. In a small world, the specification of acts, states and consequences entails that everything that might happen at all subsequent stages of a decision problem is laid out from the outset. The result is a highly attenuated rendering of time in which “genuine” learning—such as uncovering new, formerly unimagined, possibilities—is precluded (Binmore, 2008; Feduzi & Runde, 2014). SWRs impose no such requirement on their analogue acts, states and consequences, and are therefore consistent with the ongoing, contingent and adaptive nature of strategic decisions in which it is rarely possible to know everything relevant about what the future might hold from the outset (Alvarez et al., 2013; Packard et al., 2017). SWRs therefore offer room for the acquisition of new information over time and thus the possibility of updating the SWR over the course of the implementation of a strategic decision.

Decision-makers’ ability to construct SWRs in a way that provides a useful guide to action (Levinthal, 2011; Csazar & Levinthal, 2016) depends on both the cognitive attributes of the decision-maker and the nature of the informational environment (Gary et al., 2012; Gavetti & Rivkin, 2007; Maitland & Sammartino, 2015a,b). Foremost among the former are capacities and dispositions such as imagination, creativity and curiosity, that are crucial to envisioning possible futures, as well as an actor’s prior knowledge, experience and expertise, and the reasoning and problem-solving methods used (associative thinking, analogical reasoning, fast-and-frugal thinking, etc.) (Alvarez & Porac, 2020; Arikan, Arikan, & Koparan, 2020; Brandenburger, 2017; Felin & Zenger, 2017; Gavetti & Levinthal, 2000; Grégoire et al. 2010; Lovallo et al., 2012;

from that associated with the absence of numerically definite probabilities or the possibility of unforeseen *consequences*.

Maitland & Sammartino, 2015a; Pontikes & Rindova, 2020; Porac & Tschang 2013; Rindova & Courtney, 2020; Simon, 1947).

The quality and quantity of the information available to the decision-maker is also of paramount importance to the construction of effective SWRs (Gavetti & Rivkin, 2007; Maitland & Sammartino, 2015b). At the very early stages of an entrepreneurial endeavour, after a major shock or, more generally, in rapidly changing environments, the information accessible to strategic decision-makers will often be highly incomplete, non-definitive and ambiguous (Alvarez & Barney, 2007; Forbes, 2007; Gavetti & Rivkin, 2007; Maitland & Sammartino, 2015b; Miller, 2008; Shepherd et al., 2015). Moreover, the limited information that is available is liable to change constantly, not least because actors' responses to changes in their environment may lead to it changing still further (Alvarez & Porac, 2020; Berglund, Bousfiha, & Mansoori, 2020; Gavetti et al., 2017; Packard & Clark, 2020; Patvardhan & Ramachandran, 2020).

Thus far, however, the literature has had little to say about issues concerning the updating of SWRs, in particular the prescriptive question of how decision-makers should update their SWRs over the course of the implementation of strategic decisions. On the one hand, contributors taking a forward-looking, predictive approach to strategic decision-making have focused primarily on the initial construction of SWRs and their impact on subsequent behaviour (Gavetti & Rivkin, 2007; Gavetti & Levinthal, 2000; Lovallo et al., 2012; Maitland & Sammartino, 2015a,b). On the other, contributors emphasising the adaptive nature of strategic decision-making have downplayed the role of cognitive representations by adopting a non-predictive, "backward-looking" approach (Gavetti et al., 2012: 26; Patvardhan & Ramachandran, 2020), perhaps because adaptations are seen as "comparatively ad hoc, serendipitous, and/or messy" rather than the product of careful thought (Klingebiel & De Meyer, 2013: 133; Cyert & March, 1963; Langley et al., 1995).

There is a significant gap in the literature, therefore, especially in situations of Knightian uncertainty, where the ability to anticipate the future and construct adequate SWRs at the outset is highly limited, and decision-makers often have difficulty adapting their cognitive representations effectively over time (Tripsas & Gavetti, 2000: 1148). This issue is a pressing one even when decision-makers create, at least in part, the future through their own actions (Alvarez & Barney, 2007; Gavetti & Porac, 2018) by adopting a “shaping-oriented forward-looking” approach (Patvardhan & Ramachandran, 2020). Our aim in what follows is to take some steps towards filling this gap, by recognizing the merits of engaging in forward-looking behaviour while adopting the view that strategic decision-making under Knightian uncertainty is nevertheless a dynamic, adaptive process in which decision-makers are called on to systematically update their cognitive representations as additional information emerges.

THEORIZING THE SCENARIO SPACE

In addressing how decision-makers should update their SWRs we will focus specifically on the task of updating the SWR analogues of small world states and the state space, which to avoid confusion we henceforth refer to as scenarios and the scenario space respectively. We do so because, as we shall show, the way in which decision-makers update the scenario space is pivotal to the process of updating SWRs. By scenarios we mean relatively coarse and partial descriptions of future states of affairs that may be imagined by a decision-maker in the course of constructing and updating a SWR. Scenarios are therefore hypotheses about what may occur in the future, that specify how currently unknown and contingent parts of reality might unfold (Jungermann & Thuring, 1988; Porter, 1985; Schoemaker, 1993). The scenario space at a given point in time is then the set of scenarios the decision-maker currently regards as possible in relation to a particular decision.

There are three general points to highlight about scenarios and the scenario space. The first

is that scenarios have to be constructed and are thus the product of managerial imagination and other cognitive attributes. We will assume this is done on the basis of projections from some set of imagined influences, an influence being any category of event or condition that could contribute causally to the realisation of a state of affairs. For example, when deciding whether to introduce a new product in a particular market, relevant influences might include future events such as competitor moves and regulatory changes, and background conditions such as prevailing technologies and existing market demographics. The second point, one we have already emphasised, is that scenario spaces are almost always incomplete (Gilboa, Minardi, & Samuelson, 2017; Shackle, 1979), typically running to only a relatively small number of scenarios (Maitland & Sammartino, 2015a; see also Kahneman & Lovallo, 1993; Lovallo & Kahneman, 2003).

The third and final point is that while an incomplete scenario space rules out the use of classical probabilities, since these require a complete list of mutually exclusive eventualities (Bradley & Drechsler, 2014; Ghirardato, 2001; Feduzi & Runde, 2014), it does not preclude scenarios being the subject of qualitative judgements of probability. That is to say, strategic decision-makers may yet be able to arrive at judgements to the effect that a scenario is possible (has a non-zero probability) or is impossible (has a zero probability), or to make judgements of the form “Scenario A is highly probable”, “Scenario A is more probable than Scenario B” or “Scenario A is as probable as Scenario B” on the basis of the available evidence (see Keynes (1921) for a theory of qualitative probability on these lines). We will proceed on the basis that strategic decision-makers are able to make judgements of this kind when assessing their scenarios, and to avoid the risk of these being confused with probability judgements as standardly interpreted we henceforth refer to them as judgements of *credibility*.

We are now in a position to theorize how the scenario space may be constructed and updated over time. To facilitate the analysis, we draw on some foundational concepts introduced by

Faulkner, Feduzi and Runde (2017). Faulkner et al. define a *known* as “any feature of the world that an individual has knowledge of”, and an *unknown* or *gap in knowledge* as “any feature of the world that an individual lacks knowledge of” (2017: 1282). A *feature of the world* is any fact “about past, present and future reality” (2017: 1281) such as the current state of the shale gas market in the US, Euro CO emission standards applying to diesel vehicles in present-day London, or the closing level of the FTSE on 1 July 2025. They further distinguish between a *known unknown*—“a gap in knowledge that an individual knows about and is aware of at the relevant time”—and an *unknown unknown*—“a gap in knowledge that an individual is not aware of at that time, either because they do not know about that gap in knowledge or because, despite knowing of it, they are unaware of it” (2017: 1283).

Building on these concepts Faulkner et al. introduce the notion of *hypothetical values* associated with a gap in knowledge, where a hypothetical value is “any value—outcome, state of affairs, result, quantity and so on—that could conceivably be thought to be a candidate for the actual or true value of the unknown under consideration” (2017: 1283). Thus, if the unknown concerns the Euro CO emission standards applying to diesel vehicles in London, for example, the associated hypothetical values would likely include numerous different quantities such as 1.00 g/km, 2.72 g/km, and so on. Finally, the *set of hypothetical values* associated with an unknown is “the entire collection of values that could conceivably be regarded as the true value of that unknown by any person within the group or community concerned” (2017: 1283). This set therefore represents the universe of conceivable hypothetical values at a given point in time, comprising all values that could potentially be imagined by a decision-maker from that community, given the prevailing background conditions (state of knowledge, cognitive abilities, etc.).

To relate these concepts to scenarios and scenario spaces, consider a decision-maker facing a strategic decision problem. The future course of events as they pertain to that decision is a known

unknown to the decision-maker, with hypothetical values associated with this unknown consisting of alternate futures that might be envisioned. The crucial link to the SWR framework is that these hypothetical values are potential scenarios, that is, imagined future states of affairs that might be included in the decision-maker's scenario space. The set of hypothetical values with respect to the unknown future is the set of all conceivable scenarios, comprising each and every scenario that any person in the broader community, put in the position of decision-maker, might conceivably imagine at the present time. We will call this set, *the set of imaginable scenarios*.

To further characterize the scenario space, we make three assumptions about the scenarios a specific decision-maker facing a strategic decision problem comes up with:

1. The scenarios consciously imagined by the decision-maker are typically fewer in number than conceivably could have been imagined. In particular, the number of items in the set of imaginable scenarios will generally far exceed the number of items in the decision-maker's scenario space.
2. The scenarios consciously imagined by the decision-maker could include scenarios that they regard as impossible.
3. Each of the scenarios consciously imagined by the decision-maker is either genuinely possible or impossible at any point in time, and where which of the two they are is something that depends on the way the world is rather than on what the decision-maker believes about the world.³

³ While scenarios are only coarse and partial descriptions of the future, we assume they have sufficient content to be classified as descriptions of states of affairs that are either genuinely possible or genuinely impossible at any given point in time.

Given these assumptions, at any given point in time during a decision problem each element of the set of imaginable scenarios can be classified according to three attributes: whether or not the decision-maker has consciously considered that scenario; if they have, whether they regard it as possible or not; and whether that scenario is genuinely possible or not. The resulting classification of scenarios is depicted in Table 1, which assigns every conceivable scenario to one of six cells.

Insert Table 1 about here

Cells 1 and 2 contain scenarios imagined and regarded as possible by the decision-maker, split between those that are correctly (cell 1), and those that are mistakenly (cell 2), regarded as such. Taken together, the scenarios in these two cells therefore constitute the decision-maker's *scenario space*, indicated by the shaded area. The scenarios in cells 3 and 4 are those consciously imagined by the decision-maker but excluded from the scenario space because they are regarded as impossible, correctly so in cell 4 but mistakenly so in cell 3. The scenarios in cells 5 and 6 also lie outside the scenario space, having not even entered the imagination of the decision-maker.

As a benchmark against which to assess the accuracy of a decision-maker's scenario space, we define a scenario space as *correct* when it contains all, and only, those conceivable scenarios that, at the present time, are genuinely possible future states of affairs. In Table 1, a correct scenario space therefore requires that cells 3 and 5 in the left-hand column, and cell 2 in the right-hand column, are empty. This standard is unlikely ever to be met in practice, the scenario space falling short of being correct whenever a decision-maker commits an "error of omission" in failing to include a genuinely possible scenario space, and/or an "error of commission" by including a genuinely impossible scenario.

In light of the earlier discussion of the factors influencing decision-makers' ability to construct and update SWRs, it is the awareness of relevant influences that is likely to be crucial to

their capacity to populate the scenario space with genuinely possible scenarios. Recall that a scenario is a projected future state of affairs that flows from a set of imagined influences. It then follows that a decision-maker's gaps in knowledge about particular influences implies corresponding shortcomings in the scenario space, most notably the omission, either through being unimagined or regarded as impossible, of scenarios that might genuinely arise from these unrecognized influences. The scale of the impact on the scenario space depends in part on whether such influences are known unknowns, or unknown unknowns, for the decision-maker. In the former case, the decision-maker will at least be in a position to contemplate hypothetical values associated with the relevant influence, even if he or she may not be able to identify all or even a representative spread of them and construct the associated scenarios. The latter case is likely to be more damaging, since a decision-maker who is unaware of having a gap in knowledge about an influence will not be in a position even to begin contemplating its range of associated hypothetical values and the scenarios these may lead to.

While Table 1 provides a useful way of representing the scenario space, and of evaluating the scope and accuracy of its contents at a given point in time, the static nature of the analysis limits what can be said about the updating of scenario spaces during the implementation of a strategic decision and how the scenario space might change as a result. By *updating* the scenario space, we mean the process by which the decision-maker revises their beliefs about the future in response to new or newly-considered information and thereby may be led to include new and/or exclude existing scenarios in/from the scenario space.⁴ From the perspective of Table 1, a decision-maker

⁴ The updating process could leave the scenario space unchanged if the new or reconsidered information does not warrant the inclusion of a new scenario or the exclusion of an old one. Note also that our notion of updating differs from that of Bayesian updating, where the latter only

updating their beliefs implies the vertical movement of scenarios between cells, that is, from a cell in one row of the table to a cell in another row within the same column. The possibilities are as follows:

1. scenarios previously unimagined are now considered and regarded as possible (scenarios in cells 5 and 6 move to cells 1 and 2);
2. scenarios previously considered possible are now forgotten, rendering them unimagined (scenarios in cells 1 and 2 move to cells 5 and 6);
3. scenarios previously considered impossible are now considered possible (scenarios in cells 3 and 4 move to cells 1 and 2);
4. scenarios previously considered possible are now considered impossible (scenarios in cells 1 and 2 move to cells 3 and 4);
5. scenarios previously unimagined are now considered and regarded as impossible (scenarios in cells 5 and 6 move to cells 3 and 4);
6. scenarios previously considered impossible are now forgotten, rendering them unimagined (scenarios in cells 3 and 4 move to cells 5 and 6).

Of these six possibilities only the first four are associated with changes to the decision-maker's scenario space. The first and third involve additions, augmenting the scenario space with scenarios now considered possible that had previously not been imagined by the decision-maker or had been imagined but disregarded as impossible. The second and fourth involve eliminations

involves revising probability judgements about scenarios already included in the scenario space and thus excludes the possibility of adding new scenarios (Feduzi & Runde, 2014).

from the scenario space, removing scenarios previously considered possible but now either regarded as impossible or disregarded entirely.

These findings lead to Table 2, a dynamic version of Table 1 that depicts ways the scenario space may change during the updating process. The double-headed arrows represent the different movements of scenarios that may occur as a result of the decision-maker revising their beliefs during the process of updating, with the double-headed nature of each arrow reflecting that, for each, movement may be in either direction. The grey arrows denote movements implying changes to the membership of the scenario space, while the white arrows indicate movements that leave the contents of the scenario space unaltered.

Insert Table 2 about here

It is worth highlighting two further points concerning changes to the scenario space and the updating process, both arising from the implementation of strategic decisions taking place over (possibly long periods of) time. The first is that developments in the world may result in conceivable scenarios that were formerly genuinely impossible becoming possible or *vice versa*.⁵ In Tables 1 and 2 this would imply horizontal movement of scenarios between cells, that is, from

⁵ Note that such developments may be either exogenous or endogenous to the strategic decision-maker under consideration. The former case is illustrated by an exogenous shock to an industry, such as the imposition of a new government regulation that renders impossible scenarios that were formerly possible. The latter case is illustrated by the action of an entrepreneur who, in developing a new product, endogenously creates the conditions that render possible scenarios that were formerly impossible. For a discussion of the two cases in the field of entrepreneurship, see Alvarez and Barney (2007).

a cell in one column to a cell in the other column within the same row. Crucially, however, such developments imply no change to the membership of the decision-maker's scenario space. At most, and unbeknown to the decision-maker, the effect on the scenario space is to shift scenarios between cell 1 and cell 2, either because a scenario the decision-maker (until now wrongly) regards as possible is now rendered genuinely possible (moving from cell 2 to cell 1), or else a scenario the decision-maker (until now rightly) regards as possible is now rendered genuinely impossible (moving from cell 1 to cell 2). Other cases, either where scenarios move between cells 3 and 4, or between cells 5 and 6, have no impact on the scenario space.

The second point is that the world may develop in ways leading either to the emergence of novel scenarios, meaning possible future states of affairs that were previously unimaginable given the prevailing background conditions and so previously outside the set of imaginable scenarios, or to the outright elimination of scenarios, meaning possible future states of affairs that were previously imaginable that then become entirely unimaginable given the prevailing background conditions.⁶ In effect, our earlier analysis assumes that the set of imaginable scenarios is fixed and unchanging over time, with the members of this set distributed between the six cells in Tables 1 and 2. Once it is acknowledged that this set may change over time, however, entirely new scenarios may enter, or existing scenarios exit, the table.

The implication for the scenario space during updating is that scenarios that were previously wholly unimaginable may be added to the scenario space, or scenarios that were formerly in the

⁶ The former case is illustrated by the discovery of a new substance with hitherto unimagined properties that leads people to be able to conceive of previously unimaginable scenarios. The latter case, while perhaps less likely to occur in practice, is illustrated by the loss of shared knowledge resulting from the burning of the library of Alexandria.

scenario space may become unimaginable. To allow for these possibilities we now amend Table 2, adding double-headed arrows from cells 1 and 2 that comprise the scenario space, to an area outside the table denoting previously unimaginable scenarios. We have also removed the two white arrows that denoted changes that would leave the scenario space unaltered. The resulting Table 3 therefore fully characterizes the ways the scenario space might change during updating.

Insert Table 3 about here

We conclude this section with some brief comments on the decision-maker's objective in updating the scenario space. In simple terms, updating aims to improve the accuracy of the scenario space, drawing on new and existing information to i) re-evaluate earlier assessments that are likely to have been generated under significant information, time and other constraints, and ii) take account of developments in the world that have occurred since any previous updating. The key issue here, however, is what improving the accuracy of the scenario space entails.

Since a perfectly accurate scenario space is one that is correct in the sense defined earlier, improving the accuracy of the scenario space, in principle at least, involves moving closer towards the correct scenario space. This means seeking to include a greater number of genuinely possible scenarios in the scenario space (i.e., reducing errors of omission) by moving scenarios in the left-hand column of Table 3 upwards via the arrows into cell 1, while also trying to reduce the number of genuinely impossible scenarios in the scenario space (i.e., reducing errors of commission) by moving them downwards in the right-hand column of the table from cell 2 to cell 4 (in the first instance). By definition this approach to updating leads in the limit to a correct scenario space.

Of course, given strategic decision-makers' cognitive make up, the nature of the informational environment in which they are operating and other factors that may be in play, it is unlikely that updating could achieve anything close to a correct scenario space in practice. But that

is not to say that strategic decision-makers could not be more effective in updating their scenario spaces in many cases, or that the procedures they might use to do so are equally effective. To pursue these matters we now turn to the psychology literature on human reasoning, in particular to its findings in relation to hypothesis-testing behaviour.

DISCONFIRMATION AND COUNTERFACTUAL REASONING IN HYPOTHESIS- TESTING

Hypothesis testing—a key function of the mind aimed at establishing whether, based on the available evidence, a hypothesis is supported and should therefore be retained, or is refuted and should therefore be revised or abandoned (Evans, 2007)—is an important area of research for psychologists studying human reasoning (Poletiek, 2001). Since scenarios are themselves hypotheses about alternative futures, updating the scenario space is *de facto* an exercise in hypothesis testing. It is therefore natural to ask what might be learned from the psychology of reasoning about how decision-makers may revise their scenario spaces during the course of implementing strategic decisions.

In what follows we focus on two simple hypothesis-testing heuristics—disconfirmation and counterfactual reasoning—that people are found to actually employ when solving decision tasks and that have been extensively studied in the psychology of reasoning (Cherubini, Castelvechchio & Cherubini, 2005; Farris & Revlin, 1989a,b; Gale & Ball, 2009; Gorman & Gorman, 1984; Gorman, Stafford, & Gorman, 1987; Klayman & Ha, 1987; Oaksford & Chater, 1994; Tukey, 1986). In the present section we characterize the two heuristics, using Wason’s famous 2-4-6

experiment (Wason, 1960; Evans, 2016) to illustrate the steps involved in their implementation.⁷ In the following section we then transpose the two heuristics to the strategic decision-making context, developing procedures for updating the scenario space based on each heuristic.

Disconfirmation and counterfactual reasoning

Wason's (1960) 2-4-6 task is a rule-discovery problem widely used in experimental studies of hypothesis-testing behaviour. In this task the experimenter privately chooses a rule—"three numbers in increasing order of magnitude" in Wason's original study—that generates particular triples of numbers. Subjects are then presented with a triple consistent with that rule—2-4-6 in Wason's study—and invited to uncover the rule. To do so, they are required to generate hypotheses about what the rule might be and to test those hypotheses by providing triples of their own to the experimenter, who then tells them whether or not the triples conform to the chosen rule. When subjects believe they have identified the rule, they are instructed to declare it to the experimenter, who tells them whether they are correct or not. If the rule the subject declares is incorrect the task continues. The task ends when a subject correctly identifies the rule, the allotted time runs out, or the subject gives up.

Subjects employ a variety of hypothesis-testing strategies in attempting to solve Wason's task and its variants (Tukey, 1986), including disconfirmation and counterfactual reasoning (Farris & Revlin, 1989a,b; Gorman & Gorman, 1984). These last two strategies represent an *eliminative* approach in being aimed at rejecting, rather than confirming, a hypothesis under consideration.

⁷ As different contributors provide slightly different accounts of what these methods entail, some of which contain internal contradictions (Farris & Revlin, 1989a,b; on this point, see Oaksford & Chater, 1994), we will attempt to interpret their main tenets as typically understood within the psychology literature (Evans, 2016) but in way that avoids the aforementioned contradictions.

Disconfirmation is a basic form of Popperian-style falsification that recommends that the subject start with a working hypothesis explaining some observed phenomenon, assume this hypothesis is correct, and then perform a series of tests that involve deliberately looking for new evidence that is inconsistent with the working hypothesis. If such evidence is found, then the working hypothesis is eliminated, a new hypothesis generated, and the process begins again. The higher the number of tests passed by any working hypothesis, in the sense of not being eliminated on the basis of new evidence, the greater the subject's confidence in it.

Table 4 specifies the steps involved in disconfirmation and classifies each in terms of the three broad activities involved in hypothesis-testing, namely hypothesis formation, evidence collection and feedback evaluation. The final column illustrates the use of the heuristic in Wason's 2-4-6 task, enumerating the implementation of each step from the beginning of the experiment.

Insert Table 4 about here

Counterfactual reasoning also directs the subject to attempt to eliminate the current working hypothesis. However, unlike disconfirmation it does so by looking for evidence in favour of a specific alternative hypothesis (Farris & Revlin, 1989a,b; Gorman, Stafford, & Gorman, 1987; Oaksford & Chater, 1994; Tukey, 1986). The subject again starts with a working hypothesis explaining some observed phenomenon and performs a series of tests. But in this case the subject assumes that the working hypothesis is false. In particular, they generate an alternative to the working hypothesis that is consistent with the existing evidence, and then deliberately look for new evidence that would be consistent with the alternative hypothesis and inconsistent with the original working hypothesis. If such evidence is found, then the original working hypothesis is eliminated and the alternative becomes the new working hypothesis, and the process begins again. The higher the number of tests passed by any working hypothesis, the greater the subject's confidence in it.

The steps and activities involved in counterfactual reasoning are given in Table 5, together with their implementation from the beginning of Wason's 2-4-6 task.

Insert Table 5 about here

USING DISCONFIRMATION AND COUNTERFACTUAL REASONING IN UPDATING SCENARIO SPACES

We now transpose disconfirmation and counterfactual reasoning as represented in the psychology literature into a strategic decision-making context, addressing in particular the question of how a decision-maker might update their scenario space over the course of the implementation of a strategic decision. To this end Table 6 describes procedures that operationalise each heuristic method for use in this context.

Insert Table 6 about here

The disconfirmation-based (henceforth D-based) procedure begins with the decision-maker assuming for the purpose of the exercise that their existing scenario space is correct in the sense of containing all, and only, those scenarios that are genuinely possible future states of affairs. The decision-maker then selects one of these scenarios and looks for new evidence against it, meaning information that, if uncovered, would reduce the credibility associated with that scenario. Once this search is completed the decision-maker then re-evaluates the credibility of all scenarios in the scenario space in light of the new evidence. When the evidence is so strong that a scenario can no longer be considered possible and must be removed from the scenario space, the decision-maker is obliged to generate an alternative scenario not currently in the scenario space, but consistent with the now expanded evidence, by imagining a new influence they had not considered before. The

procedure is then repeated as resources allow or until the updated scenario space is deemed adequate.

In the case of the counterfactual reasoning-based (henceforth CR-based) procedure, the decision-maker begins by assuming the existing scenario space is incorrect in the sense of containing one or more scenarios that are in fact impossible and/or omitting one or more scenarios that are genuinely possible. The decision-maker must then generate an alternative scenario not currently in the scenario space, consistent with the existing evidence, by imagining a new influence they had not considered before. Once done, the decision-maker then looks for new evidence in favour of this newly created scenario, that is information that, if uncovered, would increase the credibility associated with that scenario. Having completed this search, the decision-maker then evaluates the impact of the new evidence on both the new and existing scenarios, retaining the new scenario and/or eliminating existing scenarios as justified by the evidence. The procedure is then repeated as resources allow or until the scenario space is deemed adequate.

EVALUATING THE COMPARATIVE PERFORMANCE OF DISCONFIRMATION- AND COUNTERFACTUAL REASONING-BASED PROCEDURES IN UPDATING SCENARIO SPACES UNDER KNIGHTIAN UNCERTAINTY

We now examine the relative performance of the D- and CR-based procedures in updating scenario spaces under conditions of Knightian uncertainty. To prepare the ground, we begin with findings from the psychology literature on the impact of ecological considerations on the effectiveness of disconfirmation and counterfactual reasoning in Wason's 2-4-6 task and some of its variants. After characterizing the environment in which strategic decisions take place and the ways in which this differs from the experimental settings of Wason-style tasks, we evaluate the two procedures relative to three sets of criteria germane to improving the accuracy of the scenario space and so promoting more effective decision-making in strategic settings.

On the performance of disconfirmation and counterfactual reasoning

In his original study Wason found that subjects performed poorly in the 2-4-6 task, something he attributed to their tendency to rely on what we will refer to as a *confirmation* heuristic of testing hypothesized rules by generating triples consistent with those rules.⁸ For example, subjects presented with the triple 2-4-6 often come up with the hypothesis that the experimenter's rule is "three even numbers", and proceed to test this hypothesis by generating further triples consistent with this rule such as 12-14-16, 18-20-22 and 24-26-28. Given that such triples are also consistent with the experimenter's rule of "three numbers in increasing order of magnitude", they elicit a series of positive responses from the experimenter that lead subjects to the erroneous conclusion that they have correctly identified the experimenter's rule.

⁸ Wason's (1960, 1966) subjects' repeated or even exclusive use of tests aimed at confirming their current hypotheses is now commonly referred to as the "confirmation bias". However, we use the term confirmation heuristic here for two reasons. First, the use of a heuristic method does not necessarily reflect bias, if the latter is understood to imply a flawed or irrational approach to problem solving (Klayman, 1995; McKenzie, 2004). This point underpins the distinction we make between heuristics and biases in the present paper and is consistent with the broader literature on the ecological rationality of using heuristics in situations of Knightian uncertainty (Mousavi & Gigerenzer, 2017). Second, we will use the term confirmation bias in a different, broader, way later in the paper. In referring to a confirmation heuristic we also acknowledge that the term *positive test heuristic* (Evans, 2014) would be a more accurate label for subjects' tendency "to test cases that are expected to have the property of interest" (Klayman & Ha, 1987: 211). For simplicity, however, we treat the terms confirmation heuristic and positive test heuristic as equivalent.

Wason (1960: 139) proposed that the key to success in the 2-4-6 task is to attempt to eliminate rather than confirm the hypothesis under consideration, a suggestion that generated a stream of studies exploring whether use of some form of disconfirmation or counterfactual reasoning might improve subjects' hypothesis-testing performance (Farris & Revlin, 1989a,b; Gorman, Stafford, & Gorman, 1987; Oaksford & Chater, 1994; Tukey, 1986). The main finding of this work concerns the critical importance of ecological considerations: that rather than one heuristic dominating in all circumstances, the performance of different heuristics in Wason-style tasks varies with the *structure of the task environment* (McKenzie, 2004).

This finding emerges from the results of two quite different types of study. In the first, particular attributes of the 2-4-6 task, most notably the nature of the relationship between the experimenter's true rule and the subject's working hypothesis, are shown to affect the efficacy of different heuristics (Klayman & Ha, 1987; Oaksford & Chater, 1994). Thus, in the case described above in which the subject's hypothesised rule is a restricted version of the experimenter's true rule, it follows as a matter of logic that both disconfirmation and counterfactual reasoning can be effective in facilitating the discovery of the experimenter's true rule by promoting elimination of an incorrect working hypothesis. In contrast, in the opposite case in which the experimenter's rule is a restricted version of the subject's hypothesised rule, disconfirmation and counterfactual reasoning are invariably ineffective. Between these two extremes lies a continuum of cases in which the efficacy of the heuristics varies according to the form of the relationship between the true and hypothesized rules.

The importance of the task environment is also evident in the results of studies of the second kind, looking at the performance of the two heuristics in an experimental setting. Here minor variations in the procedures adopted have been found to have considerable impact on their effectiveness. For example, while Tweney et al. (1980) observed no difference between the

performance of subjects using disconfirmation and those using confirmation with Wason's original task protocol, Gorman and Gorman (1984) found that subjects instructed to use disconfirmation perform better than those instructed to use confirmation, or simply choose their own strategy, with the protocol altered so that the experimenter provides feedback on rules guesses only at the end of the experiment. Similar variations arose in experiments in which subjects were induced to employ counterfactual reasoning through the use of a facilitatory manipulation of the 2-4-6 task called "dual goal instructions" in which subjects must discover two rules. Thus, while Vallée-Tourangeau, Austin, & Rankin (1995) found that counterfactual reasoning typically improves performance over subjects choosing their own strategy in this case, Wharton, Cheng, & Wickens (1993) found that this was not necessarily true, the variation in results being the product of minor differences between the two studies in the information given to subjects regarding the relationship between the two rules (on this point, see Gale & Ball, 2009).

Strategic situations of Knightian uncertainty

The overriding message we take from the literature just reviewed is the need to be sensitive to the environment in which strategic decision-making takes place when assessing the merits of the D- and CR-based procedures for updating scenario spaces. The present section is accordingly devoted to characterizing environments of this kind, which we will henceforth refer to as strategic situations of Knightian uncertainty. We focus on three characteristics we regard as having most bearing on our subsequent evaluation of the two procedures, explaining each with reference to differences between strategic situations of Knightian uncertainty and the environments of Wason-style experiments.

The first characteristic concerns the nature of the task facing the strategic decision-maker. Unlike Wason's subjects who are called on to discover what they know to be a single true rule and to render this rule with complete accuracy, the aim of strategic decision-makers is to come up with

a more accurate scenario space. Doing so means simultaneously entertaining multiple scenarios in order to increase the number of genuinely possible, while decreasing the number of genuinely impossible, scenarios in the scenario space, and where scenarios are necessarily coarse and partial rather than wholly accurate descriptions of possible futures. Of course, to commit to the existence of a number of genuinely possible scenarios at any one point in time is to commit to a world in which the course of history is not fully determined, and on this point we share the intuitions of Knight (1921: 219-222) himself (see also, Lawson, 1988: 45).

The second characteristic is that, unlike Wason's 2-4-6 task in which the true rule remains constant throughout the course of the experiment, in strategic situations of Knightian uncertainty the set of genuinely possible scenarios a decision-maker is aiming to incorporate in their scenario space may change over time. This has two important implications for decision-making. First, it may happen that none of the scenarios included in the scenario space at the outset of a decision problem (say, time t) remain genuinely possible and thus a relevant guide to action at a later point of the implementation phase (say, time t_{+1}), even if the scenario space at time t was correct. Second, as the body of evidence upon which the scenario space was built at time t might have become completely irrelevant by time t_{+1} , decision-makers updating their scenario spaces might have to generate wholly new scenarios that are disjoint from all of the initial scenarios.

The final characteristic we highlight is that, in contrast to Wason's subjects who receive perfectly reliable, "deterministically accurate", feedback from the experimenter during the task (Klayman & Ha, 1987), the new information available to decision-makers in strategic situations of Knightian uncertainty will usually be partial, subject to doubt, open to multiple and possibly conflicting interpretations, and at odds with at least some of the existing information. That is to say, the evidence decision-makers have access to when updating their scenario spaces will often be profoundly ambiguous and non-definitive, two qualities we regard as major sources of Knightian

uncertainty.

By evidence being ambiguous we mean that it leaves a decision-maker uncertain about whether it increases or decreases the credibility of a scenario. The greater the degree of ambiguity, the more uncertain a decision-maker is in this regard. The significance of ambiguity for the procedures described in Table 6 lies in that it leaves decision-makers susceptible to errors when analysing how new evidence bears on the credibility of scenarios. These errors may take the form of false positives, where new evidence is wrongly judged to increase the credibility of a scenario, or false negatives, where new evidence is wrongly judged to decrease the credibility of a scenario.

By evidence being non-definitive we mean that, while it may lead to a change in the credibility of a scenario, it falls short of being able to justify the exclusion of a scenario in the “all-or-none manner” (Evans, 2016: 2082) characteristic of the evidence obtained in Wason-style tasks. The significance of non-definitive evidence thus lies in that it precludes the analogue of what Bacon (1620) called *crucial* experiments that would allow scenarios to be quickly and definitively eliminated from the scenario space. Note that the two qualities identified here are not independent: ambiguity (of any degree) is sufficient, but not necessary, for evidence to be non-definitive, while non-definitiveness is neither necessary nor sufficient for ambiguity (of any degree).

Evaluating the D- and CR-based procedures

We now evaluate our D-and CR-based procedures for updating scenario spaces in strategic situations of Knightian uncertainty. To do so we need to specify appropriate criteria against which to judge their performance. This is a complex issue, since the effectiveness of any method for improving the accuracy of scenario spaces might meaningfully be assessed against a number of different criteria. In what follows we concentrate on three that seem to us particularly relevant to effective decision-making in strategic settings: (1) counteracting the confirmation bias, (2) promoting the exploration of the set of imaginable scenarios, and (3) facilitating action to mitigate

or exploit the consequences of what would otherwise have been Black Swans.

Counteracting the Confirmation Bias: The Ambiguity Effect. Although the term is often used in more specific ways (Nickerson, 1998), we adopt a broad definition of the confirmation bias as a general “inclination to retain, or a disinclination to abandon, a currently favored hypothesis” (Klayman, 1995: 386). We have selected counteracting the confirmation bias as our first criterion in view of management research finding that strategic decision-makers, no less than Wason’s subjects, are prone to this bias in their reasoning (Bazerman & Moore, 2013; Ng, Westgren, & Sonka, 2009; Gavetti & Rivkin, 2005; Miller, 2008; Russo & Schoemaker, 1992; Weick & Sutcliffe, 2007). The paucity of research on how SWRs should be updated notwithstanding, this literature provides strong *prima facie* evidence that strategic decision-makers will display the same tendency when updating their scenario spaces, namely, to be inclined towards retaining the current scenario space by seeking and interpreting information in ways favourable to existing scenarios.

While both disconfirmation and counterfactual reasoning may be effective in counteracting the confirmation bias in situations of low ambiguity such as Wason’s original experiment, the effectiveness of both of our procedures in this regard is likely to be reduced when ambiguity is high. The reason for this lies in the reluctance of decision-makers to attach significance to evidence they believe to be ambiguous, for fear of it leading to errors in updating the credibility of scenarios. In other words, the possibility of error may impede decision-makers from incorporating ambiguous evidence into their judgements to the extent they would have were that evidence known to be unambiguous (Gorman, 1986, 1989; Klayman, 1995; Koehler, 1993; McKenzie, 2004).

Crucially, however, the two procedures differ in how much their performance in counteracting the confirmation bias is affected by ambiguity. The D-based procedure starts by assuming that the existing scenario space is correct and then attempting to eliminate scenarios from the scenario space on the basis of negative evidence. Ambiguity implies the possibility of false

negatives, which has been shown to lead to “hypothesis perseveration” in experiments where decision-makers are aware of the possibility of error in the feedback they receive (Gorman, 1989: 389). Here decision-makers tend either to immunize their hypotheses against disconfirmation by treating disconfirming data as errors, or to correctly recognize that there was no error but then devote so much time to replicating experiments to be sure that they failed to test their hypotheses adequately (Gorman, 1986, 1989). In short, the threat of false negatives provides grounds for decision-makers to resist giving up on their existing scenarios immediately in the face of negative evidence. The more ambiguous the evidence, the less inclined decision-makers relying on the D-based procedure will be to give up on scenarios in their existing scenario space, and thus the less effective the procedure will be in offsetting the confirmation bias.

Ambiguity also reduces the extent to which the CR-based procedure offsets the confirmation bias, but the reduction is likely to be less significant in this case. The CR-based procedure begins by assuming that the existing scenario space is incorrect and then looking for positive evidence in favour of alternative scenarios. Consequently, ambiguous evidence once more renders decision-makers’ susceptible to errors when analysing the bearing of evidence on the credibility of a scenario, in this case to false positives rather than false negatives. The reluctance of decision-makers to accept such evidence again suggests that the more ambiguous the evidence, the less inclined decision-makers relying on the CR-based procedure will be to reduce the credibility of, and possibly give up on, scenarios in their existing scenario space, and the less effective the procedure will be in offsetting the confirmation bias.

However, unlike the D-based procedure, there are two closely-related tendencies that counteract the ambiguity of evidence reducing the power of the CR-based procedure to offset the confirmation bias: first, individuals being more inclined to interpret ambiguous evidence in a way that confirms rather than disconfirms their beliefs (McKenzie, 2004), and second, individuals being

more willing to incorporate ambiguous evidence when it confirms rather than disconfirms their beliefs (Klayman, 1995). Recall that, by design, the CR-based procedure instructs decision-makers to assume their existing scenario space is incorrect at the point scenarios are tested, and to then seek to confirm the newly generated alternative scenario. So long as decision-makers are faithful to the instruction to find evidence that confirms the alternative scenario, they will be more likely to interpret and then accept ambiguous evidence as being in favour of it. Any inclination to dismiss alternative scenarios on grounds of the potential falsity of evidence in their favour will thus be tempered.

Of course, false positives may lend alternative scenarios more credibility than they deserve, and thus lead to an excessive reduction of the credibility of the existing scenarios. But given that our starting point was that strategic decision-makers suffer from the confirmation bias, this effect is actually welcome in the present context for encouraging them to question their existing scenarios. We thus arrive at the following proposition.

Proposition 1: In strategic situations of Knightian uncertainty in which evidence is highly ambiguous, the CR-based procedure helps strategic decision-makers offset the confirmation bias and thereby question their existing scenario space more acutely than does the D-based procedure.

Exploring the set of imaginable scenarios. Prompted by the fact that, in strategic situations of Knightian uncertainty, the set of genuinely possible scenarios is liable to change significantly during the implementation of a decision, as well as evidence that decision-makers tend to be disinclined to generate alternative hypotheses, and when they do do so, to generate an unduly low number of them (Dougherty, Gettys, & Thomas, 1997; Fischhoff, Slovic, & Lichtenstein, 1978; Gettys, Pliske, Manning, & Casey, 1987; Heath, Larrick, & Klayman, 1998),

the second criterion against which we assess the performance of our two procedures is their ability to promote exploration of the set of imaginable scenarios. In particular, we will consider how the rate at which new scenarios are generated and included in the scenario space when using each procedure is likely to be affected by the non-definitiveness and the high degree of ambiguity of evidence.

Before we proceed it is worth noting that the two procedures are specified to counteract a closely related tendency exhibited by decision-makers when generating hypotheses, that of coming up with alternatives that are overly similar to one another and that occupy the same general frame (Heath et al., 1998). Thus, both procedures stipulate that when the decision-maker is required to generate a new scenario, this is done by imagining an influence they had not been considered before rather than simply by contemplating an additional scenario associated with influences already considered. By forcing the decision-maker to consider a new influence, the two procedures aim to promote the generation of a more diverse and heterogenous body of scenarios in the scenario space.

Turning to the rate at which new scenarios are generated and included in the scenario space, the D-based procedure is fundamentally *reactive* in that, only when sufficient negative information is found, does it move from the elimination of existing scenarios from the scenario space to the formulation of new scenarios. Rather than being a necessary element of the process of scenario-testing, the generation of new scenarios thus emerges only as a by-product of the process of eliminating existing scenarios. The CR-based procedure is instead *proactive* as it reverses the order of these activities, moving from the generation of new scenarios to the elimination of existing ones when justified by the evidence. The generation of new scenarios thus is integral to the process of scenarios-testing.

This difference in the order of activities fundamentally affects the relative performance of the two procedures in promoting exploration of the set of imaginable scenarios when information

is non-definitive. Using the D-based procedure, decision-makers generate a new scenario only when the search for evidence against a scenario in the current scenario space yields evidence sufficient to eliminate an existing scenario. However, while this eventuality may sometimes arise, the non-definitive nature of information means that new (negative) evidence will often only throw doubt on a scenario, and that it will take time to acquire the disconfirming evidence required to reduce the credibility of a scenario sufficiently for it to be eliminated. This is all the more so, the higher the degree of credibility decision-makers initially attach to an existing scenario. The speed of this process of generation, and accordingly the overall number of new scenarios contemplated and eventually included in the scenario space over time, is thus significantly—and negatively—affected by the non-definitive nature of the available evidence.

In contrast, decision-makers using the CR-based procedure begin with the assumption that their existing scenarios are deficient and are required to generate a new scenario that currently lies outside their scenario space and then look for evidence in support of it. As before, the acquisition of new (positive) evidence will often lead to only gradual increases in the credibility of a new scenario and corresponding decreases in the credibility of existing ones. However, since this procedure does not require an existing scenario to be eliminated for a new one to be generated, the non-definitive nature of information does not impact on the rate at which decision-makers contemplate alternative scenarios.

Moreover, the non-definitive nature of information should not dramatically affect the rate at which newly imagined scenarios are actually included in the scenario space. For unlike the D-based procedure, where the non-definitive nature of information is likely to significantly prolong the time it takes to find sufficient disconfirming evidence to eliminate a scenario, in the case of the CR-based procedure this adverse effect is largely ameliorated by the relatively small amount of

positive evidence needed to include a new scenario in the scenario space. And with decision-makers explicitly directed to look only for positive evidence in favour of new scenarios, the CR-based procedure pushes decision-makers to take them seriously and to find support for their inclusion in the scenario space. We thus arrive at the following proposition.

Proposition 2a: In strategic situations of Knightian uncertainty in which evidence is non-definitive, the CR-based procedure will lead to a more rapid accretion of new scenarios in the scenario space than the D-based procedure.

Turning to the highly ambiguous nature of the information typically available to strategic decision-makers, this reduces the effectiveness of both procedures in promoting the exploration of the set of imaginable scenarios. However, the reduction due to ambiguity is likely to be less significant for the CR-based procedure than it is for the D-based procedure.

Decision-makers using the D-based procedure generate a new scenario only when the newly acquired negative evidence reduces the credibility of a scenario already included in the scenario space sufficiently for it to be eliminated. The existence of ambiguity reduces the speed at which the degree of credibility is decreased since, given decision-makers' reluctance to attach significance to evidence they believe to be ambiguous, disconfirming data tend to be discounted when there is the possibility of false negatives. The fear of committing an error of omission will thus lead decision-makers to stick with their existing scenarios until data become more reliable. The consequence is that ambiguity reduces the speed at which new scenarios are generated, thereby having a further depressing effect on the rate at which new scenarios are contemplated and eventually included in the scenario space.

In contrast, the CR-based procedure encourages decision-makers to generate new scenarios that currently lie outside their scenario space and to look only for evidence that might increase their credibility. As the proactive generation of new scenarios is integral to the process of testing existing

scenarios, ambiguity does not affect the rate at which new scenarios are generated. However, the possibility of false positives arising from evidence being ambiguous might reduce the rate at which the degree of credibility of an alternative scenario is increased, thereby preventing its inclusion in the scenario space. The fear of committing an error of commission may thus lead decision-makers to dismiss the alternative scenarios unless and until data become more reliable. Ambiguity might thus reduce the rate at which new scenarios are included in the scenario space.

Nevertheless, there are two reasons the depressing effect of ambiguity on the rate at which new scenarios are included in the scenario space is likely to be less significant for the CR- than for the D-based procedure. First, as discussed above, decision-makers are more likely to take into account confirming data than disconfirming data, when data might be subject to error, and to interpret ambiguous evidence in a way that confirms rather than disconfirms their beliefs. Again, this tendency implies that the inclination to dismiss new scenarios on grounds of the potential falsity of evidence in their favour will be tempered. Of course, ambiguity might lead decision-makers to give new scenarios more credibility than they deserve. But this effect is actually welcome in the present context as it positively affects the speed at which decision-makers using the CR-based procedure include new scenarios in the scenario space.

Second, decision-makers who suffer from the confirmation bias are more likely to tolerate the possibility of errors of commission—including new scenarios on the basis of false positives—than the possibility of errors of omission—excluding existing scenarios on the basis of false negatives. This is because a new scenario can be added to the scenario space as long as a minimum amount of positive evidence is found and, unless the credibility of an existing scenario is reduced enough for it to be eliminated, without excluding any existing scenario. This leads to the following proposition.

Proposition 2b: In strategic situations of Knightian uncertainty in which

evidence is highly ambiguous, the CR-based procedure will lead to a more rapid accretion of new scenarios than the D-based procedure.

Black Swans and taking action. In strategic situations of Knightian uncertainty decision-makers' scenario spaces are prone to rapid obsolescence, and the likelihood of Black Swans (Taleb, 2007) and their associated risks and opportunities will be high. The question then arises whether, under such circumstances, the D- or CR-based procedure will be more effective in promoting the uncovering of what would otherwise have been Black Swans and thereby prompting action aimed at mitigating or exploiting their consequences.

According to Taleb's well-known formulation (2007: xvii–xviii), there are three conditions for an event to qualify as a Black Swan: that it is “an outlier, as it lies outside of the realm of regular expectations, because nothing in the past can convincingly point to its possibility”; that it “carries an extreme impact”; and that “in spite of its outlier status, human nature makes us concoct explanations for its occurrence after the fact, making it explainable and predictable”. We will ignore the third condition, which seems neither an essential property of unanticipated events (Faulkner et. al, 2017; Runde, 2009) nor particularly relevant to our concern about mitigating or exploiting their consequences. A natural interpretation of Black Swans from the point of view of this paper, then, is that they correspond to events or states of affairs that occur and have an extreme impact on a strategic decision at a specific point of the implementation phase (say, time $t+1$), but were not represented in the scenario space when it was constructed or last updated (say, time t).

There are three points to note about this interpretation. First, a Black Swan as experienced by one decision-maker may not come as even a mild surprise to another. For example, the eponymous black swan (the bird) that was the Black Swan experienced by the Dutch explorer Willem de Vlamingh in his first encounter with the species in Western Australia in 1697, would have been just another black swan to the indigenous population accustomed to their presence.

Second, and contrary to many peoples' intuition, events experienced as Black Swans do not necessarily have to be highly improbable. Context, and especially changes in context, often matter here, as can be seen from de Vlamingh's first encounter with a black swan actually being highly probable once he had entered its natural habitat. Likewise, in fast-changing environments, events or situations that come as significant surprises on first exposure may in fact have been highly probable given changed conditions (and may rapidly become the "new normal"). Finally, on our understanding of Black Swans they may be positive as well as negative in their impact.

We are now in a position to compare the relative performance of the D- and CR-based procedures in relation to uncovering what would otherwise have been Black Swans. To facilitate this comparison, we will adapt the two procedures by requiring decision-makers, whenever prompted to generate a new scenario, to do so by imagining a *crucial* influence they had not considered before. By a crucial influence we mean an influence that would raise a new scenario with extreme consequences for the decision under consideration, and thereby prompt the generation of new possible courses of action.

Take first the D-based procedure, which requires decision-makers to generate new scenarios only when the newly acquired negative evidence eliminates a scenario from the existing scenario space. As summarized in propositions 2a and 2b, when information is ambiguous and non-definitive, this feature limits the extent to which the procedure promotes the inclusion of new scenarios in the scenario space and consequently reduces the likelihood of uncovering possible high impact future events or states of affairs before they occur. Further, by requiring decision-makers to look only for evidence *against* existing scenarios, the D-based procedure anchors managerial attention to prevailing expectations rather than encouraging search for information, including as yet unimagined crucial influences, that might indicate a future very different from the one expected. In this way, the procedure again works against receiving early clues about potential

Black Swans. On both counts, the D-based procedure would hinder decision-makers in their efforts to anticipate in a timely way environmental changes that might have an extreme impact, and consequently will often be left having to make *ad hoc* adaptive responses as events unfold. In other words, their ability to update SWRs in a way that allows them to systematically devise, and then implement, previously unconsidered courses of action appropriate to radically new conditions is limited.

In contrast, the CR-based procedure offers clear benefits for decision-makers exposed to Black Swans. As shown in propositions 2a and 2b, by encouraging the constant and proactive generation of alternative scenarios, the procedure leads to a larger number of new scenarios being generated and eventually included in the scenario space over time. As decision-makers' attention is now directed to imagining crucial influences and associated high-impact scenarios, the procedure also promotes the early uncovering of a larger number of potential Black Swans. The instruction to look only for evidence in favour of these new scenarios encourages decision-makers to take them seriously and to be sceptical of the existing scenarios as a guide to action. Whenever a new high-impact scenario is included in the scenario space, and whether or not one or more existing scenarios are eliminated, decision-makers are called on to update their SWRs and identify and, eventually, implement alternative courses of action to exploit or mitigate what would otherwise have been Black Swans.

We conclude that in strategic situations of Knightian uncertainty in which decision-makers' cognitive representations quickly become obsolete and are seldom adapted in a timely way, the CR-based procedure is superior to the D-based procedure in helping update SWRs in a way that favours the early uncovering of what would otherwise have been Black Swans and the systematic identification and implementation of previously unconsidered courses of action. We thus arrive at our final proposition.

Proposition 3: In strategic situations of Knightian uncertainty in which evidence is non-definitive and highly ambiguous, the CR-based procedure is more effective than the D-based procedure in uncovering, and so facilitating action to mitigate or exploit the extreme consequences of, what would otherwise have been Black Swans.

DISCUSSION

Our analysis contributes to the literature on how decision-makers operate in situations of Knightian uncertainty in four ways. First, we have set out a detailed conception of SWRs and of key differences between these and Savage's (1954) small worlds, an exercise that made it possible to account for and systematize various types of uncertainty that contribute to the Knightian uncertainty often faced by strategic decision-makers and to highlight various implications for learning in strategic contexts. As the notion of SWRs is gaining some traction in the behavioural strategy and entrepreneurship literature (e.g., Csaszar & Levinthal, 2016; Levinthal, 2011, 2018; Maitland & Sammartino, 2015a; Packard et al., 2017) and it is often used interchangeably with terms such as mental models, simplified cognitive representations, cognitive frames, and so on (Maitland, 2015b: 757 fn. 5; Gavetti, Levinthal, & Rivkin, 2005; Grégoire et al., 2010), we hope that this exercise will be a useful contribution to the foundations of further work on cognitive representations in management and organizational contexts. At the same time, we hope that our conceptualisation of Knightian uncertainty might be of interest to those working towards a more nuanced treatment of uncertainty in entrepreneurial studies (Ramoglou & Tsang, 2016: 425; Packard et al., 2017; Townsend et al., 2018).

Second, we have developed a theoretical framework designed to lay bare the relationship between SWRs and the scenarios they contain, and thereby to distinguish between the *set of imaginable scenarios* and the *scenario space*, and to elaborate how scenario spaces may be updated

over time. We offer this framework as a potential resource for further research into strategic decision-making under Knightian uncertainty that touches on the question of how SWRs should be updated. For example, our emphasis on the mechanisms through which the scenario space changes over time and on establishing a clear connection between the exploration of the set of imaginable scenarios and the problem of Black Swans, might complement recent research on how to operate under Knightian uncertainty (Alvarez & Barney, 2005; Townsend et al., 2018), and then especially the studies inspired by G.L.S. Shackle (1961, 1972, 1979) (e.g., Packard et al., 2017; Porac & Tschang, 2013). Further, our framework might shed light on how assumptions about the kind of uncertainty faced by strategic decision-makers play into debates on the nature of entrepreneurial opportunities (Alvarez & Barney, 2007; Ramoglou & Tsang, 2016) and the merits of using predictive or non-predictive approaches to strategic decision making (Packard & Clarke, 2020; Patvardhan & Ramachandranb, 2020). There are also connections to be made with the work of behavioural strategists who have embraced Gilboa and Schmeidler's (1995, 2001) case-based decision theory (e.g., Gavetti & Rivkin, 2007; Lovallo, et al., 2012), given that even some of the original authors of this theory have recently argued that, in the presence of Black Swans, it may be necessary to complement case-based reasoning with scenario-thinking (Gilboa et al., 2017).

Third, we have extended recent work on the role of heuristics in strategic decision-making under Knightian uncertainty (Bingham & Eisenhardt, 2011; Maitland & Sammartino, 2015a; Shepherd, Haynie, & McMullen, 2012) by showing how heuristic methods of inquiry from the psychology of reasoning might inform the updating of SWRs. Apart from contributing directly to research on how strategic decision-makers construct and use representations to operate in situations of Knightian uncertainty (Csaszar & Levinthal, 2016; Gavetti & Rivkin, 2007; Maitland & Sammartino, 2015a;b; Porac, et al., 1989), we believe that our work might also be of interest to researchers working on cognate topics. These topics include the benefits and costs of

comprehensive decision-making processes in turbulent environments, the meaning of rationality in situations of Knightian uncertainty, and the microfoundations of forward-looking behaviour in strategy and organizational theory (Alvarez et al., 2013; Felin & Zenger, 2017; Felin, Kauffman, Koppl, & Longo, 2014; Forbes, 2007; Gavetti, 2012; Gavetti & Porac, 2018; Miller, 2008; Patvardhan & Ramachandranb, 2020; Porac & Tschang, 2013; Shepherd & Rudd, 2014).

Fourth, we have demonstrated the importance of ecological considerations, in addition to the decision maker's cognitive makeup, when attempting to assess the heuristics decision-makers use—or should use—to operate in situations of Knightian uncertainty (Gigerenzer, Todd, & the ABC Research Group, 1999: 13; Simon, 1990: 7). In particular, we have argued that the CR-based procedure is more effective than is the D-based procedure when updating SWRs in strategic situations of Knightian uncertainty in which the information is ambiguous and non-definitive. We hope that the approach we have used here might provide a template for further research in management aimed at taking into account the informational environment when assessing the merits of using heuristics in situations of uncertainty (Artinger, Petersen, Gigerenzer, & Weibler, 2015; Sarasvathy, Ramesh, & Forster, 2014; Shepherd et al., 2012; Zhang & Cueto, 2017).

While we believe that our discussion of CR- and D-based procedures for updating scenario spaces provides some useful results for the literature on how decision-makers operate in strategic situations of Knightian uncertainty, we have only scratched the surface of the issues involved. We close with three broad avenues in which our account might be further developed. First, in this paper, we have restricted ourselves to evaluating the relative effectiveness of the CR- and D-based procedures in strategic situations of Knightian uncertainty in which the information is ambiguous and non-definitive. Our conclusion that the former is more effective than the latter is specific to this informational environment, however, and it is quite possible that the D-based procedure may be superior to the CR-based procedure in other contexts. There is considerable scope for further

research here investigating how these procedures perform across a variety of informational environments (for different characterizations of the informational environment, see Todd, Gigerenzer, & the ABC Research Group, 2012), or indeed when additional situational constraints are added such as information-acquisition costs.

Second, our analysis has ignored various ways in which differences in the distribution of judgements of credibility between scenarios in the scenario space may be important. There is scope, for example, to explore how the specification of the D- and CR-based procedures in Table 6 might be improved by taking such differences into account, such as modifying the D-based procedure in a way that requires the decision-maker to look for evidence against a high credibility scenario in the scenario space rather than simply an arbitrary scenario. Similarly, in evaluating the relative performance of the two procedures against a given criterion our analysis could be extended to examine how the results depend on the distribution of credibility among the scenarios in the existing scenario space. Finally, the set of criteria we use to assess the relative performance of the procedures could be expanded to include consideration of the degree of credibility of the new scenarios introduced into the scenario space or, more generally, of those in the resulting scenario space as a whole.

Finally, while we have focused on two specific heuristic methods of inquiry that might inform procedures for updating SWRs in the present paper, comparative exercises of the kind we have conducted could also be performed in respect of other methods that have been proposed in the philosophy of science, such as a positive-test strategy (Klayman & Ha, 1987), Bayesian epistemology (Earman, 1992) and Bacon (1620) and Mill's (1967[1875]) methods of eliminative induction (for a review of how various philosophies of science may inform individuals' hypothesis-testing behaviour, see Tukey, 1986). An exercise on these lines would be particularly useful to explore the cases in which, rather than relying on only one method, it may be necessary to adopt a

mixed methods approach to solve a particular task. For example, some studies (Tweney et al., 1980) have suggested that disconfirmation may be useful only at later stages of a task, once at least some hypotheses have been identified and confirmed. If so, it would then be necessary to study which heuristic methods may be more effective in constructing and revising SWRs at different stages of the strategic decision-making process.

REFERENCES

- Alvarez, S. A. & Barney, J. B. 2005. How do entrepreneurs organize firms under conditions of uncertainty? *Journal of Management*, 31: 776-793.
- Alvarez, S. A., & Barney, J. B. 2007. Discovery and creation: Alternative theories of entrepreneurial action. *Strategic Entrepreneurship Journal*, 1: 11-26.
- Alvarez, S. A., Barney, J. B., & Anderson, P. 2013. Forming and exploiting opportunities: The implications of discovery and creation processes for entrepreneurial and organizational research. *Organization Science*, 24: 301-317.
- Alvarez, S. A., & Porac, J. F. 2020. Imagination, indeterminacy, and managerial choice at the limit of knowledge. *Academy of Management Review* (forthcoming).
- Arikan, A. M., Arikan, I., & Koparan, I. 2020. Creation opportunities: Entrepreneurial curiosity, generative cognition, and Knightian uncertainty. *Academy of Management Review* (forthcoming).
- Artinger, F., Petersen, M., Gigerenzer, G., & Weibler, J. 2015. Heuristics as adaptive strategies in management. *Journal of Organizational Behavior*, 36: S33-S52.
- Bacon, F. 1620. *Novum organum*. English translation in: R. L. Ellis & J. Spedding (Eds.), 1905, *The philosophical works of Francis Bacon*: 212-387. Routledge.
- Bazerman, M. H., & Moore, D. A. 2013. *Judgment in managerial decision making*. Hoboken: John Wiley & Sons.
- Bewley, T. F. 2002 [1986]. Knightian decision theory. Part I. *Decisions in Economics and Finance*, 25: 79-110.
- Berglund, H., Bousfiha, M., & Mansoori, Y. 2020. Opportunities as artifacts and entrepreneurship as design. *Academy of Management Review* (forthcoming).

- Bingham, C. B., & Eisenhardt, K. M. 2011. Rational heuristics: The ‘simple rules’ that strategists learn from process experience. *Strategic Management Journal*, 32: 1437-1464.
- Binmore, K. 2008. *Rational decisions in large worlds*. *Annales d'Économie et de Statistique*, 86: 25-41.
- Binmore, K. 2009. *Rational decisions*. Princeton: Princeton University Press.
- Bradley, R. 2017. *Decision theory with a human face*. Cambridge: Cambridge University Press.
- Bradley, R., & Drechsler, M. 2014. Types of uncertainty. *Erkenntnis*, 79: 1225-1248.
- Brandenburger, A. 2017. Where do great strategies really come from? *Strategy Science*, 2: 220-225.
- Busenitz, L., & Barney, J. 1997. Differences between entrepreneurs and managers in large organizations: Biases and heuristics in strategic decision-making. *Journal of Business Venturing*, 12: 9-30.
- Cherubini, P., Castelvechio, E., & Cherubini, A. M. 2005. Generation of hypotheses in Wason's 2-4-6 task: An information theory approach. *Quarterly Journal of Experimental Psychology*, 58A: 309-332.
- Csaszar, F., & Levinthal, D. A. 2016. Mental representation and the discovery of new strategies. *Strategic Management Journal*, 37: 2013-2049.
- Cyert, R., J., & March. 1992 [1963]. *A behavioral theory of the firm*. Malden, MA: Blackwell.
- Dekel, E., Lipman, B. L., & Rustichini, A. 1998. Recent developments in modelling unforeseen contingencies. *European Economic Review*, 42: 523-542.
- Dougherty, M. R. P., Gettys, C. F., & Thomas, R. P. 1997. The role of mental simulation in judgments of likelihood. *Organizational Behavior and Human Decision Processes*, 70: 135-148.

- Drechsler, M. 2012. *Rationality, decisions and large worlds*. Unpublished doctoral dissertation, The London School of Economics and Political Science, London.
- Earman, J. 1992. *Bayes or bust? A critical examination of Bayesian confirmation theory*. Cambridge: MIT Press.
- Ellsberg, D. 1961. Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics*, 75: 643-669.
- Evans, J. St. B. T. 2007. *Hypothetical thinking: Dual processes in reasoning and judgment*. Psychology Press.
- Evans, J. St. B. T. 2016. Reasoning, biases and dual processes: The lasting impact of Wason (1960). *Quarterly Journal of Experimental Psychology*, 69: 2076-2092.
- Farris, H. H., & Revlin, R. 1989a. Sensible reasoning in two tasks: Rule discovery and hypothesis evaluation. *Memory & Cognition*, 17: 221-232.
- Farris, H. H., & Revlin, R. 1989b. The discovery process: A counterfactual strategy. *Social Studies of Science*, 19: 497-513.
- Faulkner, P., Feduzi, A., & Runde, J. 2017. Unknowns, Black Swans and the risk / uncertainty distinction. *Cambridge Journal of Economics*, 41: 1279-1302.
- Feduzi, A., & Runde, J. 2014. Uncovering unknown unknowns: Towards a Baconian approach to management decision making. *Organizational Behavior and Human Decision Processes*, 124: 268-283.
- Felin, T., & Zenger, T. R. 2017. The theory-based view: Economic actors as theorists, *Strategy Science* 2: 258-271.
- Felin, T., Kauffman, S. A., Koppl, R., & Longo, G. 2014. Economic opportunity and evolution: Beyond landscapes and bounded rationality. *Strategic Entrepreneurship Journal*, 8: 269-282.

- Fischhoff, B., Slovic, P., & Lichtenstein, 1978. Fault trees: Sensitivity of estimated failure probabilities to problem representation. *Journal of Experimental Psychology: Human Perception and Performance*, 4: 330-344.
- Forbes, D. P. 2007. Reconsidering the strategic implications of decision comprehensiveness. *Academy of Management Review*, 32: 361-376.
- Gale, M., & Ball, L. J. 2009. Exploring the determinants of dual goal facilitation in a rule discovery task. *Thinking & Reasoning*, 15: 294-315.
- Gary, M. S., Wood, R. E., & Pillinger, T. 2012. Enhancing mental models, analogical transfer, and performance in strategic decision making. *Strategic Management Journal*, 33: 1229-1246.
- Gavetti, G., & Levinthal, D. 2000. Looking forward and looking backward: Cognitive and experiential search. *Administrative Science Quarterly*, 5: 113-137.
- Gavetti, G., & Rivkin, J. W. 2005. How strategists really think. *Harvard Business Review*, 83: 54-63.
- Gavetti, G., Levinthal, D., & Rivkin, J. W. 2005. Strategy making in novel and complex worlds: the power of analogy. *Strategic Management Journal*, 26: 691-712.
- Gavetti, G., & Rivkin, J. W. 2007. On the origin of strategy: Action and cognition over time. *Organization Science*, 18: 420-439.
- Gavetti G. 2012. Toward a behavioral theory of strategy. *Organization Science*, 23: 267–285.
- Gavetti, G., Greve, H. R., Levinthal, D., & Ocasio W. 2012. The behavioral theory of the firm: assessment and prospects. *The Academy of Management Annals*, 6: 1-40.
- Gavetti, G., & Menon, A. 2016. Evolution cum agency: Toward a model of strategic foresight. *Strategy Science*, 1: 207-233.

- Gavetti, G., Helfat, C. E., & Marengo, L. 2017. Searching, shaping, and the quest for superior performance. *Strategy Science*, 2: 194-209.
- Gavetti, G., & Porac, J. F. 2018. On the origin of great strategies, *Strategy Science*, 3: 352-365.
- Gettys, C. F., Pliske, R. M., Manning, C., & Casey, J. T. 1987. An evaluation of human act generation performance. *Organizational Behavior and Human Decision Processes*, 39: 23-51.
- Ghirardato, P. 2001. Coping with ignorance: unforeseen contingencies and non-additive uncertainty. *Economic Theory*, 17: 247-276.
- Gigerenzer, G., Todd, P. M., & the ABC Research Group. 1999. *Simple heuristics that make us smart*. New York: Oxford University Press.
- Gigerenzer, G., & Gaissmaier, W. 2011. Heuristic decision making. *Annual Review of Psychology*, 62: 451-482.
- Gilboa, I., & Schmeidler, D. 1995. Case-based decision theory. *Quarterly Journal of Economics*, 110: 605–639.
- Gilboa, I., & Schmeidler, D. 2001. *A theory of case-based decisions*. Cambridge: Cambridge University Press.
- Gilboa, I., Minardi, S., & Samuelson, L. 2017. *Cases and scenarios in decisions under uncertainty*. Working paper No. ECO/SCD-2017-1200, HEC Paris.
- Gorman, M. E. 1986. How the possibility of error affects falsification on a task that models scientific problem-solving. *British Journal of Psychology*, 77: 85–96.
- Gorman, M. E. 1989. Error, falsification and scientific inference: An experimental investigation. *Quarterly Journal of Experimental Psychology*, 41: 385-412.

- Gorman, M. E., & Gorman, M. E. 1984. A comparison of disconfirmatory, confirmatory and a control strategy on Wason's 2-4-6 task. *Quarterly Journal of Experimental Psychology*, 36: 629-648.
- Gorman, M. E., Stafford, A., & Gorman, M. E. 1987. Disconfirmation and dual hypotheses on a more difficult version of Wason's 2-4-6 task. *Quarterly Journal of Experimental Psychology*, 39: 1-28.
- Grégoire, D. A., Barr, P. S., & Shepherd, D. A. 2010. Cognitive processes of opportunity recognition: The role of structural alignment. *Organization Science*, 21: 413-431.
- Heath, C., Larrick, R. O., & Klayman, J. 1998. Cognitive repairs: How organizational practices can compensate for individual shortcomings, *Research in Organizational Behavior*, 20: 1-37.
- Hertwig, R., & Pachur, T. 2015. Heuristics, history of. In J. D. Wright (Ed.), *International encyclopedia of the social & behavioral sciences*: 829-835. Oxford: Elsevier.
- Jungerman, H., & Thuring, M. 1988. The labyrinth of the experts' minds: Some reasoning strategies and their pitfalls. *Annals of Operation Research*, 16: 117-130
- Kahneman, D., & Lovallo, D. 1993. Timid Choices and bold Forecasts: A cognitive perspective on risk taking. *Management Science*, 39: 17-31.
- Keynes, J. M. 1921. *A Treatise on probability*. Reprinted in: The Collected Writings of John Maynard Keynes, Vol. VIII, 1973. London: Macmillan.
- Keynes, J. M. 1937. The general theory of employment. *Quarterly Journal of Economics*, 51: 209-223.
- Klayman, J., & Ha, Y. W. 1987. Confirmation, disconfirmation and information in hypothesis testing. *Psychological Review*, 94: 211-228.

- Klayman, J., 1995. Varieties of confirmation bias, *Psychology of Learning and Motivation*, 32: 385-418.
- Klingebiel, R., & De Meyer, A. 2013. Becoming aware of the unknown: Decision making during the implementation of a strategic initiative, *Organization Science*, 24: 133-153.
- Knight, F. H. 1921. *Risk, uncertainty and profit*. Chicago: University of Chicago Press.
- Koehler, J. J. 1993. The influence of prior beliefs on scientific judgments of evidence quality. *Organizational Behavior and Human Decision Processes*, 56: 28-55
- Langley, A., Mintzberg, H., Pitcher, P., Posada, E., & Saintmacary, J. 1995. Opening up decision making: The view from the black stool. *Organization Science*, 6: 260-279.
- Lawson, T. 1988. Probability and uncertainty in economics analysis. *Journal of Post Keynesian Economics*, 11: 38-65.
- Levinthal, D. A. 2011. A behavioral approach to strategy — What's the alternative? *Strategic Management Journal*, 32: 1517-1523.
- Levinthal, D. A. 2018. From strategy to strategic organization. In M. Augier, C. Fang, & V. P. Rindova (Eds.), *Behavioral strategy in perspective (Advances in strategic management)*: 71-77. Bingley, UK: Emerald Publishing Limited.
- Lien, Y. W., & Lin, W. L. 2011. From falsification to generating an alternative hypothesis: Exploring the role of the new-perspective hypothesis in successful 2-4-6 task performance. *Thinking & Reasoning*, 17: 105-136.
- Lovullo, D. & Kahneman, D. 2003. Delusions of success: How optimism undermines executives' decisions. *Harvard Business Review*, 81: 56-63.
- Lovullo, D., Clarke, C., & Camerer, C. 2012. Robust analogizing and the outside view: two empirical tests of case-based decision making. *Strategic Management Journal*, 33: 496-512.
- Maitland, E., & Sammartino, A. 2015a. Decision making and uncertainty: The role of heuristics

- and experience in assessing a politically hazardous environment. *Strategic Management Journal*, 36: 1554-1578.
- Maitland, E., & Sammartino, A. 2015b. Managerial cognition and internationalization. *Journal of International Business Studies*, 46: 733-760.
- Manktelow, K. 2012. *Thinking and reasoning. An introduction to the psychology of reason, judgment and decision making*. Hove: Psychology Press.
- McMullen, J. S., & Shepherd, D. A. 2006. Entrepreneurial action and the role of uncertainty in the theory of the entrepreneur. *Academy of Management Review*, 31: 131-152.
- McKenzie, C. R. M. 2004. Hypothesis testing and evaluation. In D. J. Koehler & N. Harvey (Eds.), *Blackwell Handbook of judgment and decision making*: 200-219. Malden: Blackwell Publishing.
- Mill, J. S. 1967 [1875]. *A system of logic* (8th ed.). London: Longmans.
- Miller, C. C. 2008. Decisional comprehensiveness and firm performance: Towards a more complete understanding. *Journal of Behavioral Decision Making*, 21: 598-620.
- Mousavi, S., & Gigerenzer, G. 2017. Heuristics are tools for uncertainty. *Homo Oeconomicus*, 34: 361-379.
- Mukerji, S. 1997. Understanding the non-additive probability decision model. *Economic Theory*, 9: 23-46.
- Ng, D., Westgren, R., & Sonka, S. 2009. Competitive blind spots in an institutional field. *Strategic Management Journal*, 30: 349-369.
- Nickerson, R. S. 1998. Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, 2: 175-220.
- Oaksford, M., & Chater, N. 1994. Another look at eliminative and enumerative behaviour in a conceptual task. *European Journal of Cognitive Psychology*, 6: 149-169.

- Packard, M. D., Clark, B. B., & Klein, P. G. 2017. Uncertainty types and transitions in the entrepreneurial process. *Organization Science*, 28: 840-856.
- Packard, M. D., & Clark, B. B. 2020. On the mitigability of uncertainty and the choice between predictive and non-predictive strategy. *Academy of Management Review* (forthcoming).
- Patvardhan, S., & Ramachandran, J. 2020. Shaping the future: Strategy making as artificial evolution. *Organization Science* (forthcoming).
- Poletiek, F. H. 2001. *Hypothesis-testing behaviour*. Hove: Psychology Press.
- Pontikes, E. G., & Rindova, V. P. 2020. Shaping markets through temporal, constructive, and interactive agency. *Strategy Science*, 5: 149-159.
- Porac, J. F., & Tschang, F. T. 2013. Unbounding the managerial mind: It's time to abandon the image of managers as "small brains". *Journal of Management Inquiry*, 22: 250-254.
- Porac, J. F., Thomas, H., & Baden-Fuller, C. 1989. Competitive groups as cognitive communities: The case of Scottish knitwear manufacturers. *Journal of Management Studies*, 26: 397-414.
- Porter, M. E. 1985. *Competitive advantage: Creating and sustaining superior performance*. New York: Free Press.
- Ramoglou, S., & Tsang, E. W. K. 2016. A realist perspective of entrepreneurship: Opportunities as propensities. *Academy of Management Review*, 41: 410-434.
- Reymen, I. M., Andries, P., Berends, H., Mauer, R., Stephan, U., & Van Burg, E. 2015. Understanding dynamics of strategic decision making in venture creation: A process study of effectuation and causation. *Strategic Entrepreneurship Journal*, 9: 351-379.
- Rindova, V. P., & Courtney, H. 2020. To shape or adapt: Knowledge problems, epistemologies and strategic postures under Knightian uncertainty. *Academy of Management Review* (forthcoming).
- Runde, J. 1998. Clarifying Frank Knight's discussion of the meaning of risk and

- uncertainty. *Cambridge Journal of Economics*, 22: 539-546.
- Runde, J. 2009. Dissecting the Black Swan. *Critical Review*, 21: 491-505
- Russo, J., & Schoemaker, P. 1992. Managing overconfidence. *Sloan Management Review*, 33: 7-17.
- Sarasvathy, S. D., Ramesh, A., & Forster, W. R. 2014. The ordinary entrepreneur. In F. Welter & T. Baker (Eds.), *The Routledge Companion to Entrepreneurship*: 227-244. Abingdon: Routledge Press.
- Savage, L. J. 1954. *The foundations of statistics*. London: John Wiley & Sons.
- Schoemaker, P. J. H. 1993. Multiple scenario development: Its conceptual and behavioral foundation. *Strategic Management Journal*, 14: 192-213.
- Shackle, G. L. S. 1961. *Decision, order, and time in human affairs*. Cambridge: Cambridge University Press.
- Shackle, G. L. S. 1972. *Epistemics and economics: A critique of economic doctrines*. Cambridge: Cambridge University Press.
- Shackle, G. L. S. 1979. *Imagination and the nature of choice*. Edinburgh: Edinburgh University Press.
- Shafer, G. 1986. Savage revisited. *Statistical Science*, 1: 463-485.
- Shepherd, D. A., Haynie, J. M., & McMullen, J. S. 2012. Confirmatory search as a useful heuristic? Testing the veracity of entrepreneurial conjectures. *Journal of Business Venturing*, 27: 637-651.
- Shepherd, N. G., & Rudd, J. M. 2014. The influence of context on the strategic decision-making process: A review of the literature. *International Journal of Management Reviews*, 16: 340-364.

- Shepherd, D. A., Williams, T. A., & Patzelt, H. 2015. Thinking about entrepreneurial decision making: Review and research agenda. *Journal of Management*, 41: 11-46.
- Simon, H. A. 1997 [1947]. *Administrative behavior* (4th ed.). New York: Free Press.
- Simon, H. A. 1990. Invariants of human behavior. *Annual Reviews of Psychology*, 41: 1-19.
- Taleb, N. N. 2007. *The black swan: The impact of the highly improbable*. London: Penguin books.
- Todd, P. M., Gigerenzer, G., & the ABC Research Group. 2012. *Ecological rationality: Intelligence in the world*. New York: Oxford University Press.
- Townsend, D. M., Hunt, R. A., McMullen, J. S., & Sarasvathy, S. D. 2018. Uncertainty, knowledge problems, and entrepreneurial action. *Academy of Management Annals*, 12: 659-687.
- Tripsas, M., & Gavetti, G. 2000. Capabilities, cognition and inertia: Evidence from digital imaging. *Strategic Management Journal*, 21: 1147-1161.
- Tweney, R. D., Doherty, M. E., Worner, W. J., Pliske, D. B., Mynatt, C. R., Gross, K. A., & Arkkelin, D. L. 1980. Strategies of rule discovery in an inference task. *Quarterly Journal of Experimental Psychology*, 32: 109-123.
- Tukey, D. D. 1986. A philosophical and empirical analysis of subjects' modes of inquiry in Wason's 2-4-6 task. *Quarterly Journal of Experimental Psychology*, 38: 5-33.
- Vallée-Tourangeau, F., Austin, N. G., & Rankin, S. 1995. Inducing a rule in Wason's 2-4-6 task: A test of the information quantity and goal-complementarity hypotheses. *Quarterly Journal of Experimental Psychology*, 48A: 895-914.
- Wason, P. C. 1960. On the failure to eliminate hypotheses in a conceptual task. *Quarterly Journal of Experimental Psychology*, 12: 129-140.
- Weick, K. E., & Sutcliffe, K. M. 2007. *Managing the unexpected. Resilient performance in an age of uncertainty* (2nd ed.). San Francisco: John Wiley & Sons.

Wharton, C. M., Cheng, P. W., & Wickens, T. D. 1993. Hypothesis-testing strategies: Why two goals are better than one. *Quarterly Journal of Experimental Psychology*, 46: 743-758.

Zhang, S. & Cueto, J. 2017. The study of bias in entrepreneurship. *Entrepreneurship: Theory and Practice*, 41: 419-454.

TABLE 1

The Set of Imaginable Scenarios and the Scenario Space (adapted from Faulkner et al. 2017)

Scenario is at present...

		a genuine possibility	not a genuine possibility
Scenario is at present...	consciously imagined by the decision-maker and ...	1. Scenarios imagined and correctly regarded as possible	2. Scenarios imagined and incorrectly regarded as possible
	regarded as possible		
	regarded as impossible	3. Scenarios imagined and incorrectly regarded as impossible	4. Scenarios imagined and correctly regarded as impossible
not consciously imagined by the decision-maker		5. Scenarios not imagined	6. Scenarios not imagined

TABLE 2

Updating Beliefs: Movements within the Table

Scenario is...



		a genuine possibility	not a genuine possibility
		Scenario is...	consciously imagined by the decision-maker and ...
regarded as possible			
regarded as impossible	3. 		4. 
not consciously imagined by the decision-maker		5. 	6. 

TABLE 3

Updating the Scenario Space


		Scenario is...	
		a genuine possibility	not a genuine possibility
Scenario is...	consciously imagined by the decision-maker and ...	1.   	2.   
	regarded as possible		
	regarded as impossible	3.   	4.   
not consciously imagined by the decision-maker	5.   	6.   	
previously unimaginable			

TABLE 4
Disconfirmation

Activities	Steps	Illustration in Wason's original 2-4-6 task (where the correct rule is "three numbers in increasing order of magnitude" and the initial triple is 2-4-6)
Hypothesis formation	Establish working hypothesis consistent with the <i>existing</i> body of evidence and assume that this hypothesis is correct.	Establish working hypothesis consistent with the <i>existing</i> body of evidence (2-4-6), e.g. the rule is "three even numbers" and assume that this hypothesis is correct.
Evidence collection	Look for <i>new</i> evidence that is inconsistent with the working hypothesis.	Provide to the experimenter a previously untried triple that is inconsistent with the rule "three even numbers" e.g., 3-5-7.
Feedback evaluation	If no such new evidence is found, retain the working hypothesis and restart the procedure.	If the new triple does not conform to the experimenter's rule, retain the rule "three even numbers" and restart the procedure.
	If such new evidence is found, generate a <i>new</i> working hypothesis consistent with the now expanded body of evidence, and restart the procedure.	If the new triple conforms to the experimenter's rule, eliminate the rule "three even numbers", generate a new working hypothesis consistent with the now expanded body of evidence (2-4-6; 3-5-7), e.g., "three numbers increasing by 2", and restart the procedure.

TABLE 5
Counterfactual Reasoning

Activities	Steps	Illustration in Wason's original 2-4-6 task (where the correct rule is "three numbers in increasing order of magnitude" and the initial triple is 2-4-6)
Hypothesis formation	<p>Establish working hypothesis consistent with the <i>existing</i> body of evidence and assume that this hypothesis is incorrect.</p> <p>Generate an alternative hypothesis consistent with the existing body of evidence.</p>	<p>Establish working hypothesis consistent with the <i>existing</i> body of evidence (2-4-6), e.g. the rule is "three even numbers" and assume that this hypothesis is incorrect.</p> <p>Generate an alternative hypothesis consistent with the existing body of evidence, e.g., the rule is "three numbers increasing by 2".</p>
Evidence collection	Look for new evidence that is consistent with the alternative hypothesis and inconsistent with the working hypothesis.	Provide to the experimenter a previously untried triple consistent with the rule "three numbers increasing by 2" and inconsistent with the rule "three even numbers", e.g., 3-5-7.
Feedback evaluation	If no such new evidence is found, retain the working hypothesis and restart the procedure.	If the new triple does not conform to the experimenter's rule, retain the rule "three even numbers" and restart the procedure.
	If such new evidence is found, make the alternative hypothesis the new working hypothesis and restart the procedure.	If the new triple conforms to the experimenter's rule, eliminate the rule "three even numbers", make the rule "three numbers increasing by 2" the new working hypothesis, and restart the procedure.

TABLE 6
 Disconfirmation-based and Counterfactual Reasoning-based Procedures for Updating
 Scenario Spaces

Disconfirmation	Counterfactual Reasoning
<ol style="list-style-type: none"> 1. Assume that the existing scenario space is correct. 2. Select one of the scenarios S in the existing scenario space and look for new evidence <i>against</i> this scenario. 3. Analyse the bearing of this new evidence on all existing scenarios in the scenario space. 4. Whenever a scenario is eliminated from the scenario space, generate an alternative scenario S' not currently in the scenario space, consistent with the now expanded evidence, by imagining a new influence you had not considered before. <p>1.' Restart the procedure.</p>	<ol style="list-style-type: none"> 1. Assume that the existing scenario space is incorrect. 2. Generate an alternative scenario S' not currently in the scenario space, consistent with the existing evidence, by imagining a new influence you had not considered before. 3. Look for new evidence <i>in favour</i> of this new scenario S'. 4. Analyse the bearing of this new evidence on S' and all existing scenarios in the scenario space, adding S' to the scenario space and/or eliminating existing scenarios as justified by the evidence. <p>1.' Restart the procedure.</p>

Alberto Feduzi (a.feduzi@jbs.cam.ac.uk) is Senior Faculty in Management Practice at Cambridge Judge Business School and Director of Studies in Management Studies at Murray Edwards College, Cambridge. He holds a PhD from the University of Cambridge and his research focuses on judgment and decision making under uncertainty.

Philip Faulkner (pbf1000@cam.ac.uk) is Reddaway Fellow, Director of Studies and Senior College Teaching Officer in Economics, at Clare College, Cambridge, and a Fellow of Cambridge Judge Business School. A co-editor of the *Cambridge Journal of Economics*, his research focuses on the nature of technology and decision making under extreme uncertainty.

Jochen Runde (j.runde@jbs.cam.ac.uk) is Professor of Economics & Organisation at Cambridge Judge Business School and Professorial Fellow at Girton College, Cambridge. He is co-editor of the *Cambridge Journal of Economics* and his main areas of research are social ontology and the ontology of technology, decision making under uncertainty, and explanation in the social sciences.

Laure Cabantous (laure.cabantous.1@city.ac.uk) is Professor of Strategy and Organization at the Business School (formerly Cass), City, University of London. She studies decision making and leadership practices in organizations using a performative perspective. She also has an interest for embodied ways of knowing and modes of reasoning in organizational settings.

Christoph H. Loch (c.loch@jbs.cam.ac.uk) is Professor in the Technology and Operations Group, and the Dean of the Cambridge Judge Business School. His research examines innovation processes, project management, strategy execution, and the motivation of professional employees.