



City Research Online

City St George's, University of London

Citation: Hampton, J. A., Passanisi, A. & Jonsson, M. L. (2011). The modifier effect and property mutability. *Journal of Memory and Language*, 64(3), pp. 233-248. doi: 10.1016/j.jml.2010.12.001

This is the unspecified version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://city-ref-test.eprints-hosting.org/id/eprint/1004/>

Link to published version: <https://doi.org/10.1016/j.jml.2010.12.001>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, 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. For full details of reuse please refer to [City Research Online policy](#).

In press, *Journal of Memory and Language*, 2011

The Modifier Effect and Property Mutability

James A. Hampton

City University London

Alessia Passanisi

Università di Catania

Martin L. Jönsson

Lund University

Address for correspondence:

James A. Hampton

Psychology Department

City University London

Northampton Square

London EC1V OHB

UK

Keywords: concepts, modifier, mutability, compositionality

Abstract

The modifier effect is the reduction in perceived likelihood of a generic property sentence, when the head noun is modified. We investigated the prediction that the modifier effect would be stronger for mutable than for central properties, without finding evidence for this predicted interaction over the course of five experiments. However Experiment 6, which provided a brief context for the modified concepts to lend them greater credibility, did reveal the predicted interaction. It is argued that the modifier effect arises primarily from a general lack of confidence in generic statements about the typical properties of unfamiliar concepts. Neither prototype nor classical models of concept combination receive support from the phenomenon.

The modifier effect is a phenomenon reported by Connolly, Fodor, Gleitman and Gleitman (2007) in which the truth of a generic statement is considered less likely when the subject noun is preceded by an atypical modifier. In their study, a bare plural generic statement such as “ravens are black” was rated as increasingly less likely to be true when the subject noun “ravens” was qualified by a typical modifier (“feathered ravens are black”), an atypical modifier (“jungle ravens are black”), or two atypical modifiers (“young jungle ravens are black”). This effect has proven to be highly robust (Jönsson & Hampton 2006; 2011), and is of considerable potential theoretical interest. For example, Connolly et al. used the effect to argue against prototype theories of concept modification, which, they suggested, should predict that common properties of a concept are inherited with unaffected strength when the concept is modified.

The purpose of this article is not directly to take issue with Connolly et al.’s interpretation of the effect, which we have done elsewhere (see Jönsson & Hampton, 2006; 2008; 2011; Jönsson, 2008; Hampton & Jönsson, 2011, for our position on the theoretical issues involved, and empirical explorations of the effect.) Rather, the article presents a new empirical investigation of the phenomenon in order to explore three possible explanations. Before presenting the particular research question we addressed, the modifier effect will be explained in a little more detail.

To set the context of this investigation, let us briefly consider what the effect shows. Depending on the relative frequencies of the objects involved, there is often only a weak constraint between the likelihood that a property is true of a class in general and the likelihood that it is true of a subclass. Most ravens may be black, while jungle ravens may be green. Unless the subclass is a substantial majority of the class, or the property is true of almost all of the class, the two proportions (ravens that are black versus jungle ravens that are

black) are free to vary independently. Only in the case of a universally quantified statement does a logical constraint appear (“all ravens are black”, and “all jungle ravens are ravens” entails that “all jungle ravens are black”). We explored this case of universally quantified statements in a previous paper (Jönsson & Hampton, 2006) where we demonstrated a robust tendency to fallacious reasoning, in which people’s judgments violated this logical constraint.

Given the lack of any strong constraint for generic sentences, there may nonetheless be ways of using world knowledge to predict a difference in likelihoods. A modified concept may be more likely to possess a property if the modifier “supports” the property. Thus fresh vegetables are more likely to be good for your health than vegetables alone. Fast cars are more likely to consume a lot of gasoline than are cars in general. Similarly a modifier may contradict a property, as in cases such as “broken clocks tell the time” (see for example Springer & Murphy, 1992). The basis of the modifier effect is that one should select modifiers with no *prima facie* connection to the property in question. For instance, in the sentence “dry cellars are dark”, although “dry” is an atypical modifier of “cellars”, being dry has no direct bearing on the question of whether cellars are dark or not. Connolly et al. (2007) based their selection of materials on this principle of independence, and so it is in this context that the reduction in judged likelihood is to be understood.¹ If, they argued, a modified concept inherits the typical properties of the unmodified noun by default, there should be no reduction in judged likelihood. The finding of the modifier effect, they concluded, is evidence against the default inheritance of prototype properties, and therefore is

¹ Connolly et al.’s (2007) original items were examined more closely in the light of this principle in Jönsson and Hampton (2011), where people were asked to justify their decisions about relative likelihood of the same set of materials. There were very few items that showed any systematic deviation from the principle of independence required here. The full set of sentences are presented in the Appendix so that the reader can verify this for themselves.

also evidence against models that suggest that prototypes combine to form the meaning of complex noun phrases.

Having set the context, we now turn to the main purpose of the paper. In order better to understand the modifier effect, the question we address is how the effect interacts with the mutability of the asserted predicate, (“are black” in the ravens example above). The notion of mutability, as developed by Sloman, Love and Ahn (1998) (see also Sloman & Ahn, 1999), is a key variable in theories of conceptual memory. A highly mutable property is one that could easily be imagined to be different. “Is black” is a mutable property of a raven, since one could easily suppose that ravens were brown rather than black, without any consequent changes to their other properties of being birds, scavengers, intelligent, a certain size and so forth. By contrast a *central* property is one that has causal dependencies with other properties of the concept. For example, having wings is less mutable and more central for a raven than is being black. Since their wings enable them to fly, and it is on account of their flying that they can roost in trees, find their food and so forth, if ravens in general had no wings many other properties would also have to be different. So having wings is a central property of ravens whereas being black is a mutable property. According to causal model theory (Sloman, 2005), mutability/centrality is a consequence of the dependencies that exist between the different properties that compose the concept’s schema or prototype. The more dependencies that link to a given property, the more central the property will be. The fewer dependencies, the more mutable it will be. Mutability is thus a matter of degree. Note that mutability is not the same as the prevalence of a property. There may happen to be just as many black ravens as there are ravens with wings. It is the links with other properties that make a property central rather than mutable.

In their paper, Connolly et al. (2007) demonstrated that the modifier effect is moderated by the typicality of the modifier. “Young jungle ravens are black” was rated as less likely than “Feathered ravens are black”. However no research has yet considered whether the effect is also moderated by the centrality versus mutability of the predicate. Would the effect of an atypical modifier be greater for a mutable property like “is black” than for a central property of ravens such as “has wings”? The answer to this question has the potential to shed light on the basis of the modifier effect, (more particularly as Connolly et al. used primarily mutable properties of the concepts).

How should mutability of a property affect the size of the modifier effect? We will discuss three accounts. We first consider theories of prototype combination. According to psychological models of concept modification (e.g. Hampton, 1987; 1988; Murphy, 1988; 2002; Smith et al., 1988) modifying a concept such as “raven” to produce the concept “jungle raven” involves an operation on its conceptual content. Assuming that “jungle” is taken to refer to a habitat, the concept of raven will be modified to include the property [HABITAT = *jungle*]. The models differ in detail, but we can use Smith et al.’s model to represent the general case. It has been argued (e.g. by Connolly et al. 2007) that the modifier effect is evidence against such models, since according to this model, while the dimension HABITAT receives additional weight, other prototypical properties of the concept which are unrelated to habitat (such as [COLOR = *black*]) should remain in the schema unchanged. Jönsson and Hampton (2008) pointed to a number of problems with this argument. For example, given that habitat now has extra importance, and assuming that importance is a *relative* measure, other properties must then have reduced importance, which would lead to a general modifier

effect.² But, for models like Smith et al. (1988) it is also the case that a central property is less likely to suffer a reduced importance than a mutable property. Because of the causal dependencies among central properties, the change in a single dimension, such as HABITAT, is likely to have little effect on a central property such as “has wings”. The network of interlinked properties that form the central part of the concept protects individual central properties from radical change. On the other hand mutable properties such as COLOR are more likely to be affected by a shift in weight resulting from modification. Since it is easy to imagine the color of ravens being other than black, one is more easily tempted to entertain the possibility of jungle ravens not being black. Hence confidence in the likelihood of the property being true goes down. In sum, prototype models lead naturally to the prediction that mutability/centrality will moderate the effect of a modifier.

Our second account takes its inspiration from Connolly et al.’s own interpretation of their result, and its roots in the work of Fodor and Lepore (2002). Connolly et al. propose two stages in conceptual combination. The first stage is a classical logical combination of the two concepts, and involves no consideration of properties or prototypes. Thus a jungle raven is a raven and has some relation to jungles, and that is all. They then propose that in a second stage pragmatic and knowledge-based reasoning may be used to derive the probability that a jungle raven will have the same properties as other ravens or as other jungle creatures. The classical account of conceptual combination is a simple and direct way of preserving the principle of Compositionality (Fodor and Lepore, 2002; Machery et al., 2011). This principle holds that the meaning of a complex concept should be derived only from the meanings of its parts, and their mode of combination. Connolly et al.’s account makes no prediction

² There is also considerable empirical evidence, backed up by detailed theoretical models, that the features of prototypes are not inherited without alteration – see for example Hampton (1987; 1997), Murphy (1988), and Medin and Shoben (1988).

concerning how the modifier effect should be moderated by mutability, since they give no account of how the second phase of a combination might work. Our second theoretical account is therefore based simply on the first, classical, phase of their account. If concepts are combined classically, then the prototype information associated with each concept (including whether a property is central or mutable) will not be inherited by the combined concept. In that case, we would predict that both central and mutable properties should become less likely to be true in equal measure when the concept is modified. In passing from a negative claim (non-inheritance of prototype features) to a positive prediction (reduction in likelihood under modification), we follow a simple principle of regression to the mean. As unpredictable noise is added to the system during modification, so the tendency will be for those properties with higher likelihood at the start to end up with reduced likelihood at the end. Thus, while this prediction of non-interaction of mutability with the modifier effect is not derivable from Connolly et al. (and should not be attributed to them), we take it to be the prediction most in keeping with the notion that modified concepts do not inherit the default attributes of their constituent concepts (the inference that they *do* draw from their results). Since mutability is about the conceptual content *within* a prototype, if the prototype is not inherited, then mutability should not interact with the modifier effect. We will term this the Non-Inheritance account.

Our third and final account of the modifier effect is based on more general pragmatic considerations. Connolly et al. noted that (perhaps surprisingly) even a typical modifier produced a reduction in likelihood for prototypical properties. Thus “feathered ravens are black” was considered less likely than “ravens are black”. They suggested that this result could reflect the infelicity of the expression. Why would the speaker choose to include the typical modifier, if the statement was in any case generally true of the whole class? Estes and

Glucksberg (1999) demonstrated a similar effect. Sentences with a prototypical feature and an atypical modifier such as “peeled apples are round” were generally slow for people to verify relative to more pragmatically relevant sentences such as “peeled apples are white”, where the property is true only of the subclass. This difference was reversed when the sentence was changed to “even peeled apples are round”, where the pragmatics now favors the more general property.

Evidence for the Pragmatic account comes from a study by Jönsson and Hampton (2011, Experiment 2) in which participants first judged whether modified sentences were more or less likely to be true than unmodified versions of the same sentence, and were then asked to justify their choices. A large proportion of justifications referred to pragmatic issues, not just for typical but for atypical modifiers as well. In the case that the modifier works by reducing the pragmatic well-formedness of the statements and of the experimental task in general, then we would expect to find the effect working equally strongly for mutable and central statements. The prediction of the Pragmatic account is thus in line with the Non-Inheritance account: there should be no interaction between the modifier effect and mutability of the predicates. (These two accounts will be further differentiated as the paper progresses.)

The first experiment was a very simple study designed to obtain pairs of properties that varied in mutability for a set of concepts, and to test whether more mutable properties were considered less likely to be true of the modified concept (e.g. jungle ravens), compared to more central properties. The non-inheritance position outlined above would predict that, since prototypes of unmodified concepts are not inherited by modified concepts, then mutable and central properties should be considered equally unlikely of the modified concept. Similarly, the Pragmatic account suggests that the modifier will affect all properties

equally as a consequence of introducing general uncertainty about the unfamiliar phrase. On the other hand, according to the Prototype Combination account, because of their inherited causal dependencies, central properties should be considered more likely to be true than mutable properties of the modified concepts. The experiment therefore assessed which of two sentences “MN are P_1 ” and “MN are P_2 ” would be considered more likely to be true where MN is a modifier + noun phrase and the predicates P_1 and P_2 differ in terms of mutability for the concept N.

A pretest (Experiment 1a) was used to find pairs of predicates P_1 and P_2 that differed in their relative mutability for the same subject noun N. A variant of Sloman, Love and Ahn’s (1998) imagination task was used to measure mutability, and 20 pairs with the clearest difference in mutability were selected for further testing.

Experiment 1a: Pretest

Method

Participants. Twenty-three students (undergraduate and graduate students) at City University London participated voluntarily for no reward.

Materials. Materials were adapted from Connolly et al. (2007), who provided a set of materials comprising 40 nouns, each with a familiar property and a set of modifiers. Feature norms (Cree & McRae, 2003) were used to find a second property for each of 39 of these nouns (one was omitted in error). Since Connolly et al. used broadly mutable properties (e.g. ravens are black, sofas are comfortable) the new properties were chosen to be more central. Since the task required people to imagine a concept with a missing feature, all properties were then expressed in negated form.

Design and Procedure. The instructions were as follows:

“On each line below you will see the name of a kind of thing on the left and then two

negative statements on the right. We want you to try to imagine an example of the thing for which the first statement is true, and then an example of the thing for which the second statement is true. Then circle whichever statement was easier to imagine. When thinking of the examples, try to assume that the example is just like a normal instance in all other respects except for the one denied in the statement. For instance, for the combination

Parka : Is not warm Cannot be used for clothing

think of a normal parka that is not warm, then think of a normal parka that cannot be used for clothing. If you think that it is easier to imagine a parka that is just like a normal parka but is not warm than it is to imagine a parka that is just like a normal parka except that it cannot be used as clothing, circle “not warm”. Otherwise circle "cannot be used as clothing". Remember, circle the negative statement that is easiest to imagine being true, and try to think of things that are normal in all other respects except for the one denied in the statement.”

Beneath the instructions the 39 unmodified nouns each appeared together with the two negated properties. Order of nouns was randomized, and order of the two properties for each noun was balanced between groups. The task took about 10 minutes to complete.

Results

Overall, 34 out of the 39 new properties were judged to be less easy to imagine when negated than the original properties, in keeping with our intuitions. The 20 pairs with the strongest difference between properties were chosen for the subsequent experiment (all 20 of the differences were significant on a 1-tailed sign test, $\alpha = .05$.) The proportion of participants who found it easier to imagine the mutable property being changed varied across the 20 selected items from .70 to .91, with a mean of .81. The nouns, modifiers and pairs of

properties may be seen in Appendix A.

Experiment 1b: Main test

Method

Participants. Twenty-four students (undergraduate and graduate students) and members of the administrative staff at City University London participated voluntarily.

Materials. Each booklet consisted of 20 target and 30 filler pairs of sentences, each member of a pair having the same atypically modified head noun (atypical modifiers were taken from Connolly et al.) and one of two predicates. For instance, for the noun “doves” the following pair was created,

- A) Brazilian doves are white. (*mutable*)
- B) Brazilian doves have wings. (*central*)

Filler pairs were randomly distributed through the booklet, with 6 fillers at the start to avoid warm-up effects. Filler pairs were similar to the targets in both structure and likelihood, and provided opportunities to use all 3 possible responses while also disguising the nature of the manipulation within the list.

Design and Procedure. Order of sentences within pairs was balanced across two versions of the booklet. Three options appeared to the right of each 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. Participants circled the option they found most appropriate for each sentence pair. The task took about 20 minutes.

Results

All ANOVAs in this paper were run with items and with participants as random effects. The results of the two analyses (F_1 and F_2) were then combined into a Min F' statistic (see Clark, 1973) which tests for generalizability to both populations. As predicted, property

mutability had a strong effect on the judged likelihood of the sentences. The central property sentence was selected as being more likely 42% of the time, and the mutable only 14%, with the sentences otherwise being judged equally likely. Nineteen out of 20 sentence pairs showed the effect. ANOVA confirmed a very significant difference in proportions ($Min F'(1, 35) = 16.2, p < .001$). Of the 23 participants, 12 consistently chose the central sentence, another 7 judged most of the pairs equally likely, and the remaining 4 divided their responses equally between the 3 options.

Discussion

As predicted by the prototype combination account, central properties of the unmodified prototype were judged more likely to be true of a modified concept than were mutable properties. There was therefore evidence that the modified noun concept had inherited the causal dependency structure from its parent noun concept, as predicted by prototype combination models. However since the experiment did not include a control condition with *unmodified* sentences, the results do not tell us whether the degree to which the modifier is having an effect is greater for the mutable sentences. Perhaps the central sentences were already more likely to be true, before the modifier was applied. Experiment 2 pursued this possibility.

Experiment 2

Having established that mutability affects the rated likelihood of properties for modified concepts, we next turned to the main research question outlined in the Introduction. In Experiment 2 both modified and unmodified versions of the sentences were included in order to test whether mutable properties would show a greater modifier effect than central properties, as predicted by prototype combination models, or whether there is no interaction as predicted by our Non-Inheritance and Pragmatic accounts.

Method

Participants. Thirty-six Masters students at City University London participated voluntarily by completing a booklet in a classroom setting.

Materials. Four more items were added from Experiment 1a to arrive at 24 pairs of properties, one central and one mutable, for 24 nouns. Four sentences were then constructed for each subject noun by crossing (a) whether the noun was modified with an atypical modifier or left unmodified, and (b) whether the property was mutable or central. The quadruples were divided into 4 sets of 6 at random, and these sets were rotated across conditions across 4 booklets, so that each booklet contained 6 sentences in each condition, and each subject noun occurred only once in each booklet. Thirty-two filler sentences, some clearly true, others clearly false, and others borderline true, were distributed at random through the booklet. About half of the fillers had modified subject nouns.

Design & Procedure. Two factors were manipulated within-subjects and within-items: whether or not the subject noun was modified (with an atypical modifier), and whether the property was mutable or central. Sentences were rated on a 10 point likelihood rating scale. The following instructions were given, together with examples of a clearly true, a doubtful and a clearly false sentence:

“The following pages contain a list of English sentences. The numbers 1 through 10 appear to the right of each sentence. We would like you to circle the number to the right of each sentence that you feel best captures how likely *you* think it is that the sentence is true. 1 = very unlikely and 10= very likely (the numbers in between represent intermediate levels of certainty).”

The task took about 15 minutes.

Results

Figure 1 shows mean likelihood ratings for each condition. Unmodified sentences (black bars) were judged more likely than modified (a difference of 1.5) and central sentences were judged more likely than mutable (a difference of 1.2). ANOVA was run with factors of Modifier and Mutability. Both main effects were significant (for Modifier, $Min F'(1, 50) = 15.2, p < .001$, and for Mutability, $Min F'(1, 34) = 9.2, p < .005$). There was no evidence of an interaction (F_1 and F_2 both < 1). Contrary to the predictions of prototype combination models, but in keeping with our Non-Inheritance and Pragmatic accounts, the size of the modifier effect was nearly identical for properties that were central or mutable for the head noun concept.

Discussion

The results of Experiment 2 confirmed the finding in Experiment 1 that modified concepts were judged more likely to possess properties judged to be central for unmodified concepts than those judged to be mutable. However we failed to find any evidence that mutability affects the size of the modifier effect, since these central properties were *also* judged more likely for the unmodified concepts, and to just the same extent. The results were therefore consistent with the idea that concepts are combined in a classical fashion, and that *all* prototypical properties are subsequently considered less likely for a modified concept, regardless of centrality. At the same time the results are consistent with a general pragmatic view that proposes that people generally become suspicious of unfamiliar modified concepts.

One possible way to save the prototype combination account is to suggest that the central properties used were not sufficiently central. Or perhaps the imaginability task was not picking out the right dimension. Experiment 3 used a new pretest for materials asking participants to choose whether properties were necessary, central or mutable, and these data were used to select a new set of materials.

Experiment 3 was also designed to test a differential prediction of the Pragmatic and Non-Inheritance accounts of the effect. To this end we included category properties as predicates in the experiment. The idea was that according to the Pragmatic account there is a general suspicion aroused by the unfamiliarity of the phrases and oddness of the statements and of the rating task. Hence the modifier effect should continue to be found in equal measure, regardless of whether the property is mutable, central or categorical. However, on the Non-Inheritance view, at least if it is combined with the classical theory of concept composition that inspired it, there should be a difference between category properties and other properties, since on this view some logical truths do follow. Consider for instance the statements “ravens are birds” and “jungle ravens are birds”. According to concept composition, one can be confident that jungle ravens are ravens. That is what compositionality guarantees (aside from non-compositional idioms such as a *nest egg*, which is not an egg). However if one is certain that jungle ravens are ravens, and equally certain that ravens are birds, there should be no reason to doubt that jungle ravens are birds. It should certainly be a lot more likely that they are birds than that they are black. A further reason for expecting that category predicates would be immune to the modifier effect was that Hampton (1987) found evidence that properties considered necessary for a concept should be inherited as necessary for a subclass. So even if the Non-Inheritance model predicts no general interaction with mutability, it does predict differences between categorical properties and other properties.

It may be argued that according to Fodor’s non-definitional theory of concepts (Fodor, 2000) the fact that ravens are birds is not a “part” of the concept of RAVEN, which is simply an unanalyzed symbol that points to the class of ravens. Yet even with this way of individuating conceptual content, a person who holds the belief that “All ravens are birds”,

and who understands from compositionality that “All jungle ravens are ravens”, is very likely to assent to the statement that all jungle ravens are birds. Thus while a non-definitional account of concepts, like Fodor’s, does not give the result *directly*, the account only needs to be minimally enriched with the transitivity of class inclusion to generate the prediction that category statements should be less affected by a modifier than mutable properties.

Experiment 3a: Pre-test

Experiment 3 was designed to test for the interaction of mutability and modification with three kinds of property: mutable properties, a new set of central properties, and categorical properties. A pre-test was run asking people to differentiate properties into three levels of importance – mutable, central and necessary. Categorical properties (which we expected to be judged necessary) were also included in the pre-test, so there were actually four kinds of pre-selected sentences for each concept which were then judged according to the three levels of importance. For practical reasons, the pre-test was run in Italian at an Italian University. Because the materials presented no problems for translation, and included only everyday concepts familiar in both cultures, it was assumed that the judgments would be valid for the consequent main experiment which was conducted in London. The ratings in the main experiment to be reported below confirmed this assumption.

Method

Participants. 46 students at Facoltà di Scienze della Formazione, Catania participated voluntarily for no reward.

Materials. 40 nouns were used based in part on the previous materials, and four properties were generated for each noun, which we deemed in advance to be mutable, central, necessary and categorical for the noun. Materials and instructions were in Italian, but

English translations are given here.

Procedure. Each participant received instructions followed by a list of nouns each with 4 properties to judge. The instructions were as follows:

“In this booklet you will find concepts on the left (in bold) each of which have four property statements next to them. We want you to decide whether each statement is **necessary**, **central** or **mutable** for the relative concept on the left. For each statement please choose one option (necessary, central, or mutable) that you think is most appropriate, by ticking the box on the right”.

Items were displayed thus:

<i>crocodiles</i>	<i>a) have big jaws</i>	<i>Necessary</i> ☐ / <i>Central</i> ☐ / <i>Mutable</i> ☐
	<i>b) have at least one lung</i>	<i>Necessary</i> ☐ / <i>Central</i> ☐ / <i>Mutable</i> ☐
	<i>c) are reptiles</i>	<i>Necessary</i> ☐ / <i>Central</i> ☐ / <i>Mutable</i> ☐
	<i>d) are dangerous</i>	<i>Necessary</i> ☐ / <i>Central</i> ☐ / <i>Mutable</i> ☐

Instructions included definitions of Necessary, Central and Mutable as follows, (with examples):

“A property of a concept is **mutable** if it can change without having any effect on other characteristics describing that concept.....

A property of a given concept is **central** when, in some way, it causes or explains another characteristic of the same concept...

A property of a given concept is **necessary** if it is always true.....

The order of nouns was randomized, and the order of the four properties for each noun was balanced. The task took about 10 minutes.

Results

Responses were scored using a three-point scale of Mutability (1=necessary, 2=central,

3=mutable). Overall, the statements with categorical properties were chosen as least mutable followed by the necessary, central and mutable ones (means on the scale from 1 to 3 were respectively 1.34, 1.72, 1.89 and 2.46). The 7 quadruples with the least differentiation between properties were discarded, leaving 33 sets. Since the preselected central and necessary properties were not well differentiated, showing considerable overlap, the 4 sentences for each concept were reduced to 3 by selecting either the central or the necessary property (whichever was closest to half way between the categorical and the mutable sentences). There were therefore 3 levels of mutability in the experiment: “mutable”, “central”, and “category”, with mean mutability ratings (and standard deviations) of 2.46 (.28), 1.84 (.25) and 1.34 (.21) respectively.

Experiment 3b: Main Study

Method

Participants. Seventy-two undergraduate students at City University London participated voluntarily.

Materials. Six sentences for each noun concept were constructed depending on whether the noun was modified or not, and on whether the property was mutable, central or categorical. Each booklet contained 11 sentences for each level of mutability, with materials rotated across the 3 levels of mutability across 3 booklets. Presence or absence of the atypical modifier was manipulated between-subjects, so each booklet with unmodified head nouns was matched by a second booklet in which the head nouns were all modified. Each subject noun occurred only once in each booklet. Target sentences were randomly embedded in 33 filler sentences, some clearly true, others clearly false, and others borderline true. About half of the fillers had modified subject nouns. Target sentences can be found in Appendix B.

Design. Two factors were manipulated, Mutability within-subjects with 3 levels, and

Modification between-subjects. Both factors were within-items.

Procedure. Each participant completed a booklet in a classroom setting. Sentences were rated on a 10-point likelihood rating scale as in Experiment 2. The task took about 15 minutes.

Results

Mean likelihood ratings are shown in Figure 2. Unmodified sentences were judged more likely than modified, and category sentences were judged more likely than central sentences, which in turn were judged more likely than mutable. ANOVA was run with factors of Modifier (present or absent) and Mutability of the property. Both main effects were significant (for Modifier, $\text{Min } F'(1, 70) = 34.4, p < .001$, and for Mutability, $\text{Min } F'(2, 140) = 38.9, p < .005$). There was no evidence of an interaction (F_1 and F_2 both < 1). Again contrary to expectations, the size of the modifier effect was the same at each of the three levels of centrality for the unmodified concept. Remarkably, even category membership statements (pigeons are birds) were considered less likely when the subject noun was modified (solitary pigeons are birds).

A number of additional analyses were undertaken to explore the data. Examination of individual items revealed no obvious bad items or other trends discernible that could have been responsible for masking an interaction. Furthermore, fillers were reliably answered – for example the clearly true filler statements had mean ratings of around 9 out of 10. We tried dichotomizing the data by treating ratings of 1-5 as False, and 6-10 as True, but likelihood of a True rating defined in this way showed exactly the same pattern of results and lack of interaction as before.

A power calculation was made for a planned contrast between the modified and unmodified conditions for the size of the difference between ratings for category and mutable

properties. The observed difference in the effect was only 0.12. Power was estimated to be over 90% to detect a difference in the size of effect of as much as 0.75 on the 10 point scale.

Discussion

The results of Experiment 3b again provided no evidence that mutability affects the size of the modifier effect, even when properties with carefully pretested levels of mutability were used. Remarkably, statements about category membership were reduced in likelihood in just the same way as other properties. The fact that category statements were also affected by the modifiers suggests that neither the Prototype Combination nor the Non-Inheritance accounts of the effect are correct, and that in fact the modifier effect reflects some quite general consequence of the pragmatics of the experimental task, leading to reduced confidence in all sentences relating to modified concepts.

Before accepting this conclusion, our next experiment considered whether the lack of an interaction could be explained by the type of rating that was used. Perhaps rating the “likelihood” of something being true on a 10 point scale invites people to signal the unfamiliarity of the modified head noun phrase by giving a lower rating, even when they believe the statements to be actually true. In Experiment 4 we therefore tried a different, more qualitative response scale.

Experiment 4

Experiment 4 used an ordinal response scale with just 4 possible responses labeled: “necessary, important, replaceable, not true”, instead of the 1-10 numerical rating. The aim was to discourage a mapping of familiarity on to the rating scale, and to encourage participants to consider the actual truth of the statements. We expected that a statement that was considered “necessary” for an unmodified head noun should continue to be necessary for its modified form (Hampton, 1987), whereas other properties may shift down the scale when

the subject noun is modified. In order to reduce the overall sense of otherworldliness about the list, we also chose filler sentences so that 50% of them used typical modifiers, as in Connolly et al.'s original study.

Method

Participants. Seventy-two undergraduate students at City University, London participated for reward by completing a booklet in a classroom setting.

Materials. We used the same materials as in Experiment 3b changing only the fillers and response scale as described above.

Design. The design was the same as in Experiment 3b, with modification between-subjects and type of property (category, central and mutable) within-subjects.

Procedure. Each booklet contained instructions and 3 pages of sentences. Next to each sentence was a choice of 4 responses, *necessary*, *important*, *replaceable*, and *not true*. A participant had simply to circle whichever of the four they felt best applied to the sentence. For example if they believed it to be necessary that pigeons are birds, they should circle the “necessary” response. Instructions explained the meaning of the four terms, using examples as in Experiment 3a. The task took about 15 minutes.

Results

Although the response scale was only ordinal, for a simple test of the interaction, and to compare the results with the previous experiment, the responses were mapped to a 1-4 numerical scale and averaged. Results are shown in Figure 3. Once again unmodified sentences received responses further up the scale of necessity/truth than modified ones, and the ordering of type of sentences of category, central and mutable was again seen. Both main effects were significant (for Modifier, $Min F'(1, 70) = 35.9, p < .001$, and for Mutability, $Min F'(2, 140) = 148.0, p < .005$), but again there was no interaction between the two factors.

Analysing the data by frequency of responses led to the same conclusion. For example, category statements about unmodified concepts were considered necessary 63% of the time, but the same statements about modified concepts were judged necessary only 47% of the time. Similarly, unmodified statements were judged “not true” 3% of the time, but this rose to 12% for modified statements. The lack of interaction in Experiment 3 was therefore not solely attributable to the response scale used.

Discussion

At this point, three studies have failed to show any influence of mutability on the modifier effect. People were as likely to downgrade the likelihood of a category membership statement (roosters/stupid roosters are birds) as they were the likelihood of a central property (roosters/stupid roosters have wings) or a mutable property (roosters/stupid roosters crow at dawn). These results suggest that the source of the modifier effect is unlikely to be found in the semantic interaction of modifier and head noun, and is more likely to reflect an overall non-specific bias in confidence for any sentence with an atypically modified head noun.

In Experiment 5 we searched for the elusive interaction by giving the same participants both modified and unmodified versions of the same sentence. In the studies reported so far, a participant only saw one or the other version of each sentence. The use of rating scales can be anchored differently when the surrounding context changes, and although we always used at least 50% filler sentences that were identical in each booklet, our participants were never confronted with the contrast between modified and unmodified versions of the same sentence within the same booklet. Suppose that in the first half of a booklet, you judge that “roosters are birds” is very likely to be true. When faced with “stupid roosters are birds” later on in the booklet, you may now consider that the statement is still just as likely. On the other hand a mutable property for “stupid roosters” such as “crow at dawn” may show less influence of

having seen the earlier unmodified sentence, so that the long-awaited interaction may now appear. To reduce the length of the booklet, we dropped the middle level of “central” properties, and just compared mutable with category properties in a 2x2 design.

Experiment 5

Method.

Participants. Thirty-two participants took part as unpaid volunteers. They were recruited from the residents of a student accommodation facility in London, and all had fluent English.

Materials. Thirty-two head nouns and atypical modifiers were selected based on previously used materials, together with a category and a mutable property for each. Four booklets were constructed, each notionally divided into four quarters. In the first quarter were 16 target sentences, four in each condition. These same 16 targets were then repeated in the third quarter of the booklet, but with the modifier condition changed. Thus the 8 sentences presented as unmodified in the first quarter were presented as modified in the third quarter, and those that were modified were presented in the third quarter as unmodified. The same procedure was applied to the second and fourth quarters of the booklets. A further 32 filler sentences were randomly interspersed throughout the booklet, (8 in each quarter), so that there was always a minimum lag of 32 sentences between the two different versions (modified and unmodified) of the same sentence. As in previous experiments, the fillers were a mix of modified and unmodified sentences some true and some false.

Design and Procedure. The design was 2x2 within-subjects and within-items. Materials were divided into four sets which were then rotated across the four conditions across four booklets. Ratings were collected as in Experiment 2.

Results.

Figure 4 shows the results. The left panel shows the mean ratings for those sentence pairs

which appeared in the booklet with the modified version first. Thus for example, people first rated “stupid roosters are birds/crow at dawn” and then later on “roosters are birds/crow at dawn”. In this order of presentation, the previous results were replicated. There were significant main effects of modifier and of centrality, and no hint of an interaction. The right hand panel shows the same results but for the other order – where the unmodified sentence was rated first, followed by the modified version. Here a much reduced effect of the modifier can be seen. Comparing the two panels, unmodified sentences represented by the black filled bars were unaffected by being presented first or second, with mean ratings of around 9 for category and 7.77 for mutable properties. On the other hand modified sentences were rated more highly when presented second ($M = 8.56, 7.19$) than when presented first ($M = 7.73, 6.13$).

Three-way ANOVA with within-subjects factors of order, modifier and mutability confirmed this interpretation of the results. There were main effects of order ($\text{Min } F' (1, 61) = 9.8, p < .005$), mutability ($\text{Min } F' (1, 62) = 28.4, p < .001$), and modifier ($\text{Min } F' (1, 61) = 26.7, p < .001$), and a significant interaction of modifier and order ($\text{Min } F' (1, 59) = 5.16, p = .027$). No other interactions were significant ($\text{Min } F' < 1$). The breakdown of the interaction by order showed that although the modifier effect was stronger in one order than the other, there was still a significant modifier effect for each order (Modified-Unmodified, $\text{Min } F' (1, 57) = 30.54, p < .001$, and Unmodified-Modified, $\text{Min } F' (1, 61) = 5.01, p < .05$). Notably however there was no evidence at all for any interaction between the mutability and modifier factors in either order. (The same was true if the first and second halves of the booklet were analyzed separately.)

Discussion.

The search for conditions under which the modifier effect would show an interaction with

mutability failed once again in Experiment 5. We had supposed that if exposed to both modified and unmodified versions of the same sentence, people would become aware of the contrast involved, and that this might trigger the kind of semantic considerations that could give rise to the interaction. On the other hand, the manipulation we used did have an interesting effect. When the modified sentence was presented *after* the unmodified version, then the modifier effect was significantly reduced (from 1.43 to 0.53). Modified sentences were rated as more likely when they were in the second half of the booklet. We are led to conclude even more strongly that the modifier effect is probably driven not by the semantics of the task, but by some other pragmatic consideration involving the unfamiliarity or oddness of the head noun phrases. Having established that (say) *roosters* are part of the common discourse, then the phrase *stupid roosters* is more likely to indicate a proper subset of roosters, and hence the subset's properties (both categorical and mutable) are more likely to be considered the same as its parent class.

To test this account, our final study attempted to manipulate the pragmatics of the rating task more directly, by asking participants in one condition to read a short piece of text (a “wiki”) about each modified noun phrase before doing the ratings. Attention to the text was ensured by having a short memory test at the end of the procedure. The aim of the text was to establish that the subset named by the modifier-noun combination was a bona fide class, with the words being used in a literal sense. Thus each text introduced the modified concept in a context that made it clear that it was of the same kind as the unmodified concept. In addition, directing attention to the later memory test reduced the pragmatic burden on the ratings themselves, so that the trivial truth and falsehood of many of the sentences should be recognized in an unproblematic way.

Experiment 6

Method

Participants. Eighty students at City University London participated in the experiment, 20 randomly allocated to each condition.

Materials. Booklets were constructed based on 16 concepts, each with an atypical modifier and two critical properties, one categorical and one mutable. In addition two filler properties were created for each concept, one possibly true (e.g. pigeons are grey) and one most probably false (pigeons can smile). The booklets were laid out with an instruction cover sheet, followed by all four properties listed for each of the concepts in turn. The four sentences were in a new random order for each concept. For example:

pigeons are birds
pigeons are grey
pigeons live in parks
pigeons can smile

Half of the concepts in each booklet were unmodified, as in the example here, and half were modified (e.g. *Tibetan pigeons are birds*). Modified versus unmodified sentences were rotated across booklets, so that half the participants judged the modified version and half judged the unmodified version for each concept. The same order of the 16 concepts was used in each booklet. The second factor was the presentation of a wiki-context. These contexts were obtained from Google searches of the web. Half the booklets provided a brief context for each concept, and half had no context. For example, for pigeons, the following text appeared before the list of four sentences:

Tibetan pigeons mate for life and rear their broods together, although if one dies the other will take a new mate. Once the simple nest is built, the female lays an egg and then another a day or so later. Once the eggs hatch, both parents feed the young squabs.

The context did not provide a direct answer to the truth of any of the four sentences, but established that the modified concept was non-idiomatic. The same context was provided for both modified and unmodified versions of the concepts. Thus, across 4 booklets, each concept appeared in 4 sentences, one unmodified without context, one unmodified but with a context, one modified without context, and one modified with a context. Further examples of context passages are to be found in Appendix C.

Design and procedure. Three factors were manipulated in the design. The provision of a context was a between-participants factor. Provision of context was deliberately confounded with telling participants to expect a memory test for the context material, in order to encourage them to read it carefully. Modification of the concepts was a within-participants within-items factor. Finally, two types of property were used, category and mutable. This sentence factor was within participants and within items. In addition each concept was paired with two filler sentences as described above.

For the No Context condition, the coversheet contained the following instructions:

“The following pages contain 16 groups of sentences. Each group contains 4 sentences. The numbers 1 through 10 appear to the right of each sentence. We would like you to circle the number to the right of each sentence that you feel best captures how likely *you* think it is that the sentence is true. 1 = very unlikely and 10= very likely (the numbers in between represent intermediate levels of certainty).”

Examples were given of a set of sentences concerning Paris with a clearly true statement, two possibly true statements and a false statement, and the circling of numbers was illustrated. The instructions on the cover sheet for the Context condition were adapted from these as follows:

“The following pages contain 16 groups of sentences, each of them preceded by a piece of information taken from “wikipedia” or some newspaper. The numbers 1 through 10 appear ...”

Following the illustrations of true and false sentences about Paris, the instructions continued:

“After this task we will test your memory about the information that appeared before each group of sentences. We want you to say if they are true or false. For instance, with the sentence “Paris is one of the world's leading business centers”, if you read the information in the previous exercise, you should circle TRUE”

Results

Performance in the memory test was satisfactory, showing good attention was paid to the text. Likelihood ratings were averaged across participants and items for each condition. Twenty-two missing data points (<1%) were replaced by the mean across participants for that item and condition. Means are displayed in Figure 5. The left panel of the figure shows that without a context there was a somewhat greater modifier effect for Mutable ($M = 0.9$) than for Category ($M = 0.6$) sentences. The right panel shows the results of adding in a context and memory test. The interaction between mutability and the modifier was enhanced, with modifier effects of 0.8 for mutable and only 0.2 for category sentences.

A 3-way ANOVA was conducted on the data with factors of Modifier (present or absent), Mutability (mutable vs category sentences) and Context (Context provided or absent). There were significant main effects of Modifier (Min $F'(1,35) = 13.3$, $p < .001$) and Mutability (Min $F'(1,19) = 10.8$, $p < .005$), and significant two-way interactions between Mutability and Context (Min $F'(1,43) = 4.3$, $p < .05$), and between Mutability and Modifier (Min $F'(1,64) = 4.53$, $p < .05$). No other effects approached significance ($F < 1$).

The two significant two-way interactions may be understood as follows. The interaction

of Mutability with Context reflects the fact that in the Context condition, the plausibility of the Category statements (regardless of modification) increased (from 8.54 to 8.80), while the plausibility of the Mutable statements decreased (from 7.79 to 7.36). (Neither shift was significant alone on a t-test). Thus provision of a context increased the difference in likelihood between Category and Mutable properties. The second interaction, that of Mutability with Modifier, is the effect that failed to emerge in any of the previous experiments. The modifier effect was less overall for the Category statements (mean 0.4) than for the Mutable statements (mean 0.9).

Was it the provision of the context that led to the emergence of the interaction between the modifier effect and mutability? Although the three-way interaction was not significant, it was found that when analyzed separately, the interaction in question was significant in the Context condition ($F_1(1,78) = 7.8, p < .01, F_2(1,15) = 6.6, p < .05$), but not in the No-Context condition ($p > .2$ in each analysis). In fact the interaction was twice as great in the Context condition. Similarly, when planned comparisons were used to identify the presence of modifier effects in each condition, they were found to be present in both category and mutable statements in the No Context condition ($t(39) > 3.0, p < .01$ for both subjects and items analyses), but only present in the mutable statements in the Context condition ($t(39) > 3.0, p < .01$ for subjects and items for mutable statements, $t_1(39) = 1.36, t_2(15) = 0.64$, for category statements).

A final check was made to test whether the interaction of modifier and mutability observed in the Context condition could be owing to a ceiling effect. Perhaps the modified category rating (8.71) in that condition was already too close to ceiling for the unmodified category rating (8.90) to go any higher. To rule out a ceiling effect, a median split was applied to the average of the modified and unmodified category statement ratings for the 16

sentences in the Context condition. The analysis was re-run separately for the 8 higher rated sentences, and the 8 lower rated sentences. If a ceiling effect was generating the interaction, then the modifier effect should be seen for the lower rated sentences, but not for the higher, which would be constrained by the ceiling. For the higher rated category sentences unmodified sentences were rated at 9.75 versus 9.20 for the modified, giving a modifier effect of +0.55. For the lower rated category sentences mean ratings were 8.04 (unmodified) vs 8.22 (modified). So in fact the modifier effect (-0.18) was completely absent for the lower rated sentences where the effect should have been stronger according to the ceiling effect hypothesis. The interaction of Mutability with Modifier was still significant for the lower rated sentences on the items ANOVA ($F(1,7) = 5.84, p < .05$), but not for the higher rated sentences (where the ceiling effect would have generated an interaction). There was therefore no evidence that a ceiling effect was responsible for the interaction in the Context condition.

Unfortunately the power of the study was insufficient to obtain the significant three-way interaction that would have most clearly demonstrated the influence of the context on the interaction of mutability and modifiers. (A problem exacerbated by a non-significant interaction in the same direction seen in the No Context condition). However the pattern of the results, coupled with the repeated failure to find the interaction in previous studies in this series, strongly suggest that the provision of a context (and memory test) was the decisive factor.

Discussion

After a long series of experiments, we were finally able to demonstrate a significant interaction in which the modifier effect was seen in mutable statements but not in category statements. Although the 3-way interaction was not significant, there was clear evidence that this interaction was present in the context condition, while in the No context condition the

effect was not significant. Similarly, planned t-tests showed that the modifier effect was no longer significant for category membership statements when a context was provided. It is therefore reasonable to conclude that it was the provision of the context (and other changes in procedure introduced in this condition) that led to the elimination of a modifier effect specifically for category membership statements.

The overall effect of providing a context in the form of a short wiki was to increase confidence in the category membership of the concepts, while at the same time decreasing confidence in possession of the mutable properties. It is therefore plausible to conclude that a major component in the modifier effect when no context is provided is simply suspicion about the compositionality of the modified noun phrase. If people question whether a jungle raven is a bird, then they must logically also question whether it is a raven. Given the wiki context that establishes the existence and nature of the modified concepts as being the right kind of thing, then suspicions about category membership disappear. On the other hand people are consequently less certain about whether the mutable property should now apply.

Experiment 6 introduced a number of procedural changes, each of which may have played a role in obtaining the Modifier x Mutability interaction. In addition to providing a context, we also listed four sentences together for each concept, and provided a memory test as an explanation of the aims of the experiment. Further research may tease apart which of these was the most critical factor. Our aim here was simply to establish the existence of a procedure in which the elusive interaction could be made to appear. In that aim, the experiment was successful.

It may be considered that providing participants with a snippet of knowledge will naturally bias them into giving the desired result. However it should be noted that the wiki texts never mentioned the properties being rated, their sole function being to establish that

the subclass named by the modified noun phrase was in the right general domain, so that there was no reason to suspect an idiomatic or figurative use of language. By showing that in this instance the modifier effect does after all affect mutable properties more than categorical ones provides a much needed boundary condition to contrast with the failure to find the interaction in the previous experiments. Whether the effect can still be obtained with more subtle or less direct textual contexts is a question for future research.

General Discussion

Our research began with a desire to understand the modifier effect. Why is it that people rate the likelihood of generic sentences as lower when an atypical modifier is applied to the subject? Three possible answers were considered. One was to consider that concepts combine classically and, additionally, that prototypical features are not inherited directly by the modified concepts. When two concepts are combined, as when an adjective modifies a noun, rules of compositionality require that the resulting complex concept is composed just of the meanings of the two component concepts and their mode of combination. Hence a BROWN COW is just brown and a cow. Once the concept of BROWN COW has been formed compositionally, one has to look once more into the world and discover the prototypical properties of this new set, which may or may not show inheritance of the prototypes of the individual concepts. This was our Non-Inheritance account.

A second explanation was that the effect was the result of processes of prototype combination as outlined in models such as Hampton (1987, 1988) or Smith et al. (1988). When a concept is modified, according to this approach, the prototype information of each concept is integrated into a novel representation or composite prototype. There is a change in the set of property information stored with that concept in memory, leading to a shift in the weights associated with any particular property and consequently with its judged likelihood.

In more complex conceptual combinations, the process can lead to both the loss of some attributes and the emergence of new attributes (Hampton, 1987; Medin & Shoben, 1988), but in the case of the simple independent modifiers used here, the inheritance of features could be expected to be fairly direct.

It was argued that these two accounts differ in their predictions about how central and mutable properties would respond to the modifier effect. While the Non-Inheritance account would be consistent with the finding that both properties should suffer an equal reduction in likelihood, the prototype combination view predicts a greater effect for mutable properties. Round 1 went to the Non-Inheritance account, when Experiment 2 showed equivalent modifier effects for the two kinds of properties.

Both of these accounts however would predict that categorical properties of a concept should be less likely to show a modifier effect. The Non-Inheritance account would have it that one knows for certain that a brown cow is brown and that it is a cow. Hence anything that follows with equal certainty should be equally unaffected by the modification. People know that cows are mammals. That is not part of the prototype of cow, in the sense that a prototypical property is a defeasible property which contributes to differences in typicality within the class. So people should judge “cows are mammals” and “brown cows are mammals” as equally likely, whereas “cows go moo”, and “brown cows go moo” should differ in likelihood.

Neither of these accounts can therefore explain the results of Experiments 3 to 5 in which it was found that in every case the effect of a modifier was to reduce the likelihood of category and mutable properties in equal measure. We therefore need to turn to a third possible explanation of the effect, namely pragmatics. Connolly et al. (2007) themselves noted that pragmatics was the most likely explanation for why even a typical modifier can

produce the effect. In their study, sentences like “quacking ducks can swim” were considered less likely than the unmodified “ducks can swim”. To explain the modifier effect for atypical modifiers however they appealed to “a further set of pragmatic-inferential processes that draw on general knowledge of the world”, without providing a convincing account of why the effect should work in the direction of reducing likelihood across the board, or why the effect should increase with atypicality.³ If these processes are largely knowledge dependent and unpredictable, then there should be no systematic relation between the properties of unmodified and modified concepts.

Hampton et al. (2009; see also Jönsson & Hampton, 2011) showed that when asked to account for their judgments, people did in fact frequently refer to pragmatic factors. There is something odd or ill-formed about sentences that state general truths about more specific subsets (see also Estes & Glucksberg, 1999; Springer & Murphy, 1992). In the present experiments, a pragmatic explanation would point to the unfamiliarity of the atypically modified concepts. Jungle ravens and albino crocodiles are not familiar categories. People’s response is apparently to treat them with a cautious (if not healthy) suspicion. Not only do they doubt whether albino crocodiles have tails and four legs, they are also inclined (to an equal extent) to doubt whether they are reptiles at all. There is no other explanation for the fact that category statements were just as strongly affected by modifiers as were mutable statements.

As evidence for this pragmatic account, we can emphasize two particular results. In Experiment 5 people saw both modified and unmodified versions of the same sentence. In this case, it was apparent that the second time around, they were more confident in the status

³ In their discussion, Connolly et al. make the argument that people’s knowledge of language and its use would guide them to assume that modified concepts will differ from unmodified concepts in relevant ways. They suggest that Default to Prototype is a poor bet to make. We have made an argument to the contrary in Jönsson and Hampton (2008).

of the modified concepts. That is to say that both category and mutable properties were now rated more highly for the modified concepts, although still slightly below the ratings for unmodified concepts. It can be argued that having first made a judgment about an unmodified concept (ravens or crocodiles) participants became more willing to accept the modified concepts as being subsets of these concepts. The unmodified concepts established the domain within which the modification could be seen to lie. The second result that points to a pragmatic account was the effect of the manipulation introduced in Experiment 6. Here we provided half of the participants with a short wiki extract in which the subtypes were referred to as being the right kind of thing to belong in the conceptual class. (We also distracted attention from the judgments by including a recall test for the wiki information). The principle effect of this manipulation was to increase confidence in the category membership of the modified concepts, while decreasing confidence in the mutable properties. Having set people's suspicions about the modified concepts to rest, the semantic effects that had been predicted from the start were finally evident. As both the Non-Inheritance and Prototype Combination accounts would predict, the category membership of a concept was not significantly affected by the presence of the modifier, whereas the mutable prototypical property was significantly reduced in likelihood, which lends some support to the prototype account which incorporates the dependency relations between features that determine mutability,.

There are other results in the literature where people's beliefs in category membership have been shaken. In a study using universal quantifiers, Sloman (1998) showed that people do not fully accept that "all iron is pentavalent" should follow necessarily from the statement "all metals are pentavalent". Similarly Jönsson and Hampton (2006) found that the modifier effect can be found with universally quantified sentences. People considered that "all sofas

have backrests” was more likely to be true than that “all uncomfortable handmade sofas have backrests”. These results suggest that people do not consider statements of class inclusion between concepts to be analytically true. In an early demonstration of this effect, Hampton (1982) found that people would happily endorse a sentence such as “a chair is a kind of furniture”, while at the same time considering that carseats or skilifts were chairs but were not furniture. More recently, Calvillo and Revlin (2005) showed that even for relatively clear categorical bare generic statements (e.g. iron is a metal) there was significant variation across items in the degree to which people felt certain about their truth, and that this variation predicted judgments of argument strength in Sloman’s (1998) deductive categorical inference task. The difference between these earlier results and those reported here is that they largely involved universally quantified statements, which thus invoke deductive reasoning rather than the more inductive style of reasoning involved in our studies.

In conclusion, our studies have served to establish that the modifier effect has multiple causes. A primary cause of the effect lies in the pragmatics of the task, and only once the fact is established that the modified concepts do in fact exist and are not strange idiomatic misnomers, can its semantic basis be investigated. It would appear from our final experiment that given the right pragmatic context, categorical statements are *not* affected by a modifier, whereas mutable statements are. Further tests of the modifier effect should therefore be careful to explore it within a supportive pragmatic context.

References

- Calvillo, D. P. & Revlin, R. (2005). The role of similarity in deductive categorical inference. *Psychonomic Bulletin and Review*, 12, 938-944.
- Clark, H. H. (1973). The language-as-fixed-effect fallacy: A critique of language statistics in psychological research. *Journal of Verbal Learning and Verbal Behavior*, 12, 335-359.
- 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.
- Estes, Z. & Glucksberg, S. (1999). Relevance and Feature Accessibility in Combined Concepts. In M.Hahn & S. C. Stoness (Eds.), *Proceedings of the Twenty-first Annual Conference of the Cognitive Science Society* (pp. 149-154). Mahwah, N.J.: Erlbaum.
- Fodor, J. A. & Lepore, E. (2002). *The Compositionality Papers*. Oxford: Clarendon Press.
- Hampton, J. A. (1982). A demonstration of intransitivity in natural categories. *Cognition*, 12, 151-164.
- 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. (1997). Emergent attributes in conceptual combinations. In T.B.Ward, S. M. Smith, & J. Viad (Eds.), *Creative Thought: An Investigation of Conceptual Structures and Processes* (pp. 83-110). Washington DC: American Psychological Association 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* (pp. x). Oxford: Oxford University Press. (in press)
- 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.
- Jönsson, M. L. (2008) *On Compositionality: Doubts about the Structural Path to Meaning* Doctoral Dissertation, Lund University
- Jönsson, M. L. & Hampton, J. A. (2006). The Inverse Conjunction Fallacy. *Journal of Memory and Language*, 55, 317-334.
- Jönsson, M. L. & Hampton, J. A. (2008). On prototypes as defaults. (Comment on Connolly, Fodor, Gleitman and Gleitman, 2006). *Cognition*, 106, 913-923.
- Jönsson, M.L. & Hampton, J.A. (2011). The modifier effect in within-category induction: Default inheritance in complex noun phrases. *Language and Cognitive Processes*, in press.
- Medin, D. L. & Shoben, E. J. (1988). Context and Structure in Conceptual Combination. *Cognitive Psychology*, 20, 158-190.
- 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. (1981). On the adequacy of prototype theory as a theory of concepts. *Cognition*, 11, 35-58.
- Osherson, D. N. & Smith, E. E. (1997). On typicality and vagueness. *Cognition*, 64, 189-206.
- Rips, L. J. (2001). Necessity and Natural Categories. *Psychological Bulletin*, 127, 827-852.

- Smith, E. E., Osherson, D. N., Rips, L. J., & Keane, M. (1988). Combining Prototypes: A Selective Modification Model. *Cognitive Science*, 12, 485-527
- Sloman, S. A. (1998). Categorical inference is not a tree: The myth of inheritance hierarchies. *Cognitive Psychology*, 35, 1-33.
- Sloman, S. A. (2005). *Causal models: How we think about the world and its alternatives*. New York: Oxford University Press.
- Sloman, S. A. & Ahn, W. K. (1999). Feature centrality: naming versus imagining. *Memory & Cognition*, 27, 526-537
- Sloman, S. A., Love, B. C., & Ahn, W. K. (1998). Feature centrality and conceptual coherence. *Cognitive Science*, 22, 189-228.
- Springer, K. & Murphy, G. L. (1992). Feature availability in conceptual combination. *Psychological Science*, 3, 111-117.

Author Notes

This research was conducted in part while the second author was supported by the European Erasmus scheme. Correspondence should be addressed to the first author at hampton@city.ac.uk.

Appendix

A) Materials (N = 20) selected for use in Experiment 1b..

MODIFIER	NOUN	MUTABLE PROPERTY	CENTRAL PROPERTY
speckled	catfish	have whiskers	can breathe underwater
dry	cellars	are dark	are found in basements.
Brazilian	doves	are white	have wings
male	hamsters	live in cages	need to drink water to survive
inexpensive	limousines	are long	have wheels.
blue	napkins	are made of paper	can be used for wiping one's mouth
bitter	nectarines	are juicy	grow on trees
acrylic	parkas	are warm	can be used as clothing
solitary	penguins	live in cold climates	have feet
Tibetan	pigeons	live in parks	have wings
furry	pigs	live on farms	have legs
jungle	ravens	are black	have wings
handmade	saxophones	are made of brass	require air to produce sound

baked	seaweed	is green	grows in water
itchy	shirts	have buttons	can be worn on torso
uncomfortable	sofas	are found in living rooms	are used for sitting on.
domestic	storks	have long legs	need to eat to survive
painted	thimbles	are made of metal	can be worn on one's fingers
futuristic	wagons	are used by pulling them	have wheels
Namibian	zebras	are fast	have four legs

B) Materials (N = 33) used for Experiments 3b and 4. Expts 5 and 6 used subsets of these materials, with very minor modifications.

MODIFIER	NOUN	MUTABLE PROPERTY	CENTRAL PROPERTY	CATEGORY
dangerous	caterpillars	have many legs	crawl to move around	are insects
striped	catfish	have whiskers	contain bones	are fish
heavy	coins	are round	can be used for buying things	are money
albino	crocodiles	are dangerous	have at least one lung	are reptiles
flightless	doves	are white	have a beak	are birds
imported	dresses	can be elegant	are made of fabric	are clothing

native american	gloves	are made of wool	have fingers	are clothing
Scandinavian	guitars	are made of wood	can be played	are musical instruments
hairless	hamsters	live in cages	are kept as pets	are rodents
miniature	kettles	can be made of plastic	can hold water	are kitchenware
inedible	lambs	are white	are warm blooded	are mammals
Mongolian	limousines	are expensive	are long	are vehicles
vine-grown	nectarines	are sweet	have stones	are fruit
hunting	parkas	are warm	are man-made	are clothing
oval shaped	pearls	are white	are smooth	are jewels
solitary	pigeons	live in parks	need to eat to survive	are birds
long-haired	pigs	live on farms	are omnivorous	are mammals
gas-powered	refrigerators	are used for storing food	are less cold than a freezer	are kitchen appliances
homegrown	rhubarb	is used for pies	grows in the ground	is a plant
antique	rifles	are dangerous	shoot bullets	are weapons
unintelligent	roosters	crow at dawn	have a heart	are birds
old	saxophones	are made of brass	can be used to play music	are musical instruments

ceremonial	scarves	are made of wool	are flexible	are clothing
Baltic	seaweed	is green	grows in sea water	is a plant
opulent	shacks	are made of wood	can be used for storage	are buildings
bright pink	shirts	are made of cotton	are made for wearing	are clothing
short	skyscrapers	are made of concrete	have at least one entrance	are buildings
handmade	sofas	are comfortable	are found in living rooms	are furniture
Korean	squirrels	eat nuts	have bushy tails	are rodents
short-legged	storks	are black and white	have had mothers	are birds
purple	strawberries	can be eaten	have seeds	are fruit
8-wheeled	wagons	can be loaded	are used for carrying things	are vehicles
speckled	zebras	are fast	are warm blooded	are mammals

C) Examples of wiki-stories used in Experiment 6. (These were presented without the initial modifier-noun phrase, which is included for information).

Speckled Catfish: Catfish have deeply forked tails. When small, their smooth-skinned bodies are usually speckled; however, these spots can disappear in older exemplars. They have a small dorsal fin with stiff spine standing high on their back.

Antique Rifles: A rifle is designed to be fired from the shoulder, with a barrel that has a helical groove or pattern of grooves ("rifling") cut into the barrel walls. The origins of rifling are difficult to trace, but some of the more antique rifles seem to have occurred in Europe during the fifteenth century.

Opulent Shacks: In Australia, particularly in Tasmania, shacks were originally located on crown land such as along river banks. They were roughly built with no legal title on the land they were located on. Now, there are quite grand and opulent shacks often used during the summer season.

Solitary Penguins: Penguins in New Zealand inhabit coastal forests and neighboring southern islands. Unique in appearance and behavior, these solitary creatures have experienced population decline in the last 50 years due to habitat loss and predation by introduced species.

Baltic seaweed: German marine biologists are doing their bit to improve Germany's culinary reputation with a campaign to get Baltic seaweed recognized as a delicacy. A seaweed wine is already retailing at 22,50 euros (\$28) a bottle.

Hairless Hamsters: Hamsters are stout-bodied creatures, with tails much shorter than body length and have small ears, short stocky legs, and wide feet. They vary in color from white to shades of gray and black; they can even be hairless.

Figures

Fig 1. The Effect of an Atypical Modifier on Rated Likelihood of Central and Mutable Sentences in Experiment 2.

Fig 2. The Effect of an Atypical Modifier on Rated Likelihood of Category, Central and Mutable Sentences in Experiment 3b.

Fig 3. The Effect of an Atypical Modifier on Rated Likelihood of Category, Central and Mutable Sentences based on Categorical Response Scales in Experiment 4

Fig 4. The Effect of an Atypical Modifier on Rated Likelihood of Category and Mutable Sentences in Experiment 5. Left and Right Panels Show Different Orders of Rating the Modified and Unmodified Sentences

Fig 5. The Effect of an Atypical Modifier on Rated Likelihood of Category and Mutable Sentences in Experiment 6. Left and Right Panels Show No Context and Context Conditions.

Figure 1




Figure 2




Figure 3




Figure 4

Order of Judgment of Modified and Unmodified Sentences




Figure 5

Sentence Likelihood With and Without a Wiki Context

