City Research Online

City, University of London Institutional Repository

RUNNING HEAD: MODIFIER EFFECT

The modifier effect in within-category induction:

Default inheritance in complex noun phrases.

Martin L. Jönsson   James A. Hampton

Lund University, Sweden   City University, London UK

Address for correspondence:

Professor James A. Hampton

Psychology Department

City University

Northampton Square

London EC1V OHB, England.

hampton@city.ac.uk

Word count approx 9326, not including Title, Abstract, References, Appendix or Tables.

Abstract

Within-category induction is the projection of a generic property from a class (*Apples are sweet*) to a subtype of that class (*Chinese apples are sweet*). The modifier effect refers to the discovery reported by Connolly et al. (Connolly, Fodor, Gleitman & Gleitman, 2007, "Why stereotypes don't even make good defaults", *Cognition, 103,* 1-22), that the subtype statement tends to be judged less likely to be true than the original unmodified sentence. The effect was replicated and shown to be moderated by the typicality of the modifier (Experiment 1). Likelihood judgments were also found to correlate between modified and unmodified versions of sentences. Experiment 2 elicited justifications which suggested three types of reason for the effect – pragmatics, knowledge-based reasoning, and uncertainty about attribute inheritance. It is argued that the results provide clear evidence for the default inheritance of prototypical attributes in modified concepts, although a full account of the effect remains to be given.

One of the most remarkable aspects of human language is that speakers can communicate successfully using familiar words in novel combinations. An important proposal about how this creativity is possible is that people derive the meaning of the novel combinations from the meanings of the words that constitute them and the syntactic structure in which they are embedded. Asked to interpret a relatively novel phrase such as "blue apple", people retrieve the meanings of "blue" and "apple", and use the syntax of adjective-noun phrases to provide a semantic interpretation. Explanations such as this, in terms of the *compositionality* of natural languages, have had a fundamental impact on philosophical research into meaning and concepts during the 20[th] century, from seminal work such as Frege (1923) and Davidson (1967) through to recent papers by Fodor and Lepore (2002) (see Jönsson, 2008, for a book length critical examination of this tradition). Psychological research into how meanings and concepts are mentally represented has similarly become increasingly interested in producing theories that address the problem of compositionality (see Murphy, 2002, for an overview of psychological work, and Machery, Hinzen and Werning, 2011, for a collection of contributions to the philosophical debate including Hampton & Jönsson, 2011).

A critical question for the problem of compositionality in novel noun phrases is the extent to which the complex concept corresponding to the phrase "inherits" the default prototypical attributes of the head noun concept. To what extent should a blue apple be considered to have the other common attributes of apples? Prototypical attributes are those that are *generally* true of the members of a conceptual category, although they may not be universally true (Krifka et al., 1995; Pelletier, 2009). People happily assent to statements such as "birds fly" and "apples are sweet", even though counterexamples exist (such as penguins, or sour apples). In fact general truth conditions for these kinds of generic statements have proven surprisingly difficult to elaborate in any precise manner. There appears to be no simple quantificational rule for determining which generic sentences are considered true and which

false. (For example "birds lay eggs" is considered true while "birds are female" is not, although there are more female than egg-laying birds in the population). For our purposes, and without going into the complexities of different accounts of generics, it is assumed that an important class of these sentences express prototypical attributes of the concepts in question, of the kind that Rosch and Mervis (1975) identified as constituting the content of prototype concepts (see also Hampton, 1979, 2006; Khemlani, Leslie & Glucksberg, 2009).[1]

Given that an attribute is not universally true of a concept (for example, many birds do not fly), how should one determine whether the predicate should also be considered generically true of the complex concept formed when an adjectival or nominal modifier is applied to the noun. In simple terms, if birds typically fly should one therefore conclude that pet birds typically fly, or that Tibetan mountain birds do? The question concerns a form of inductive reasoning in which properties of the general category are projected to subsets of the category.

Two theoretical positions can be contrasted. On the one hand, according to simple models of prototype modification, such as the selective modification model (Smith et al., 1988) *all* of the prototypical attributes of the noun concept are inherited unchanged by the modified concept with the exception of the specific dimension to which the modifier applies. Thus the concept BLUE APPLE would have all the normal default attributes of apples (roundness, peel, seeds, taste) but its color would now be represented as "blue" rather than red, green, brown or yellow. In addition the weight of COLOR in determining typicality of exemplars would be increased. This model therefore predicts that a modified concept should inherit all of the prototypical attributes of its unmodified parent.

Two important exceptions to this modification process however need to be specified. Default attribute inheritance only applies, according to Smith et al. (1988), if the noun phrase is "novel" and the modifier is "compatible". If the modifier-noun phrase names a familiar

type of object, such as "cooking apple", then a concept representation for the phrase may already exist, incorporating known attributes of the items in question (such as that they are larger and less sweet than regular apples, at least in the UK). Similarly, if the modifier were not "compatible" it could trigger a wider range of conceptual modifications. Hence "glass apple" would not only have the value GLASS for the MADE OF dimension, but would also consequently have different values across a range of other dimensions – skin and taste for example. Being made of glass is not compatible with having a peel, or being crunchy. In sum, as long as the modifier noun phrase is novel and the modifier compatible, then simple models of prototype combination suggest that the complex concept should inherit all the prototypical attributes of the noun by default.

The alternative position to the problem of attribute inheritance is one advocated by Fodor (1998) (see also Connolly et al., 2007). Motivated (in part) by the impossibility of providing clear general rules for attribute inheritance in prototype concept combination, Fodor argued that concepts should not be identified with prototypes. We will not attempt to do full justice to Fodor's arguments here, but the gist of his claim is that concepts are represented in the mind as atomic mental particulars that represent classes in the world. The class of apples is associated with the word "apple" whose meaning is represented by a symbol in the mind, and similarly for the concept BLUE. The meaning of the phrase "blue apple" is obtained by forming the composite concept BLUE APPLE, which is just a matter of recruiting the meanings of the two terms and using the syntax of English to compose the meaning of the phrase. The meaning of the noun phrase will correspond to a complex concept incorporating an appropriate relation between APPLE and BLUE, which in this case would be one of set intersection (the set of things that are both apples and blue). All that one can know for sure about blue apples is that they are blue and that they are apples (and, possibly, anything that would logically follow from these two facts, such as that they are not red and not pears).

According to Fodor (1998) this is the only way in which the compositionality of concepts can be successfully incorporated into a representational theory of mind, and compositionality of concepts, he argues, is of critical importance to the explanation of central attributes of the mind such as, for instance, the systematicity and productivity of thought.

Given this very sparse "classical" theory of concept combination, the evidence that has been collected to date on the prototypes of complex concepts (Hampton, 1987; 1988, Murphy, 1988) has to be accounted for by knowledge-based processes operating *after* combination has taken place.

A closer look at some of the empirical evidence on attribute inheritance in concept combinations (e.g. Hampton, 1987, 1991, 1997a, 1997b) suggests that neither position is fully satisfactory. One could argue (with Fodor) that parsimony requires that accounts of conceptual combination should be kept simple and compositional, and that the complexity of attribute inheritance simply reflects the notorious problem of background knowledge that affects all theories of conceptual thought. However this "solves" the problem of concept combination by simply moving the problem to a later stage in processing. It is indisputable that the psychological process of interpreting a complex phrase has to start with just the two words, their stored meanings and the syntax of the phrase. That is the raw input to the process. But to claim that there is a clean compositional process, followed by a messy interpretation process, does not address the interesting issues, and moreover is less parsimonious that one that omits the first stage, and allows an interaction with stored knowledge and discourse context from the beginning. The approach also gives up too easily on the attempt to understand the degree of systematicity that can be discovered in the process of attribute inheritance and prototype composition. Some attributes of complex concepts are clearly the same as those of their elementary constituents, as when weapons that are tools are judged to be dangerous, or when sports that are games are judged to require physical exertion

(Hampton, 1987, see also Hampton, 1997b). Here, default inheritance of attributes seems to be the rule. On the other hand, contrary to simple models of prototype combination, other attributes may not be inherited by default, as is clearly the case with the knowledge-based and compatibility effects described above. Pet birds do not migrate south for the winter as other birds do, and nor are pet fish warm and cuddly as other pets typically are (Hampton, 1987; 1997a).

A more fruitful approach may therefore be to explore the extent to which prototypical attributes are inherited and, where inheritance is not complete, to identify possible explanations. In this paper, we approach the problem of default attribute inheritance with an open mind. We explore the degree to which attributes may fail to be inherited, and elicit explanations for such an effect. To the extent that systematic patterns may be found in the data, we will argue that there is a place for constructing constrained models of how prototype representations change when combined with modifiers, rather than considering the whole issue to be insoluble without a full representation of world knowledge.

The starting point for our investigation was a discovery reported by Connolly et al. (2007). They reported an effect, which we have termed the "modifier effect" which speaks directly to the question of whether the attributes of a noun prototype are inherited by a complex noun phrase in the absence of knowledge or compatibility effects. They showed that people rated bare plural generic sentences of the form "Ns are P", (where N is a noun and P is a property typically true of N), as more likely to be true than the corresponding sentence "MNs are P", where the noun N is modified by a modifier M. So, for instance, people rated sentences such as "*Penguins live in cold climates*" as being more likely to be true than modified sentences such as "*Solitary penguins live in cold climates*".[2] Sentences were assigned systematically decreasing likelihood ratings as the head noun was modified by a single typical modifier (e.g. *flightless penguins*), by a single atypical modifier (*solitary*

penguins) or by two atypical modifiers (*solitary migrant penguins*). Table 1 (see later) shows the mean judgments for the four conditions from their study on a 10 point likelihood rating scale. All means differed significantly from each other ($p < .05$ by t-tests on participant and item means). The "modifier effect" refers to this reduction in likelihood for an attribute when the head noun is modified. We will refer to the task of inferring the likelihood of an attribute in a subset of a category on the basis of its likelihood in the category as a whole as "within-category induction".

In reporting the effect, Connolly et al. (2007) noted that the result is not explicitly addressed by any of the available models of prototype combination. For example, as described above, Smith et al.'s (1988) selective modification model assumes that the modifier *selectively* modifies the relevant dimension alone. The reduced belief in the likelihood of other attributes is not therefore explained by the model. (We return to discuss the correctness of this assertion in the final section of the paper). According to the model, the concept of BLUE APPLE should inherit the default attributes of apples in all respects other than color, so "Blue apples are round" should be considered just as true as "Apples are round". We will follow Connolly et al. in referring to this assumed process as the Default to Stereotype, or DS, strategy. As long as there is no incompatibility between the modifier "blue" and the attributes being inherited (for example, roundness and peel), and no knowledge or memory of blue apples in the world that a person could call on (the combination is novel), then according to DS the attributes should be judged to be just as true of blue apples as they are of apples. This prediction is what Connolly et al. claimed to have disproved with their result.

The modifier effect clearly has important consequences for theories of prototype combination (as well as opening up further avenues of research, see Jönsson & Hampton, 2006). In the studies to be reported here we aimed to explore the phenomenon in some detail,

with a view to understanding its basis. Experiment 1 was a close partial replication of

Connolly et al.'s (2007) study, aimed at reproducing the effect. Correlational analyses were

additionally applied to the data to test a particular prediction of DS. If a modified concept

inherits by default the attributes of the simple concept prototype, but with some overall

reduction in perceived likelihood, then the relative differences in likelihood of properties for

modified concepts should reflect the differences in likelihood for unmodified concepts to

some significant degree. There should therefore be a positive correlation between the

likelihood of the attribute for the modified and the unmodified concepts. So if a given noun-

property combination was rated as more likely than others, it should still tend to be rated

above average when the subject noun is modified.

Experiment 2 then assessed the degree to which people might adopt the DS strategy,

when asked not to rate individual attribute likelihood but instead to state directly whether a

pair of N and MN sentences are equally likely or not. Experiment 2 also asked people to

provide justifications of their choices when one sentence was judged more likely than the

other, so that different explanations of the modifier effect could be considered.

<div align="center">Experiment 1</div>

Experiment 1 aimed to replicate Connolly et al.'s demonstration of the modifier effect. In

addition we tested a prediction of the DS strategy hypothesis. Connolly et al. (2007)

discussed a possible account of their result according to which the effect of an atypical

modifier is to decrease confidence in the truth of all attributes equally. If all prototypical

attributes are inherited by default, but there is a general reduction in confidence induced by

the novelty of the new noun phrase, then it should be the case that the relative ordering of

likelihood of sentences should remain the same (broadly speaking) when the head nouns are

modified. High likelihood attributes of a prototype should still be high likelihood for the

modified concept, while lower likelihood attributes should be inherited with lower

likelihood. The correlation should not be perfect since (in addition to measurement error) different modifiers may have different sizes of effect, but we nonetheless predicted that there should be a significant positive correlation for likelihood judgments for the unmodified and modified versions of each sentence.

Method

   Participants. Twenty-nine undergraduates at City University, London participated in exchange for course credit.

   Materials. Each booklet consisted of 40 target and 90 filler sentences. All were simple declarative sentences, consisting of a noun phrase and a predicate. Four versions of each target sentence were constructed as follows so that a total of 160 sentences were created. These were the same 160 target sentences used by Connolly et al. (2007). The head noun could either be (a) unmodified, (b) modified by a *typical* modifier, generally true of the class (according to frequency in Cree and McRae's (2003) feature production norms), (c) modified by an *atypical* modifier, not typically true of the class, or (d) modified by *two atypical* modifiers, one of which was that used in condition (c). For instance, for the head noun "ducks" and the predicate "have webbed feet", the following 4 sentences were constructed:

   a) Ducks have webbed feet.

   b) Quacking ducks have webbed feet.

   c) Baby ducks have webbed feet.

   d) Baby Peruvian ducks have webbed feet.

The attributes used (e.g. have webbed feet) were all typically true of the head noun concept. The atypical modifiers in conditions (c) and (d) were chosen by Connolly et al. such that (1) they did not appear in the Cree and McRae feature norms, (2) they would form relatively novel unfamiliar phrases and (3) (to use our term) they were compatible with the predicates. As they put it, the modifiers were "commensurable with the predicates, i.e., the modifier and

the predicate should be orthogonal and independent, so that the introduction of the modifier does not necessitate a change in the applicability of the predicate". A list of all materials can be seen in the Appendix.

The target sentences were rotated across four booklets so that each booklet contained 10 target sentences with different head nouns for each of the 4 conditions. The 40 target sentences in each booklet were embedded randomly in 75 filler sentences so that sentences containing the same head noun appeared in the same position in each booklet. These fillers were included to encourage the use of the full range of the scale, and to disguise the design. An additional 15 filler sentences appeared at the front of each list to avoid warm-up effects. (Connolly et al.'s 120 fillers were reduced to 90 to reduce the load on participants). Since the original filler sentences were unavailable, new ones were generated. They were similar to the target sentences in structure. There were 30 filler sentences for each of 3 levels of likely truth: (a) unlikely (*The Roman coliseum is in Moscow*); (b) intermediate likelihood (*Philadelphia is larger than Atlanta*); and (c) highly likely (*Bicycles have two wheels*).

Design and Procedure. Participants were randomly divided into 4 groups, each receiving one of the four booklets. They indicated the likely truth of each sentence using the numbers 1 through 10 appearing to the right of each sentence (1 = very unlikely; 10 = very likely). The task took from 20 to 45 minutes.

Results

Mean likelihood ratings. Table 1 shows that the results were very similar to those obtained by Connolly et al. (2007). ANOVA was run by participants (F_1) and items (F_2), with modifier condition as a within-subjects/within-items factor with 4 levels. Min F' was used for tests of significance in all reported analyses. The effect of modifier condition on rated likelihood was significant (*Min F'*(3, 198) = 17.83, *p* < .001). Effect size η^2 was .64 for subjects and .42 for items. Post-hoc pairwise comparisons between conditions were all

significant ($p < .001$) except for that between atypical and double atypical modifier conditions. Modification of the subject noun did therefore reduce judged likelihood of sentences, and this effect was greater with atypical than with typical modifiers.

Correlations. Having replicated the pattern of means from the earlier study, an analysis tested the hypothesis that, although rated likelihood is reduced by a modifier, the *relative* rated likelihood of different attributes is still maintained to some significant degree. If people are using the DS strategy, then the rated likelihood of attributes for the modified concepts should reflect their likelihood for the unmodified prototypes, resulting in a significant positive correlation across items between the modified and unmodified forms of the sentences. To test this hypothesis, correlations were calculated between mean likelihood ratings for modified and unmodified versions of each of the 40 concepts for each of the three levels of modifier typicality. (Estimated pooled split-half reliability of the scales was 0.6). For typical modifiers, the correlation was positive but not significant ($r(38) = .262$, $p = .10$, while for atypical and double atypical modifiers, the correlations were both significantly positive ($r(38) = .406$ and $.422$, respectively, both $p < .01$). There was therefore evidence that, at least with atypical modifiers, the likelihood of an attribute for the noun prototype was predictive of its strength for the modified noun phrase concepts.

Discussion

Experiment 1 found the modifier effect to be replicable and robust. We also confirmed that the effect was greater for atypical than for typical modifiers. In addition, there was evidence that the relative strengths of attributes for the head noun concepts were inherited to some significant degree by the modified noun concepts. This pattern of correlation is consistent with the people using the DS strategy. If, for example, there had been no evidence for a positive correlation, the argument against the use of the DS strategy would have been strengthened. In order to explore the basis of the modifier effect further, in Experiment 2

participants were asked first to provide a comparative judgment of whether the two sentences (N and MN) were equally likely to be true or not, and were then asked to explain their judgments for those cases where they said they were not.

Experiment 2

The second experiment had two aims. The first was to test whether the perceived reduction in likelihood would still be found if participants made a direct comparison between a modified and an unmodified sentence. With the design used in Experiment 1 (taken from Connolly et al., 2007), the same participant never judged both modified and unmodified versions of the same sentence. A stronger test of the modifier effect is to set the two sentences side-by-side and ask people to judge whether or not one was more likely, and if so which.[3] This design has the advantage of drawing to the participants' attention the explicit possibility that the two sentences are equally likely (as would be predicted by simple DS models). The design also had the advantage of allowing us to fulfill our second aim, which was to explore possible reasons for the modifier effect by asking participants to justify their likelihood judgments. It would be impractical to ask people to justify the single rating judgments made in Experiment 1 ("Why did you choose a 7 for the statement that ravens are black?"), but it was perfectly sensible to ask for a justification of a preference, as in "Why did you think that ravens being black was more/less likely than jungle ravens being black?" Filler sentences were again used to reduce response bias. So that having to give the justifications did not influence the decision, participants made all their judgments first, and were then unexpectedly asked to revisit them to provide justifications.

Method

Participants. Forty undergraduates at City University, London participated either in exchange for course credit or for a small payment.

Materials. Each booklet contained 42 target and 58 filler sentence pairs. All sentence

pairs consisted of a sentence "N are P" together with a second sentence "MN are P" in which

M was one of the 3 possible modifiers used in Experiment 1, i.e. typical, atypical, and double

atypical. For instance, for the head noun "ducks" and the predicate "have webbed feet", the

following three pairs of sentences were constructed:

A)   Ducks have webbed feet.   Quacking ducks have webbed feet.

B)   Ducks have webbed feet.   Baby ducks have webbed feet.

C)   Ducks have webbed feet.   Baby Peruvian ducks have webbed feet.

Two new triples of sentence pairs were added to the 40 triples generated from the materials

from Experiment 1 (see Appendix). The resulting 42 triples were distributed evenly over the

3 booklets. Target sentence pairs were counterbalanced across the 3 booklets, and randomly

embedded in 50 filler sentence pairs so that target pairs containing the same head noun

always appeared in the same ordinal position. In addition, target pairs were always placed in

odd-numbered positions. A further 8 fillers appeared at the front of each list to avoid warm-

up effects. Filler pairs were similar to target pairs in both structure and plausibility. Half

contained one modifier and half contained two. Knowledge effects were used so that about

one third had the modified sentence as less likely ("Radiators are warm" versus

"Malfunctioning white radiators are warm"), one third had the modified sentence as more

likely ("Doors are made of metal" and "Prison doors are made of metal"), and one third had

the two sentences possibly equally likely ("Parrots are noisy" and "Green feathered parrots

are noisy"). The fillers thus ensured that there were opportunities to use all three response

options.

   Design and Procedure. Participants were randomly assigned to one of the 3 booklet

conditions. The first and last page of each booklet contained instructions, and each page in

between contained 4 sentence pairs. Three response options, one of which was to be circled,

appeared to the right of each sentence pair; 1) "the first sentence is more likely to be true", 2)

"the second sentence is more likely to be true", and 3) "the two sentences are equally likely to be true". At the end, participants were instructed to go back and justify why they answered in the way that they did, by writing a short statement next to the item. To reduce the load, they only had to justify odd numbered sentences where they had stated that one of the sentences was more likely than the other (even numbered positions always contained fillers). A decision of "equally likely" was taken to be a default judgment, not requiring any further justification. Participants proceeded at their own pace, finishing in 20 to 40 minutes.

Results and Discussion

   Frequencies. The mean proportion of answers to target pairs favoring each response option by condition is shown in Table 2. Interestingly, the "equally likely" option was the most commonly chosen for all three kinds of sentence pairs (between 60% and 69% of responses). When not equally likely, the unmodified sentence was selected as more likely 74% of the time for a typical modifier, 95% of the time for a single atypical modifier and 93% for a double atypical modifier, thus replicating the modifier effect and its interaction with modifier typicality. ANOVA was conducted on the proportions of responses given in each condition.

   Since the 3 response proportions summed to 1, the data were not independent, as required for ANOVA. A greater proportion of one response would tend to correspond to a lower proportion of the other. To break the dependence, the "equally likely" responses were omitted from the analysis. Since the majority of responses were "equally likely", the other two responses were largely independent. Correlations between the proportions of the two responses at the three levels of typicality were not significantly different from zero for the participant data, and showed only weak negative correlations for the item data, with r between -.20 and -.37. In other words, the two remaining proportions were largely independent of each other.

A two-way ANOVA was therefore run on response frequencies with factors of condition (3 levels of modifier), and response (2 levels: selecting the unmodified vs. selecting the modified). All three effects were significant. The significant main effect of condition (*Min F'*(2, 153) = 3.74, *p* < .05) corresponded to the fact that there were fewer preferences (in either direction) expressed in the typical modifier condition than in the other two. In other words the "equally likely" response was significantly more frequent (69%) in the typical modifier condition than in the others (60-62%). The significant main effect of response (*Min F'*(1,55) = 41.0, *p* < .001) confirmed that when a preference was expressed it was mainly for unmodified sentences (32%) rather than modified sentences (4%), and the significant interaction (*Min F'*(2, 158) = 8.73, *p* < .001) confirmed that the tendency to prefer the unmodified sentence was greater in the two atypical modifier conditions (a difference of 34%) than in the typical modifier condition (a difference of 15%). There was again no significant difference between the Atypical and the Double Atypical conditions.

As a double check on the ANOVA, one-way ANOVA were also run to test the effect of modifier typicality on each of the response proportions separately, and these confirmed the pattern of significant results. Atypical and double atypical modifiers did not differ, but compared to them typically modified sentences had significantly fewer responses where the unmodified sentence was selected, and significantly more responses where the modified sentence was selected. Typically modified sentences also had significantly more "equally likely" responses than the other two. All three main effects and all pairwise contrasts comparing typical with atypical and typical with double atypical conditions were significant (p < .05) on both participants and items analyses.

Justifications. Justifications were provided for 87% of responses. They were transcribed and classified by the first author and a second independent judge. (Any given justification could be classified in more than one class.) Frequency by condition is shown in Table 3, as a

function of whether the unmodified or the modified sentence was chosen as more likely.

Where the unmodified N sentence was chosen as more likely (the usual modifier effect), the

following justifications were given.

Pragmatic (45%). N was preferred as more general, while M (in the combination MN) was

considered redundant. Examples:

1) Flightless penguins live in cold climates

"All penguins live in cold climates and all penguins are flightless so to make a distinction

is arbitrary, just say penguins live in cold climates".

2) Paleolithic European ostriches had long necks

"All ostriches have long necks independent of where it (sic) *comes from otherwise it's*

not an ostrich"

Knowledge (19%). Either a knowledge-based inference or experience with individuals in the

modified noun category led people to doubt the truth of the MN sentence. Examples:

3) Edible catfish have whiskers

"Edible catfish probably do not have whiskers still attached, as they could not be eaten

like this",

4) Machine washable acrylic parkas are warm

"Machine washable acrylic is not very warm."

Uncertainty (14%). Doubt was expressed about the modified sentence. Examples:

5) Storage shacks are made of wood

"Shacks tend to be made of wood but storage shacks may not be",

6) Brazilian doves are white

"Generally doves are white, Brazil could have black doves"

Reversal (3%). The focus of the sentence was changed into a question about the likelihood of

M rather than the likelihood of P. Examples:

7) Antique sinks are found in kitchens

"The chances of finding an antique sink in a kitchen is (sic) *considerably lower than finding a non-antique sink in a kitchen",*

8) Uncomfortable handmade sofas are found in living rooms

"While some uncomfortable handmade sofas may be found, one is more likely to find a comfortable sofa in a living room".

Other (19%). The remaining justifications could not be easily classified (they were often of low relevance).

Pragmatic justifications, on the face of it, actually provided a reason for selecting the "equally likely" response – both flightless penguins and penguins in general live in cold climates. In fact, on several occasions, participants added "so I could also have said they were equally likely". Participants therefore chose the unmodified sentence on the grounds of relevance or informativeness. Pragmatic justifications were particularly common for typically modified pairs (58% of all justifications given), but were also very common for the other two conditions (about 40% of justifications given). They were evenly distributed across items.

Knowledge-based justifications were the second most frequent kind (19%) and indicated either that some of the materials were not sufficiently novel as conceptual combinations, or that people had chosen to draw inferences from broader background knowledge. They were more common in the conditions using atypically modified sentences. They were found across most items, but some items clearly had more than others – for example 8 people mentioned that bitter nectarines were likely to be unripe, and unripe fruit tended to be less juicy.

Justifications based on uncertainty were of particular theoretical interest in that they could be interpreted as implying that people were avoiding using the DS strategy, at least to the point where they had doubts about which stereotypical attributes of the concept would be inherited when it was modified. There were some 18-19% of these justifications in the two

atypical conditions, and hardly any in the typical condition. They were distributed evenly across most items. The two items most frequently justified this way were "South Sea pearls are white" and "Brazilian doves are white".

The final kind of justification, involving reversals, was specific to certain items, but is of some theoretical interest. Mostly in cases where the attribute referred to a location, people answered in a way that showed that they saw the judgment as being about the likelihood of the modifier rather than the likelihood of the attribute. "Uncomfortable sofas are found in living rooms" was considered unlikely, not because most uncomfortable sofas are likely to be found elsewhere, but because most sofas found in living rooms are likely to be comfortable. In effect the statement judged was therefore "Sofas found in living rooms are uncomfortable". This reversal of emphasis has also been noted in generics more generally (Krifka et al., 1995). For example the generic sentence "mosquitoes carry malaria" is true because most cases of malaria have mosquitoes as the vector, but not because most mosquitoes are vectors (in fact very few are). The sentence thus arguably expresses a generic attribute of malaria rather than of mosquitoes.

As can be seen in the lower half of Table 3, there were relatively few responses where the unmodified sentence was considered *less likely*. Most occurred in the Typical modifier condition, and the majority of these cases were justified using knowledge-based inferences of one of two types. First there were explanations based on experience – pet hamsters are more likely to live in cages than hamsters in general, or kitchen refrigerators are more likely to store food than refrigerators in general. The second type involved believing that the unmodified sentence was not completely true, so that the modified sentence was preferred, since it at least stood a chance of being true. For example *"Not all doves are white but I don't know what flightless Brazilian doves look like but since these are a specific kind of dove, the Brazilian ones could be all white."* In this case the Uncertainty argument was being used to

argue for the reverse of the modifier effect.

There were also a few Pragmatic justifications for selecting the *modified* sentences. These probably reflect participants mistaking which response they had given when returning to offer justifications, as they make little sense otherwise.

Discussion

As in Experiment 1 there was a modifier effect which was greater for atypical modifiers. (Neither of the experiments found any difference between the single and double atypical modifier conditions. Hence the discussion concentrates on the differences between typical and atypical (single or double) modified sentences.) A striking result was that when participants directly compared the relative likelihood of N and MN sentences, over 60% of the time they judged them equally likely – even when the modification involved either one or two atypical modifiers. It is therefore by no means automatic that modifiers will affect the likelihood of attributes. In fact on most occasions people considered the attribute to be unchanged in its likelihood, consistent with use of the DS strategy. (Remember that at this stage they were unaware that they would be asked to provide justifications, and the fillers provided ample opportunity to use all three response options). When people did differentiate the likelihood of the two sentences however, they almost always said that the modified sentence had a lower likelihood, thus explaining the small but significant drop in the mean ratings in Experiment 1 when a modifier is applied. (It is possible also that the 10 point rating scale used in that Experiment encouraged a finer discrimination of relative likelihood than the 3 response options used here – see footnote 3).

The most common justification for selecting the unmodified sentence as more likely was on the basis of the pragmatic implications of uttering each statement, in line with Grice's (1975) maxims of cooperative communication. To utter the modified statement when one knew that the more general one was also true would be to violate Grice's maxim of quantity

("Be as informative as you can"). Participants were very sensitive to this kind of consideration when they judged the sentences for relative likelihood. Although this Gricean effect may have been amplified by the contrastive task used in Experiment 2, similar effects could be expected when judging absolute likelihood in Experiment 1. What was unexpected here was that the pragmatic explanation was offered for both typically and atypically modified sentences. One might suppose for example that "quacking ducks have webbed feet" would be more infelicitous than "baby Peruvian ducks have webbed feet". The latter, after all, may serve to inform those who know little about baby ducks, Peruvian or otherwise.

For the typically modified sentences, the pragmatic justifications more or less exhausted the cases showing the modifier effect. There was no evidence that attributes of the modified concept differed substantially from those of the unmodified concept when the modifier was typical. For typically modified sentences, there were also knowledge-based justifications which more often than not justified an *increase* (29 cases) rather than a decrease (8 cases) in likelihood of the modified sentence (see Table 3).

For the atypically modified sentences, two additional types of justification were given. First, even though the modifiers were chosen to be compatible with the attributes, around 20% of the justifications showed that participants had thought of ways in which the modifier might in fact affect the attribute. (The search for coherence in conceptual combination may often lead to the construction of more elaborate scenarios – see Johnson & Keil, 2000; Kunda, Miller & Clare, 1990). Second, roughly 20% of justifications referred to uncertainty regarding the attributes of the modified concept. This latter justification is consistent with the notion of a general reduction in judged likelihood applying to unfamiliar novel noun phrases.

General Discussion

What are the implications of the modifier effect for theories of prototype combination and compositionality? In the introduction we described two contrasting positions concerning the

degree to which people may or may not use default to stereotype (DS) as a strategy for interpreting modified noun phrases. Connolly et al. (2007) argued that the modifier effect was evidence that people do not take the prototypical attributes of the head noun concept as a default for the complex concept. This claim is clearly at odds with the findings presented here. Experiment 2 demonstrated that the most common strategy for making within-category inductions was to judge an attribute to be equally true of the modified and unmodified concepts – exactly the DS strategy. Some 60% of judgments were to this effect, even when the modifier was atypical. This figure would be even higher if cases were excluded where either pragmatic or knowledge-based reasons were given for judging the modified concept sentence less likely.

In the remaining judgments where the modifier did reduce attribute likelihood, it can still be argued that the prototype for the concept had a strong influence. The typicality of the modifier moderated the degree to which likelihood judgments were reduced. But note that whether or not a modifier is typical or atypical can only be determined by knowledge of the prototype, (Fodor's atomic concept contains no information about what is typical of the concept, being simply a symbol representing the class). In addition Experiment 1 showed that the likelihood of an attribute being true of an atypically modified concept was correlated with the degree to which it was considered true of the concept alone.

Systematic patterns of attribute inheritance have been reported elsewhere. For example, Hampton (1987) demonstrated that in explicit conjunctions formed from restrictive relative clause constructions (birds that are also pets, sports that are also games), the judged importance of attributes for the conjunctive phrase was broadly predictable from their importance for each of the concepts separately. Our results here generalize this notion of importance to judgments of attribute likelihood.

Although rejecting Connolly et al.'s claim regarding the DS strategy, we should however

acknowledge that the position that they advocate is entirely consistent with our results. They proposed that, after a classical intersective concept combination stage, further inferential and pragmatics-based processes may then operate in order to arrive at a likelihood rating for each attribute. All of the justifications offered by our participants are consistent with this type of second stage process. While we argue that prototypes are inherited, it is also true that the DS model that Connolly et al. (2007) criticized – one of simple default inheritance of all prototypical attributes with unchanged likelihood – is incompatible with the effects reported. We found clear evidence for the modifier effect in both experiments, inconsistent with a simple model in which attributes are always inherited by complex concepts with their likelihood unchanged. (We have also obtained the effect in several other studies not reported here, e.g. Hampton, 2009; Hampton, Jönsson & Passanisi, 2009; Hampton, Passanisi & Jönsson, 2010). The best theoretical account of the data is probably therefore one in which attributes are inherited by default, but other factors can come into play that, overall, tend to reduce their judged likelihood.

Experiment 2 shed some light on what these factors might be. By asking participants for justifications of a judgment of relative likelihood, their attention was drawn to the intuitions that they had been using in selecting the unmodified sentence as more likely. Three main types of reason were given. One reason which applied across all modifier conditions related to the pragmatics of the sentences. Experiment 2 used a comparative truth judgment, and our participants preferred the more simply expressed sentence (doves are white) over the sentence with an apparently irrelevant modifier in it (Brazilian doves are white). Connolly et al. (2007) also noted the likelihood that Gricean effects were at work in the task, since there appeared to be no other way to explain the occurrence of the modifier effect with *typical* modifiers. We were able to confirm that almost all of the effect for typical modifiers could be attributed to pragmatics, while also showing that a similar number of pragmatic explanations

applied to the atypical and double atypical cases as well.

A second reason affecting atypical modifier sentences related to unanticipated effects of background knowledge, a well known influence on conceptual combination (Hampton, 1987; Murphy, 1988). Although the modifiers and attributes were chosen to be compatible, on a number of occasions participants were able to trace some causal link between the two, thus providing a reasoned justification for reducing (or in some cases increasing) the likelihood of the attribute for the modified concept.

The third reason given in Experiment 2 was uncertainty about the unfamiliar subset. If a modifier is atypical, then it is possible that the person will know little or nothing about the subset, and so they will have lower confidence in making an inductive judgment about its likely attributes. Never having heard before of "jungle ravens" or "albino crocodiles", participants had less confidence in these creatures having the attributes typical of regular ravens or crocodiles. This last reason is consistent with the claim (Connolly et al., 2007) that DS is likely to be a poor strategy for predicting the attributes of modified concepts, simply because the only thing one can know with confidence about a Lithuanian strawberry is that it has some connection with Lithuania and is a strawberry. The rest is guesswork. Less consistent with Connolly et al.'s claim, this justification may also reflect the suspicion that compositionality may in fact *not* apply, particularly in the case of folk names for plants. Poison oak is not an oak, and poison ivy not an ivy.

Our data suggest that this type of reasoning is relatively rare. Participants in Experiment 2 showed the modifier effect with atypical modifiers on less than 40% of trials, and when they did judge the modified sentence less likely, they used the uncertainty justification on less than 20% of those occasions, equivalent to less than 7% of the overall decisions. However Hampton, Passanisi and Jönsson (2010) did find further evidence for this kind of reasoning, when they showed that, in the absence of a supportive context, the modifier effect was just as

strong for categorical predicates ("jungle ravens are birds") as for prototypical properties

("jungle ravens are black").

Integrating the different justifications into a decision model

Putting together the different ideas generated by our participants, a tentative account of

how people process novel modified concepts might proceed as follows. We assume that the

default strategy is to attribute the properties of the unmodified prototype to the modified

concept. The likelihood of a given property being true of the modified concept would then be

subject to a number of tests, to see whether this default position should be altered for any

reason.

A first process would be to assess the assertability of the sentences. That is to say, one

tries to imagine a scenario in which a person might meaningfully assert the sentence. On a

given number of occasions, the modified sentence will prove distinctly harder to imagine in

this way. The modifier, being compatible with the attribute, appears to have little relevance

to it. Hence people state a preference for the unmodified sentence, on the grounds that it is

more likely to occur as a speech act in common conversation. Sentences with typical

modifiers were most prone to be rejected on these grounds. Since the modifier effect in

Experiment 1 was weakest for typical modifiers, we must conclude that the reduction in

likelihood generated by pragmatic oddness is relatively small compared to the other effects.

A second process would be to use the content of the modified sentence to retrieve

background knowledge in order to construct a more elaborate scenario around the modified

concept. Baby ducks may be like other creatures in which certain distinctive features only

appear later in maturity. Hence (using a background theory about biological development)

people may infer the likelihood of baby ducks not having webbed feet to be reasonably high.

These pragmatic and knowledge-based processes are likely to be mutually exclusive.

Only if the modified sentence makes some sense, does it become possible to elaborate it with

additional background knowledge. Thus "Indian seaweed is brown" may be rejected since the modifier has no relevance, whereas "jungle ravens are black" may appear pragmatically relevant, and then trigger theories about life in the jungle and how ravens might adapt to it.

The third and final reason given for rejecting the modified sentence in favor of the other was that the modified concept was really unfamiliar. Doves are white, but who knows what color Brazilian doves might be? In our process model, we would argue that this response would arise at the point where the sentence has passed a pragmatic acceptability test – a situation in which it may be sensibly uttered has been imagined – but no relevant knowledge has been found to help infer its truth or falsehood relative to the unmodified sentence. In this case, if the combination is very unfamiliar, then people may take the "safer bet" and assume that the unmodified sentence is more likely to be true. On the other hand if it is reasonably familiar, and the modifier is perceived as compatible and neutral with respect to the attribute, a default judgment of "equally likely" will be given.

This very tentative model clearly makes some interesting predictions about the time course of the decisions, such that a post-hoc division of decisions according to justification ought to differentiate them in terms of the time taken (assuming that the later justification is reasonably close to the reasoning behind the original decision). We would predict that pragmatic reasons should be fastest, followed by knowledge, followed by uncertainty, with decisions of "equally likely" slowest of all.

The modifier effect and models of conceptual combination

Having established that the modifier effect follows systematic patterns consistent with the inheritance of prototypical attributes, but subject to additional constraints, we next consider whether models of conceptual combination involving prototypes might account for the effect. That is, we reject the negative conclusion drawn by Connolly et al. (2007) that people simply do not use DS as a strategy in favor of an attempt to find an explanation for the systematic

patterns of data that have been shown. Since demonstrations of the modifier effect have so far been restricted to intersective conceptual combinations we consider two representative models (Smith et al., 1988, Hampton, 1987) that provide accounts of such combinations.

The first important point to note is that neither model says anything directly about the measure taken here – attribute likelihood judgments (see Jönsson & Hampton, 2008). In Smith et al.'s selective modification model the weight of an attribute is the degree to which that dimensional value contributes to the *relative typicality* of exemplars. It is not the subjective likelihood of the attribute being generally true. In Hampton's composite prototype model the importance of an attribute is the degree to which it contributes to similarity to the concept prototype, with observable consequences for both typicality and categorization. So while attribute weights or attribute importance are likely to correlate with judgments of likelihood there is as yet no empirical evidence of how strong this link would be. A further complication is that there is an ambiguity in the rating task. What is being judged? Is it the likelihood that the attribute is true of the whole class, or is it the proportion of members of the class for which it is true? Or is it some combination of these? This ambiguity is inherent in the semantics of generic sentences. When rating "birds fly" a person knows that the likelihood of the attribute being true of the whole class is zero, and so will tend to judge how likely it is to be true of any given bird. However when rating "Brazilian doves are white", the same person may presume that all members of the class are of the same color, and then rate the likelihood that that color is white. (Given that no participants ever asked for clarification, we suspect that in fact they respond on the basis of the associative strength between the attribute and the concept, and are unaware of the possible alternative interpretations).

Setting aside these problems, and assuming that the likelihood ratings are in fact closely related to feature weights in the models, what would the models predict about the modifier effect? Smith et al. (1988) proposed that when a concept like BLUE APPLE is formed, there

are two modifications to the representation of APPLE. First the range of possible values for color is restricted to BLUE, and second the weight of the color dimension is increased. Since (as described above) weights determine *relative* influence for determining typicality, if the weight of one dimension is increased, then the relative weight of all other dimensions will have to decrease. So the model would predict that (for example) differences in roundness will have less influence on the typicality of objects as blue apples than it will on their typicality as apples, since blueness will have come to influence a greater proportion of the variance in typicality, leaving less variation to be affected by shape. If the model is extended with the assumption that judged likelihood is based on dimensional weights, the model would predict a modifier effect. However the predicted effect would not be moderated by typicality, since formation of the concepts RED APPLE (typical) or BLUE APPLE (atypical) would generate equal shifts in weight to the color dimension (and hence away from shape). The difference between red and blue is a matter of how the values *within* the color dimension are weighted, but in both cases the increase in weight for the color dimension (and the consequent reduction in relative weight for other attributes) would be the same. Hence even the extended model fails to capture the phenomenon accurately.

The second model, the composite prototype model (Hampton, 1987; 1988) assumes that a composite prototype is formed that determines not only typicality but also membership in the complex concept category. When concepts combine, the importance of an attribute for a conjunctive concept is (approximately) the average of its importance for each conjunct. Thus if one considers "is fast" as an attribute of zebras, the attribute will have less importance for Namibian zebras, since Namibian creatures are not associated with the attribute of speed. The model does not therefore claim that attributes are inherited with unchanged importance (as the simple DS model would predict). Importance will be reduced to the extent that the attribute is not generally true of the modifier class.

Thus, if one equates attribute importance with likelihood, both models predict a modifier effect. The problem is that neither model explains why typical modifiers produce almost no effect (once pragmatic reasons are discounted), whereas atypical modifiers do. If there is nothing in memory to state that striped creatures are faster than Namibian creatures, then there is no explanation of why people consider striped zebras to be more likely than Namibian zebras to inherit the default attribute "runs fast" from zebras.

The conclusion must be that the modifier effect is not explained by either of these models. Bearing in mind the tenuous connection between attribute weights and judged likelihood, it is safest to conclude that while the modifier effect is not inconsistent with these models, neither can it be readily explained in their terms.

Explaining why the modifier effect interacts with typicality

It is therefore necessary to turn elsewhere for an explanation of the typicality effect. We suggest two possibilities, based respectively on similarity and familiarity. Developing a full account will need further experiments beyond the scope of the present study, designed to separate out these and other possible explanations.

The first possibility would be to stipulate that confidence in any attribute is reduced as a function of the dissimilarity between the modified and unmodified concepts. Hence the default attributes of "apple" would be judged less likely to be true of "blue apple" than of "red apple", just because the concept "blue apple" is less similar to "apple" than is "red apple". This move would simply apply to within-category induction the same general similarity principle that has been discovered to work well for between-category induction (Osherson et al., 1990, Sloman, 1993). Clearly, according to either of the models discussed in the previous section, atypical modifiers would generate complex concepts with less similarity to the original concept, and so the strength of the argument from concept to subconcept would be correspondingly weaker. Calvillo and Revlin (2005) showed that people tend to

have reduced confidence in class inclusion statements of the form "All A are B" when A is an atypical member of category B. In other words, confidence in Namibian zebras actually being a subclass of zebras may be lower, because of the atypicality of the modifier. If they are less likely to be zebras, then they are less likely to inherit any of the prototypical attributes of zebras.

A second possibility is that the *familiarity* of the modified concept determines confidence that the attribute has been inherited. The moderation of the modifier effect by typicality can be readily explained this way. Typically modified concepts such as quacking ducks and striped zebras are very familiar, while atypically modified concepts (baby Peruvian ducks and giant Namibian zebras) are not. Familiar modified concepts introduce knowledge effects. Rips (1995) pointed out that to study prototype combination processes in relatively pure form, it is necessary to use novel unfamiliar combinations (a proposal he referred to as the "no peeking principle"). The atypically modified noun phrases used were all accordingly unfamiliar, but necessarily the typically modified phrases were mostly not. Typical modifiers may therefore generate no modifier effect (other than pragmatic-based effects) simply because they are already familiar concepts whose attributes are known. It is not possible to create a typically modified concept that is not at least as familiar as the concept itself. One way to express this idea is to suppose that participants use an availability heuristic (Tversky & Kahneman, 1973.) If they cannot retrieve an exemplar from memory, then they will return a reduced confidence in the attributes of a concept. (Evidence for this account comes from a series of studies reported in Hampton, Passanisi and Jönsson, 2010).

In conclusion, the modifier effect presents a challenge to models of prototype combination in intersective complex concepts. We have provided evidence that modified concepts do in fact inherit prototypical attributes from their constituent concepts, and in ways that are moderated by information within the concept prototype. However while two

conceptual combination models have possible accounts of why the effect should occur,

neither has a good account of why the effect should be moderated by typicality of the

modifier. We nonetheless reject the interpretation offered by Connolly et al. that the effect

demonstrates the lack of a default to stereotype (DS) strategy. A tentative model was offered

of how the different reasons for preferring unmodified sentences might arise in the process of

making a decision. Two additional suggestions have been made of how the interaction with

typicality may work – through a generalized similarity principle, or through familiarity of

typical modifiers introducing knowledge-based effects. We suspect that more than one of

these effects may in fact be at work across different examples.

References

Calvillo, D. P. & Revlin, R. (2005). The role of similarity in deductive categorical inference. *Psychonomic Bulletin and Review, 12,* 938-944.

Connolly, A. C., Fodor, J. A., Gleitman, L. R., & Gleitman, H. (2007). Why stereotypes don't even make good defaults. *Cognition, 103,* 1-22.

Cree, G. S., & McRae, K. (2003). Analyzing the factors underlying the structure and computation of the meaning of chipmunk, cherry, chisel, cheese, and cello (and many other such concrete nouns). *Journal of Experimental Psychology: General, 132*, 163-201.

Davidson, D. (1967). Truth and Meaning. *Synthese* 17: 304-323.

Fodor, J. A. (1998). *Concepts: Where cognitive science went wrong*. Oxford: Clarendon Press.

Fodor, J.A., & Lepore, E. (2002). The Compositionality Papers. Oxford: Oxford University Press.

Frege, G. (1923). *Compound Thoughts. Logical Investigations*. Oxford: Basil Blackwell.

Grice, H. P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and Semantics* (3rd edition, pp. 41-58). New York: Academic Press.

Hampton, J. A. (1979). Polymorphous Concepts in Semantic Memory. *Journal of Verbal Learning and Verbal Behavior, 18,* 441-461.

Hampton, J. A. (1987). Inheritance of Attributes in Natural Concept Conjunctions. *Memory & Cognition, 15,* 55-71.

Hampton, J. A. (1988). Overextension of Conjunctive Concepts: Evidence for a Unitary Model of Concept Typicality and Class Inclusion. *Journal of Experimental Psychology: Learning, Memory, and Cognition 14,* 12-32.

Hampton, J. A. (1991). The Combination of Prototype Concepts. In P. J. Schwanenflugel (Ed.) *The Psychology of Word Meanings*. (pp. 91-116). Hillsdale, Lawrence Erlbaum

Associates.

Hampton, J. A. (1997a). Conceptual combination. In K. Lamberts & D. Shanks, (Eds.) *Knowledge, concepts and categories*. pp.135-162. Hove: Psychology Press.

Hampton, J. A. (1997b). Conceptual combination: Conjunction and negation of natural concepts. *Memory & Cognition, 25*, 888-909.

Hampton, J. A. (2006). Concepts as Prototypes. In B. H. Ross, (Ed.) *The Psychology of Learning and Motivation: Advances in Research and Theory, 46,* 79-113 .

Hampton, J. A. (2009). Stability in concepts and evaluating the truth of generic statements. In F.J.Pelletier (Ed.), *Kinds, Things and Stuff: Concepts of generics and mass terms. New Directions in Cognitive Science, Vol. 12, pp. 80-99.* Oxford: Oxford University Press.

Hampton, J. A., & Jönsson. M. L. (2011). Typicality and Compositionality: The Logic of Combining Vague Concepts. *In* E. Machery, W. Hinzen, & M. Werning (Eds.) *Handbook on Compositionality.* Oxford: Oxford University Press, forthcoming.

Hampton, J.A., Jönsson, M.C., & Passanisi, A. (2009). The modifier effect: Default inheritance in complex noun phrases. *In* N. A. Taatgen & H. van Rijn (Eds.), *Proceedings of the 31st Annual Conference of the Cognitive Science Society* (pp. 303-308). Austin, TX: Cognitive Science Society.

Hampton, J. A., Passanisi, A., & Jönsson, M. L. (2010). The Modifier Effect and Property Mutability. Unpublished MS: City University London (under review).

Johnson, C. & Keil, F. C. (2000). Explanatory understanding and conceptual combination. In F. C. Keil & R. A. Wilson (Eds.), *Explanation and cognition* (pp. 328-359). Cambridge, MA: MIT Press.

Jönsson, M. L. (2008) *On Compositionality: Doubts about the structural path to meaning* Doctoral Dissertation, Lund University, Sweden

Jönsson, M. L., & Hampton, J. A. (2006). The Inverse Conjunction Fallacy. *Journal of*

Memory and Language, 55, 317-354.

Jönsson, M. L. & Hampton, J. A. (2008). On prototypes as defaults. (Comment on Connolly, Fodor, Gleitman and Gleitman, 2007). *Cognition, 106,* 913-923.

Khemlani, S., Leslie, S.J., & Glucksberg, S. (2009). Generics, prevalence, and default inferences. In N. Taatgen, H. van Rijn, J. Nerbonne & L. Schomaker (Eds.), *Proceedings of the 31st Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.

Krifka, M., Pelletier, F. J., Carlson, G. N., ter Meulen, A., Chierchia, G., & Link, G. (1995). Genericity: An Introduction. In G. N. Carlson & F. J. Pelletier (Eds.), *The Generic Book* (pp. 1-124). Chicago: University of Chicago Press.

Kunda, Z., Miller, D. T., & Clare, T. (1990). Combining social concepts: the role of causal reasoning. *Cognitive Science, 14*, 551-578.

Machery, E., Hinzen, W., & Werning, M. (Eds.) (2011). *Handbook on Compositionality.* Oxford: Oxford University Press.

Murphy, G. L. (1988). Comprehending complex concepts. *Cognitive Science, 12*, 529-562.

Murphy, G. L. (2002). *The Big Book of Concepts*. Cambridge MA: MIT Press.

Osherson, D. N., Smith, E. E., Wilkie, O., López, A., & Shafir, E. (1990). Category-based induction. *Psychological Review, 97,* 185-200.

Pelletier, F.J.(Ed.) (2009). *Kinds, Things and Stuff: Concepts of generics and mass terms. New Directions in Cognitive Science*, *Vol. 12*. Oxford: Oxford University Press.

Rips, L. J. (1995). The current status of research on concept combination. *Mind and Language, 10,* 72-104.

Rosch, E., & Mervis, C. B. (1975). Family Resemblances: studies in the internal structure of categories. *Cognitive Psychology, 7,* 573-605.

Sloman, S. A. (1993). Feature-based induction. *Cognitive Psychology, 25*, 231-280.

Smith, E. E., Osherson, D. N., Rips, L. J., & Keane, M. (1988). Combining Prototypes: A
    Selective Modification Model. *Cognitive Science, 12*, 485-527.

Tversky, A. & Kahneman, D. (1973). Availability: A heuristic for judging frequency and
    probability. *Cognitive Psychology, 5*, 207-232.

Appendix

Materials used in Experiments 1 and 2 (taken with permission from Connolly et al., 2007).

HEAD	TYPICAL	ATYPICAL	DOUBLE ATYPICAL	PREDICATE
candles	scented	purple	expensive purple	are made of wax
caterpillars	furry	poisonous	Canadian poisonous	have many legs
catfish	edible	speckled	rare speckled	have whiskers
cellars	damp	dry	dry northern	are dark
coins	shiny	old	old Egyptian	are round
coyotes	carnivorous	white	old white	howl
crocodiles	egg-laying	African	albino African	are dangerous
doves	feathered	Brazilian	flightless Brazilian	are white
ducks	quacking	baby	baby Peruvian	have webbed feet
hamsters	pet	male	male hairless	live in cages
kites	colorful	silk	silk weather	have strings
lambs	fluffy	Norwegian	long-haired Norwegian	are white
limousines	expensive	inexpensive	old inexpensive	are long
napkins	white	blue	blue checkered	are made of paper

HEAD	TYPICAL	ATYPICAL	DOUBLE ATYPICAL	PREDICATE
nectarines	sweet	bitter	bitter purple	are juicy
ostriches	ugly	paleolithic	paleolithic European	have long necks
parkas	hooded	acrylic	machine-washable acrylic	are warm
pearls	expensive	South Sea	oval South Sea	are white
penguins	flightless	solitary	solitary migrant	live in cold climates
pigeons	grey	Tibetan	Tibetan hunting	live in parks
pigs	dirty	furry	furry black	live on farms
ravens	feathered	jungle	young jungle	are black
refrigerators	kitchen	commercial	inexpensive commercial	are used for storing food
rhubarb	sour	homegrown	homegrown Albanian	is used for pies
rifles	hunting	used	used army	are dangerous
roosters	red	yellow	flying yellow	live on farms
saxophones	jazz	handmade	expensive hand-made	are made of brass
scarves	knitted	inexpensive	inexpensive synthetic	are worn for warmth
seaweed	edible	baked	baked Indian	is green
shacks	storage	new	new Appalachian	are made of wood

HEAD	TYPICAL	ATYPICAL	DOUBLE ATYPICAL	PREDICATE
shirts	cotton	itchy	itchy canvas	have buttons
sinks	metal	antique	round antique	are found in kitchens
sofas	comfortable	uncomfortable	uncomfortable handmade	are found in living rooms
squirrels	tree dwelling	Nicaraguan	black Nicaraguan	eat nuts
storks	white	domestic	domestic hybrid	have long legs.
strawberries	red	Lithuanian	sour Lithuanian	have seeds
thimbles	sewing	painted	Belgian painted	are made of metal
tortoises	green	South American	South American fighting	are slow
wagons	cargo	futuristic	futuristic fruit	are used by pulling them
zebras	striped	Namibian	giant Namibian	are fast

Additional materials for Experiment 2

HEAD	TYPICAL	ATYPICAL	DOUBLE ATYPICAL	PREDICATE
elk	brown	spring	female spring	have antlers
drums	round	electronic	German electronic	have skin

Author note

Correspondence concerning this article should be addressed to James A. Hampton, Psychology Department, City University, Northampton Square, London EC1V OHB, UK. Email: hampton@city.ac.uk.

Footnotes

[1.] Generic sentences include a much wider range of examples, including attributes that are clearly not a part of the content of the concept like "mosquitoes carry the West Nile virus", and statements about the habits of individuals such as "John smokes a cigar after dinner". See Carlson & Pelletier (1995) for further discussion.

[2.] Although both "solitary" and "living in cold climates" are attributes, for clarity the term "modifier" will be used to refer to the adjective placed in front of the noun, and the term "attribute" to the property or attribute asserted in the predicate.

[3.] An anonymous reviewer pointed out that an alternative way to elicit relative rather than absolute likelihood judgments would be to place the two sentences side-by-side, but have them each rated separately on the 10 point scale. Since the size of the effect averaged only about 2 points on the scale, our procedure may have risked underestimating the effect. We appreciate this suggestion, but feel that our method of asking for a direct comparison of the two sentences was the best way to set up the task pragmatically for subsequently eliciting justifications where a preference had been expressed.

Table 1

Mean (and Standard Deviation) Likelihood Judgments from Connolly et al. (2007) and

from Experiment 1.

	Connolly et al.	Experiment 1
Unmodified condition	8.36	8.31 (3.55)
(ducks have webbed feet)		
Typically modified condition	7.71	7.51 (3.31)
(quacking ducks have webbed feet)		
Atypically modified condition	6.91	6.59 (2.53)
(baby ducks have webbed feet)		
Twice atypically modified condition	6.48	6.27 (3.05)
(baby Peruvian ducks have webbed feet)		

Table 2

Means (and Standard Deviations across Items) for Percentage Response Proportions for

Target Sentence Pairs in Experiment 2.

Comparing the unmodified sentence with:	Unmodified is more likely	Modified is more likely	Equally likely
- Typically Modified	23% (12)	8% (14)	69% (14)
- Atypically Modified	36% (14)	2% (5)	62% (13)
- Twice Atypically Modified	37% (17)	3% (5)	60% (16)
Filler Sentence Pairs	29% (8)	26% (7)	45% (13)

Table 3

Frequency of Different Types of Justifications for Judgments in Each Condition of

Experiment 2, together with Percentage of Justifications Given for that Condition.

	Typically Modified	Atypically Modified	Double Atypically Modified	Total
Reasons why unmodified more likely				
Pragmatic	61 (58%)	77 (42%)	71 (40%)	209 (45%)
Knowledge	8 (8%)	37 (20%)	42 (23%)	87 (19%)
Uncertainty	1 (1%)	32 (18%)	34 (19%)	67 (14%)
Reversal	7 (7%)	3 (2%)	5 (3%)	15 (3%)
Other (unclassified)	28 (27%)	33 (18%)	27 (15%)	88 (19%)
No response at all	21	20	26	67
Total	126	202	205	533
Reasons why unmodified less likely				
Pragmatic	6 (15%)	1 (8%)	3 (20%)	10 (15%)
Knowledge	29 (73%)	8 (67%)	10 (67%)	47 (70%)
Other (unclassified)	5 (12%)	3 (25%)	2 (13%)	10 (15%)
No response or didn't know	5	11	4	20
Total	45	23	19	87