



Cognitive Science (2012) 1–27

Copyright © 2012 Cognitive Science Society, Inc. All rights reserved.

ISSN: 0364-0213 print / 1551-6709 online

DOI: 10.1111/j.1551-6709.2012.01254.x

Conceptual and Linguistic Representations of Kinds and Classes

Sandeep Prasada, Laura Hennefield, Daniel Otap

Hunter College, CUNY

Received 17 December 2010; received in revised form 23 August 2011; accepted 6 December 2011

Abstract

We investigate the hypothesis that our conceptual systems provide two formally distinct ways of representing categories by investigating the manner in which lexical nominals (e.g., *tree*, *picnic table*) and phrasal nominals (e.g., *black bird*, *birds that like rice*) are interpreted. Four experiments found that lexical nominals may be mapped onto kind representations, whereas phrasal nominals map onto class representations but not kind representations. Experiment 1 found that phrasal nominals, unlike lexical nominals, are mapped onto categories whose members need not be of a single kind. Experiments 2 and 3 found that categories named by lexical nominals enter into both class inclusion and kind hierarchies and thus support both class inclusion (*is a*) and kind specification (*kind of*) relations, whereas phrasal nominals map onto class representations which support only class inclusion relations. Experiment 4 showed that the two types of nominals represent hierarchical relations in different ways. Phrasal nominals (e.g., *white bear*) are mapped onto classes that have criteria of membership in addition to those specified by the class picked out by the head noun of the phrase (e.g., *bear*). In contrast, lexical nominals (e.g., *polar bear*) specify one way to meet the criteria specified by the more general kind concept (e.g., *bear*). Implications for the language–conceptual system interface, representation of hierarchical relations, lexicalization, and theories of conceptual combination are discussed.

Keywords: Language–conceptual system interface; Conceptual representation; Hierarchical relations; Kind representations; Class representations; Conceptual combination

Category representations are central to human cognition; however, not all categories are created equal. Categories differ in the degree of arbitrariness and coherence they display (Keil, 1989; Markman, 1989; Murphy & Medin, 1985; Shipley, 1993). For example, the class of white things exemplifies a relatively arbitrary category wherein members are unlikely to share many properties over and above the fact that they are white. On the

Correspondence should be sent to Sandeep Prasada, Department of Psychology, Hunter College, CUNY, 695 Park Avenue, New York, NY 10021. E-mail: spasada@hunter.cuny.edu

other hand, members of a kind such as bird possess many correlated features and display a rich causal structure. Categories also differ in the manner in which they are expressed linguistically. There is a correlation between linguistic form and arbitrariness of categories, with adjectives picking out relatively arbitrary categories and nouns picking out richly structured categories (Markman, 1989). In this paper, we seek to determine whether differences in linguistic form also map onto two distinct ways of representing categories.

At first glance, the observation that adjectives pick out relatively arbitrary categories and nouns pick out richly structured categories may suggest that our conceptual systems distinguish two ways of representing categories; however, this observation is consistent with two distinct hypotheses concerning the manner in which our conceptual systems represent categories. One possibility, the *Uniform Representation Hypothesis*, is that we have a single way of representing all categories and that the differences between the categories picked out by nouns and adjectives are differences of content—differences in the properties of the members of the categories represented. On this account, kind and class representations are formally identical, differing only in their extensional properties (whether they pick out categories with arbitrary structure or rich correlational and causal structure) and their mapping to linguistic expressions. The second possibility, the *Distinct Representations Hypothesis*, is that our conceptual systems formally distinguish two types of representations of categories: kind representations and class representations, which, in turn, are more or less likely to be mapped onto more and less arbitrary categories. On this account, kind and class representations provide distinct ways of thinking and talking about categories that are not determined solely by their category structures.

The experiments in the present paper distinguish between these two possibilities by investigating whether the linguistic distinction between lexical nominals (e.g., *dog*) and phrasal nominals (e.g., *brown dog*) is mapped onto a conceptual distinction between kind representations and class representations. Phrasal nominals pick out subsets of the categories picked out by their head nouns (i.e., *brown dog* picks out a subset of dogs) and thus they necessarily pick out categories with at least as richly, if not more richly, correlated structure as members of the category picked out by their head nouns (e.g., *dog*). According to the Uniform Representation Hypothesis, all categories are represented in the same manner and kind representations are simply representations of those categories that have rich correlational and causal structure. As such, it predicts that phrasal nominals should be interpreted in the same manner as lexical nominals and that both types of expressions should be interpreted as picking out kinds as they both pick out categories with rich correlational and causal structure. On the other hand, the Distinct Representations Hypothesis allows for the possibility that phrasal nominals such as *brown dog* are mapped onto class representations despite the fact that they pick out categories with rich correlational and causal structure.

We investigated the distinction between lexical and phrasal nominals because nouns that are stored in the lexicon (1)–(3), as well as noun phrases generated via syntactic composition (4)–(5), are used to talk about categories of things.

- (1) Canaries are rare in this part of the world.
- (2) Blackbirds are not always black.
- (3) Picnic tables are often used for purposes other than picnicking.
- (4) Yellow birds are also rare in this part of the world.
- (5) Birds that Jerry's grandmother likes to talk about are likely to be yellow.

Though previous research has sometimes contrasted nouns and adjectives, adjectives are not directly mapped onto category representations, and the correlations reported between adjectives and arbitrarily structured categories are based on considering the properties of the categories picked out by phrasal nominals such as *white thing* that contain the nonspecific noun *thing*.

The experiments presented here seek to determine if lexical nominals are interpreted as picking out categories that are conceived of and spoken about as kinds, while phrasal nominals are interpreted as picking out categories that are not conceived of as kinds but as classes instead. It is important to note that we are not claiming that all lexical nominals name kinds of things, just that we use lexical nominals, not phrasal nominals, to talk about kinds. On the other hand, virtually all phrasal nominals name classes and we use phrasal nominals, not lexical nominals, to talk about classes.

In the rest of the introduction, we briefly describe three ways in which kind and class representations differ that are directly relevant to the experiments in this paper. We return to a more comprehensive consideration of kind and class representations in the general discussion.

1. Class members need not all be of the same kind

The interpretation of minimal pairs of compound nouns and phrases such as *bluebird/blue bird* and *blackbird/black bird* provide some evidence that the linguistic distinction between lexical nominals and phrasal nominals maps onto a distinction between kind and class representations. In these cases, the compounds, which are stored in the lexicon, are interpreted as naming a specific kind of bird, whereas the phrases are interpreted as naming any bird that has the color specified by the modifier. For example, there are many kinds of birds that are black (e.g., crows, ravens, blackbirds, bluebirds that have been painted black, etc.), and thus many individual birds of many kinds may be described as black birds. Consequently, the phrase *black bird* does not apply to a single kind of bird but to birds of any kind that are black. On the other hand, the compound *blackbird* applies to a single kind of bird. As such, the compound provides the means to talk about a category that is a kind, whereas the phrase provides the means to talk about a category that does not constitute a single kind but the class of things that meet the conditions specified by the phrase (that they be black and that they be birds). Experiment 1 investigates whether phrasal nominals, unlike lexical nominals, are generally understood as naming categories whose members need not be of a single kind.

2. Kind and class representations support distinct hierarchical relations

Kind and class representations also differ with respect to the hierarchical relations they support. Shipley (1989, 1993) argues that kind representations enter into both class inclusion and kind hierarchies. In the class inclusion hierarchy, individual dogs such as Fido and Rover are represented as members of the class of dogs as well as the class of animals. In the kind hierarchy, on the other hand, individual dogs are only directly represented as members of the kind dog, and the kind dog as a whole (rather than just its individual members) is represented as a member of the class of animals. One reason for making this distinction is that we make inductions both over classes of individuals (e.g., animals move) as well as classes of kinds (e.g., each kind of animal has a characteristic kind of movement) (Shipley, 1989, 1993; Waxman, Lynch, Casey, & Baer, 1997; Waxman, Shipley, & Shepperson, 1991). Further, generalizations over kinds cannot be straight-forwardly translated into generalizations over individuals. For example, we may agree that it is true that every kind of animal can locomote on its own, but not that every individual animal can locomote on its own (e.g., some individual animals may be paralyzed or born with birth defects and incapable of locomoting on their own). As such, we cannot infer that every individual animal can locomote on its own on the basis of the knowledge that every kind of animal can locomote on its own.

Experiments 2 and 3 investigate whether the concepts named by lexical nominals enter into both class inclusion and kind hierarchies (Shipley, 1989, 1993) and thus support both class inclusion (*is a*) and kind specification (*kind of*) relations. Phrasal nominals, on the other hand, are hypothesized to name class representations and thus are expected to support class inclusion relations, but not kind specification relations (since members of a class need not all be of the same kind). Wierzbicka (1984) also argues for a distinction between the *kind of* and *is a* relations by pointing out that even though we may agree that every policeman is somebody's son, we do not conceive of a policeman as a kind of son.

3. Representation and interpretation of hierarchical relations

Finally, kind and class representations also differ in the manner in which they represent hierarchical relations. Hierarchical relations between more and less general concepts (e.g., BEAR and POLAR BEAR or BEAR and WHITE BEAR) can be represented in two ways. First, the less general concept may specify criteria of category membership that include criteria *in addition* to the criteria of membership supplied by the more general concept. This type of relation is exemplified by the concepts WHITE BEAR and BEAR. As such, a white bear is understood to be something *in addition to* being a bear, namely white. A white bear is understood to have the characteristic of being a bear as well as the characteristic of being white.

Second, the hierarchically related concepts may specify the *same* criteria of category membership but at different levels of *abstractness* of description (e.g., ANIMAL and DOG) or at different levels of *granularity* of description (e.g., DOG and COLLIE). In both of these cases, we understand the lower level representations as indicating a specific way of realizing the criteria supplied by the higher level representation, rather than as providing additional

criteria (Macnamara, 1982, 1999; Prasada, 1999). Accordingly, to be a dog is understood to be a specific way of being an animal rather than to be something in addition to being an animal. Similarly, to be a collie is understood to be a specific way of being a dog rather than to be something in addition to being a dog.

Phrasal nominals are hypothesized to name classes that have criteria of category membership in addition to those specified by the head noun of the phrase. On the other hand, lexical nominals can name kinds that do not specify criteria of membership in addition to those specified by their superordinate kind concepts. Instead, they specify one way to meet the criteria specified by the more general concept. Experiment 4 tests these predictions concerning the manner in which class inclusion and kind specification relations are interpreted.

4. Experiment 1

Experiment 1 investigated whether phrasal nominals are mapped onto class representations and lexical nominals may be mapped onto kind representations by determining whether phrasal nominals, unlike lexical nominals, are understood to name categories whose members need not be of a single kind.

In the experiment, participants were asked to evaluate the truth of statements such as (6)–(9).

- (6) Sunflowers are a single kind of flower.
- (7) #Sunflowers are not necessarily a single kind of flower.
- (8) #Decorative flowers are a single kind of flower.
- (9) Decorative flowers are not necessarily a single kind of flower.

It was predicted that for phrasal nominals, participants should generally deny that the category referred to constitutes a single kind (8) and affirm that they do not necessarily constitute a single kind (9). The reverse pattern was predicted for lexical nominals.

4.1. Method

4.1.1. Participants

Twenty native speakers of English from an introductory psychology course at Hunter College participated in exchange for course credit.

4.1.2. Stimuli

In all experiments, the lexical nominals were subordinate kind terms (e.g., *collie*; *monarch butterfly*) and the phrasal nominals were adjective noun phrases (e.g., *brown dog*). As such, they both picked out lower level categories that are a subset of a basic level category. We chose adjective–noun phrases as our phrasal nominals as they are more likely to pick out kind or kind-like categories than phrases in general (e.g., the flowers John saw last year). As such, they represented a conservative choice and biased against finding a difference between phrasal and lexical nominals.

One hundred lexical nominals that we judged as likely to be familiar to participants were chosen for the study. Fifty were compounds that contained the relevant basic-level term (e.g., *monarch butterfly*), and 50 did not—47, because they were not compounds (e.g., *col-lie*), and 3 compounds that did not contain the relevant basic-level noun (*golden retriever*, *german shepherd*, *greyhound*). We refer to the first 50 items as the *repeated noun* items and the latter 50 as the *non-repeated noun* items. We included both types of items as they may potentially display differences in the ease with which they support different hierarchical relations (Experiments 2 and 3) or the manner in which the hierarchical relations are represented (Experiment 4). The items covered a range of domains, including natural kinds (animals and plants), artifacts (including tools and furniture), and vehicles, and a range of typicalities.

One hundred phrasal nominals were generated using the basic-level nouns that corresponded to the lexical nominals. The adjectives for the phrasal nominals were taken from the 8,817 adjectives of the Kucera-Francis written frequency list accessed through the MRC Psycholinguistic Database. Randomly chosen adjectives from this list were paired randomly with the basic-level nouns with two constraints. The first constraint was that the adjective–noun pairing had to create a possible description of an actual instance of the term. Therefore, adjective–noun phrases such as *hungry table* were omitted, since it is clearly impossible for an actual table to possess such an attribute as the ability to experience hunger. Few adjective–noun combinations needed to be omitted due to this constraint. The second constraint was that the adjectives had to be words that participants were likely to be familiar with. That allowed for phrasal nominals such as *decorative flower*, but not *scoriaceous flower*. Many low-frequency adjectives had to be excluded for this reason.

All the items are listed in the Appendix. In Experiment 1, these items were used to generate sentences of the form displayed in (6)–(9). Thus, each lexical nominal was used to generate a sentence that stated that the category of things picked out constituted a single kind (6) as well as a statement that the category of things picked out did not necessarily constitute a single kind (7). Parallel statements were generated using the phrasal nominals.

4.1.3. Procedure

Each of the 400 statements was presented one at a time in a randomly generated order to each participant using PsyScope software. Participants were given directions asking them to rate the “extent to which the statement strikes you as being true” on a 7-point scale ranging from “definitely not true” to “definitely true.” The participants were then presented with 10 warm-up trials to complete before moving on to the target items in the experiment.

4.1.4. Results

A 2×2 repeated measures ANOVA was performed, with statement type (single kind/not necessarily a single kind) and nominal expression (lexical/phrasal) as the factors, and the participants’ ratings as the dependent variable. The mean ratings given by participants are presented in Fig. 1. Both participant and item analyses were performed.

The main prediction was that there should be a significant interaction between statement type and nominal expression such that single kind statements should receive higher

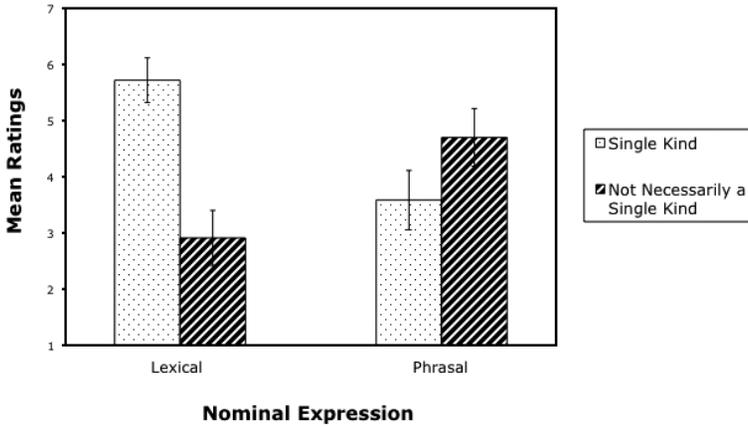


Fig. 1. Results for Experiment 1. Error bars show standard errors.

ratings than not necessarily single kind statements for lexical nominals, but the opposite should be true for phrasal nominals. This interaction was demonstrated, $F_1(1, 19) = 25.4, p < .001$; $F_2(1, 198) = 993.6, p < .001$. The ANOVA also revealed significant main effects for statement type, $F_1(1, 19) = 6.2, p = .022$; $F_2(1, 198) = 184.1, p < .001$. There was a main effect for nominal expression only in the item analysis $F_2(1, 198) = 20.9, p < .001$.

Tests of simple main effects showed that for lexical nominals, the single kind statements ($M = 5.7$) received significantly higher ratings than the not necessarily a single kind statements ($M = 2.9$), $F_1(1, 19) = 29.8, p < .001$; $F_2(1, 99) = 2.5, p < .001$. In contrast, for phrasal nominals, the single kind statements ($M = 3.6$) received significantly lower ratings than the not necessarily a single kind statements ($M = 4.7$), $F_1(1, 19) = 4.7, p = .044$; $F_2(1, 99) = 2.1, p < .005$.

Additional analyses were performed to determine if the difference between lexical and phrasal nominals may be due to differences in the complexity of linguistic form since the phrases were necessarily complex and contained the associated basic-level noun, whereas this was not necessarily the case for the lexical nominals. A 2×2 ANOVA was performed using data only from the lexical nominals that were compounds that repeated the basic level noun (e.g., *monarch butterfly*) and their matched phrasal nominals (*unfriendly butterfly*), with nominal expression and statement type as the factors, and participants' ratings as the dependent variable. The results paralleled those for the full data set, suggesting that the differences between lexical and phrasal nominals are not due to differences in their linguistic complexity. There was a significant interaction between nominal expression and statement type, $F_1(1, 19) = 22.6, p < .001$; $F_2(1, 98) = 300.1, p < .001$, a significant main effect of nominal expression, $F_1(1, 19) = 5.2, p < .034$; $F_2(1, 98) = 7.4, p < .008$, and a significant main effect of statement type, $F_1(1, 19) = 5.2, p < .034$; $F_2(1, 98) = 64.2, p < .001$. Tests of simple main effects showed that for lexical nominals, the single kind statements ($M = 5.6$) received significantly higher ratings than the not necessarily a single

kind statements ($M = 3.0$), $F_1(1, 19) = 25.9$, $p < .001$; $F_2(1, 49) = 450.1$, $p < .001$. For phrasal nominals, the single kind statements ($M = 3.7$) received lower ratings than the not necessarily a single kind statements ($M = 4.6$). This was marginally significant in the subject analysis and significant in the item analysis, $F_1(1, 19) = 3.7$, $p = .070$; $F_2(1, 49) = 36.2$, $p < .001$.

4.1.5. Discussion

Experiment 1 provides evidence that phrasal nominals are mapped onto class representations and lexical nominals may be mapped onto kind representations by showing that phrasal nominals, unlike lexical nominals, are understood as referring to categories of things that need not be a single kind of thing. The judgments involving necessity are especially interesting as they do not simply reveal factual generalizations about the extensions of kind and class concepts. With respect to lexical nominals, participants rejected the possibility that the instances need not constitute a single kind. This is what is expected if lexical nominals are mapped onto kind representations, as it is an intrinsic property of kind representations that the things that they apply to constitute a single kind. On the other hand, class representations only specify the description that members of a class must satisfy, leaving open the possibility that the members of the class are not of a single kind. In principle, it is always possible that instances of different kinds can satisfy the same description and thus be members of the same class. The results of Experiment 1 suggest that participants are aware of this difference in the formal properties of the conceptual representations that underlie lexical and phrasal nominals. As such, the experiment provides evidence for the Distinct Representation Hypothesis.

One difference between the lexical and phrasal nominals in the present experiment is that the lexical nominals were familiar, whereas the phrasal nominals were likely unfamiliar expressions. Though this difference may have influenced the data, it seems unlikely that it can account for the full pattern of responses. This conclusion is bolstered by the results of Experiment 3 which suggest that differences in the interpretation of lexical and phrasal nominals cannot be simply due to differences in familiarity.

The lack of any effects of repeatability of noun indicates that lexicalized compounds and non-compounds are both mapped onto categories that are judged to necessarily be of a single kind. This contrasted with participants' judgments concerning phrasal nominals. These results demonstrate an important difference in the formal characteristics of the conceptual combinations that underlie lexical compounds and those that underlie adjective–noun phrases and phrases more generally. This difference in the formal properties of the conceptual representations that underlie phrasal and lexical nominals has not been previously noted. In fact, focusing on content, Murphy (1988) argues for a uniform theory of conceptual combination for adjective–noun phrases and noun–noun compounds. The present findings suggest that any uniformity in the processes by which the content of conceptual combinations is specified needs to be complemented with a distinction in the formal properties of the class and kind representations that underlie phrasal and lexical nominals.

5. Experiment 2

Experiment 2 investigated whether phrasal nominals are mapped onto class representations and lexical nominals tend to be mapped onto kind representations by investigating the hierarchical relations they support. It was hypothesized that lexical nominals would support both class inclusion (*is a*) and kind specification (*kind of*) relations, as kind representations enter into both class inclusion and kind hierarchies (Shipley, 1989, 1993). On the other hand, phrasal nominals are expected to support class inclusion relations, but not kind specification relations, since, as Experiment 1 demonstrates, members of a class need not all be of the same kind.

To test this hypothesis, participants were asked to judge the truth of statements stating a class inclusion relation (10) and (12) or a kind specification relation (11) and (13) between the categories named by lexical and phrasal nominals and their corresponding basic-level categories.

- (10) Sunflowers are flowers.
- (11) Sunflowers are a kind of flower.
- (12) Decorative flowers are flowers.
- (13) #Decorative flowers are a kind of flower.

The key prediction was that the difference between class inclusion and kind specification statements would be greater for phrasal nominals which denote classes than for lexical nominals which denote kinds.

5.1. Methods

5.1.1. Participants

Thirty native speakers of English from an introductory psychology course at Hunter College participated in exchange for course credit.

5.1.2. Stimuli

The lexical nominals and phrasal nominals from Experiment 1 were used to generate sentences that stated either a class inclusion relation or a kind specification relation between the multiplicities named by the nominals and the corresponding basic-level category.

5.1.3. Procedure

Same as Experiment 1.

5.1.4. Results

A 2×2 repeated measures ANOVA was performed, with hierarchical relation (class inclusion/kind specification) and nominal expression (lexical/phrasal) as the factors, and the participants' ratings as the dependent variable. The mean ratings given by participants are presented in Fig. 2. Both participant and item analyses were performed.

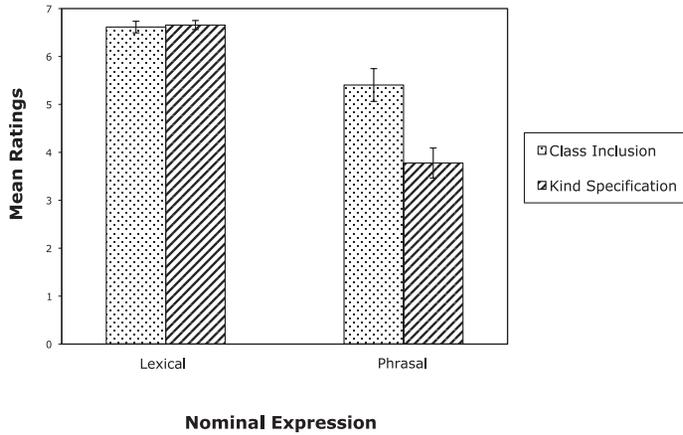


Fig. 2. Results for Experiment 2. Error bars show standard errors.

The main prediction was that there should be a significant interaction between hierarchical relation type and multiplicity expression type with lexical nominals receiving high ratings for both class inclusion and kind specification statements and phrasal nominals receiving significantly higher rating for class inclusion statements as compared to kind specification statements. This interaction was demonstrated, $F_1(1, 19) = 23.0, p < .001$; $F_2(1, 198) = 394.7, p < .001$. The ANOVA also revealed significant main effects for hierarchical relation type, $F_1(1, 19) = 14.6, p = .001$; $F_2(1, 198) = 357.3, p < .001$, and multiplicity expression type, $F_1(1, 19) = 78.1, p < .001$; $F_2(1, 198) = 673.2, p < .001$.

Tests of simple main effects showed no significant difference between class inclusion statements ($M = 6.6$) and kind specification statements ($M = 6.7$) with regard to lexical nominals. However, as predicted, class inclusion statements ($M = 5.4$) received significantly higher ratings than kind specification statements ($M = 3.8$) for the phrasal nominals, $F_1(1, 19) = 19.2, p < .001$; $F_2(1, 99) = 434.8, p < .001$.

Additional analyses were performed to determine whether the difference between lexical and phrasal nominals may be due to differences in the complexity of linguistic form. A 2×2 ANOVA was performed using data only from the lexical nominals that were compounds that repeated the basic-level noun (e.g., *monarch butterfly*) and their matched phrasal nominals (*unfriendly butterfly*), with nominal expression and statement type as the factors, and participants' ratings as the dependent variable. As with the main analyses there was a significant interaction between nominal expression and statement type, $F_1(1, 19) = 20.3, p < .001$; $F_2(1, 98) = 164.3, p < .001$, a significant main effect of nominal expression, $F_1(1, 19) = 77.6, p < .001$; $F_2(1, 98) = 263.1, p < .001$, and a significant main effect of statement type, $F_1(1, 19) = 13.8, p = .001$; $F_2(1, 98) = 166.9, p < .001$. Tests of simple main effects showed no significant difference between class inclusion statements ($M = 6.6$) and kind specification statements ($M = 6.6$) with regard to lexical nominals. However, as predicted, class inclusion statements ($M = 5.5$) received significantly higher ratings than kind specification statements ($M = 4.0$), $F_1(1, 19) = 17.4, p = .001$; $F_2(1, 49) = 224.5, p < .001$ for the

phrasal nominals. These results paralleled those for the full data set, ruling out the possibility that the differences between lexical and phrasal nominals are due to differences in their linguistic complexity.

5.1.5. Discussion

The data show that kind representations, which are expressed via lexical nominals, support both class inclusion and kind specification relations, but classes, which are expressed via phrasal nominals, only support class inclusion relations. These findings suggest that adjective–noun phrases are not interpreted as naming a subkind as they do not support kind specification statements even though they support class inclusion statements. As such, they provide further evidence for the Distinct Representations Hypothesis.

In this experiment, participants gave high ratings to kind terms in the statements that expressed both a class inclusion relation and a kind specification relation. For the phrasal nominals, participants gave a significantly higher rating to the statements that expressed a class inclusion relation. One may wonder why the kind specification statements did not get even lower ratings with the phrasal nominals if participants are not construing the categories as kinds? We suspect this is because the subset relation is very salient and important component of the kind specification relation, even though it does not exhaust the meaning of *kind of* as it does the meaning of *is a* (see Chaffin & Herrmann, 1988 for a review of evidence that semantic relations are not unitary but composed of various elements).

The results of the present experiment are in accord with the hypothesis that phrasal nominals map onto class representations which do not support kind specification relations. However, an alternative interpretation of the findings is that the differences in ratings could be attributed to participants' relative familiarity with the kind terms that constituted the lexical nominals versus the generally unfamiliar adjective–noun phrases that constituted the phrasal nominals.

There are at least two reasons why that alternative explanation is unlikely to be correct. First, lack of familiarity did not prevent participants from endorsing class inclusion relations. Second, Experiment 1 provided evidence that participants do not think phrasal nominals name kinds. Despite the unfamiliarity of the multiplicities denoted by the phrasal nominals, participants agreed that it is not necessary for the multiplicity to be a single kind. As such, it would appear that participants in the present experiment judged the multiplicities denoted by phrasal nominals not to support kind specification relations because they did not conceive of the multiplicity as a kind. Nevertheless, Experiment 3 was designed to test whether familiarity effects were responsible for the results of the present experiment.

6. Experiment 3

Experiment 3 was essentially identical to the previous experiment except that it also included novel noun–noun compounds (e.g., *umbrella flower*). The noun–noun compounds, like the phrasal nominals, were unfamiliar, but, unlike the phrasal nominals, tend to be mapped onto kind representations (Gleitman & Gleitman, 1970) and thus are predicted to

support both class inclusion and kind specification relations. Indeed, many existing kind terms are noun–noun compounds (e.g., *school bus*, *monarch butterfly*, *picnic table*).

Though novel noun–noun compounds can name kinds, they can also be used to describe arbitrary subsets (e.g., *umbrella flowers* can be used to talk about a kind of flower or to refer to the flowers that are under an umbrella on a given occasion). As such, we expect novel noun–noun compounds to support kind specification relations, but not to the same extent as the names of the established subkinds that constituted our lexical nominals. It should also be noted that though there is often a clear hierarchical relation between the kind denoted by the compound and the kind denoted by the head noun of the compound (e.g., school buses are buses and school buses are a kind of bus), this is not always the case. For example, even though *traffic light* is understood to name a certain kind of thing, it is odd to say that traffic lights are lights or that they are a kind of light. For this reason, the statements involving compounds may often sound odd. Further, given that these are entirely novel compounds, it may sometimes be difficult to infer a kind specification relation despite the fact that noun–noun compounds are understood to name kinds, not classes.

The key prediction in the present experiment is that there will be a larger difference between class inclusion and kind specification statements for phrasal nominals than for noun–noun compounds, even though they both involve unfamiliar expressions. This is because phrasal nominals are mapped onto class representations and thus support class inclusion relations but not kind specification relations, but noun–noun compounds are often mapped onto kind representations and thus can support both class inclusion and kind specification relations.

6.1. Methods

6.1.2. Participants

Twenty-one native speakers of English from an introductory psychology course at Hunter College participated in exchange for course credit. Data from one participant were omitted due to the participant's inability to complete the experiment in the allotted time.

6.1.3. Stimuli

In order to keep the length of the experiment equivalent to that of Experiments 1 and 2 and manageable for participants to complete in one 1-h session, 65 of the 100 lexical nominals used in the previous experiments were randomly selected for use in this experiment. The phrasal nominals that corresponded to the 65 lexical nominals were also included. Finally, 65 novel noun–noun compounds were created in a manner similar to that used for creating the phrasal expressions used in Experiments 1 and 2.

Nouns were taken from a randomized list of the 21,340 nouns of the Kucera-Francis written frequency list accessed through the MRC Psycholinguistic Database. Nouns from this list were paired randomly with the 65 basic-level nouns corresponding to the lexical and phrasal nominals following the same two constraints used to create the phrasal nominals (the pairing had to express a possible description of an actual instance of the term, and the nouns had to be words that the participants are likely to be familiar with). In addition, an

attempt was made to match the length of the modifying nouns ($M = 7.631$ letters) with the length of the adjectives ($M = 8.108$ letters) in the phrasal nominals. Each noun–noun compound was used to generate statements concerning class inclusion (14) and kind specification (15). These items were added to the sentences of the same form that were generated by the 65 lexical nominals and 65 phrasal nominals included in the present experiment for a total of 390 statements (see Appendix for a list of the nominal expressions used in this experiment).

(14) Umbrella flowers are flowers.

(15) Umbrella flowers are a kind of flower.

6.1.4. Procedure

Same as that of Experiments 1 and 2.

6.1.5. Results

A 2×3 ANOVA was performed, with hierarchical relation (class inclusion/kind specification) and nominal expression (lexical/phrasal/novel noun–noun compound) as the factors, and the participants' ratings as the dependent variable. The mean ratings for participants are presented in Fig. 3. Both participant and item analyses were performed.

The ANOVA showed significant main effects for hierarchical relation, $F_1(1, 19) = 13.2$, $p = .002$; $F_2(1, 192) = 371.9$, $p < .001$, and nominal expression, $F_1(2, 19) = 76.5$, $p < .001$; $F_2(2, 192) = 526.6$, $p < .001$. There was also a significant interaction between hierarchical relation and nominal expression, $F_1(2, 19) = 16.8$, $p < .001$; $F_2(2, 192) = 115.1$, $p < .001$.

Simple main effects showed no difference between class inclusion statements ($M = 6.7$) and kind specification statements ($M = 6.7$) for lexical nominals. However, for phrasal nominals, class inclusion statements ($M = 5.1$) received significantly higher ratings than

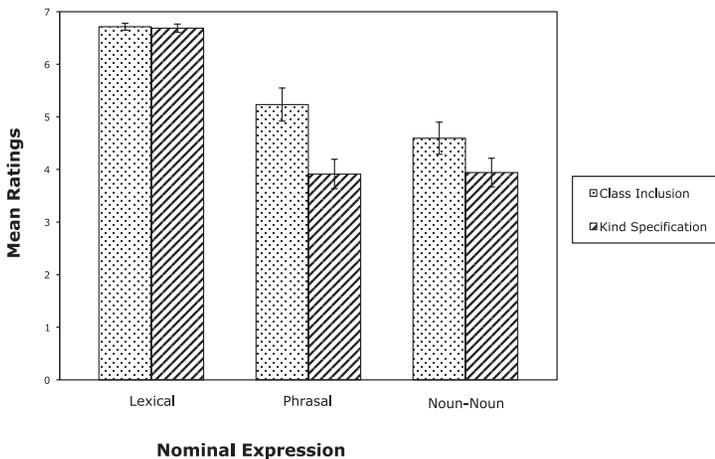


Fig. 3. Results for Experiment 3. Error bars show standard errors.

kind specification statements ($M = 3.8$), $F_1(1, 19) = 17.8, p < .001$; $F_2(1, 64) = 354.5, p < .001$. Finally, novel noun–noun compounds also received higher ratings for class inclusion statements ($M = 4.5$) than kind specification statements ($M = 3.8$), $F_1(1, 19) = 9.096, p = .007$; $F_2(1, 64) = 84.713, p < .001$.

The key prediction in the present experiment is that there will be a larger difference between class inclusion and kind specification statements for phrasal nominals than for noun–noun compounds, even though they both involve unfamiliar expressions. A 2×2 ANOVA with hierarchical relation (class inclusion/kind specification) and nominal expression (phrasal nominals/novel noun–noun compounds) as the factors was conducted to see if this interaction is found.

The ANOVA revealed a significant main effect of hierarchical relation type, $F_1(1, 19) = 15.1, p = .001$; $F_2(1, 128) = 390.7, p < .001$, but no main effect of nominal expression. Crucially, as predicted, there was a significant interaction between hierarchical relation and nominal expression, $F_1(1, 19) = 14.3, p = .001$; $F_2(1, 128) = 44.2, p < .001$. Though the class inclusion statements received a higher rating than the kind specification statements for both the phrasal nominals, $F_1(1, 19) = 17.8, p = .001$; $F_2(1, 64) = 354.5, p < .001$ and the compounds, $F_1(1, 19) = 9.1, p = .01$; $F_2(1, 64) = 84.7, p < .001$, as predicted, there was a larger difference between class inclusion and kind specification statements for phrasal nominals than for noun–noun compounds.

As in the previous experiments, additional analyses were performed to test whether differences in complexity of linguistic form are driving the results. A 2×3 ANOVA was performed using data only from the lexical nominals that were compounds that repeated the basic-level noun (e.g., *grizzly bear*) and their matched phrasal nominals (*spotty bear*) and novel noun–noun compound (*collaborator bear*), with nominal expression and statement type as the factors, and participants' ratings as the dependent variable. As with the main analyses there was a significant interaction between nominal expression and statement type, $F_1(1, 19) = 16.3, p < .001$; $F_2(2, 102) = 52.5, p < .001$, a significant main effect of nominal expression, $F_1(1, 19) = 75.8, p < .001$; $F_2(2, 102) = 269.1$, and a significant main effect of statement type, $F_1(1, 19) = 12.8, p = .001$; $F_2(1, 102) = 171.5, p < .001$.

The key prediction in the present experiment is that there will be a larger difference between class inclusion and kind specification statements for phrasal nominals than for noun–noun compounds, even though they both involve unfamiliar expressions. A 2×2 ANOVA with hierarchical relation (class inclusion/kind specification) and nominal expression (phrasal nominals/novel noun–noun compounds) as the factors was conducted to see if this interaction is found.

As with the main analyses, the ANOVA revealed a significant main effect of hierarchical relation type, $F_1(1, 19) = 15.2, p < .001$; $F_2(1, 65) = 182.0, p < .001$. There was a main effect of nominal expression only for the item analyses, $F_2(1, 65) = 6.1, p < .016$. Crucially, as predicted, there was a significant interaction between hierarchical relation and nominal expression, $F_1(1, 19) = 11.3, p = .003$; $F_2(1, 65) = 14.7, p < .001$, which was due to there being a larger difference between class inclusion and kind specification statements for phrasal nominals than for noun–noun compounds.

6.1.6. Discussion

The results for the phrasal and lexical nominals in the present experiment replicated those of Experiment 2. Furthermore, a larger difference between class inclusion and kind specification statements was found for phrasal nominals than for noun–noun compounds. Because both the phrasal nominals and noun–noun compounds were unfamiliar to the participants, this result suggests that participants’ reluctance to endorse kind specification in comparison to class inclusion statements for phrasal nominals cannot be attributed to the unfamiliarity of the expressions. As such, the results of the present experiment, as well as those of the previous experiment, suggest that phrasal nominals map onto class representations and thus support class inclusion relations, but not kind specification relations. In contrast, lexical nominals may be mapped onto kind representations and thus support both class inclusion and kind specification relations. These results provide further evidence that our conceptual systems provide two formally distinct ways of representing categories.

One might worry, however, that though both the phrasal nominals and noun–noun compounds used in the present experiment were randomly generated and expected to be novel and unfamiliar to the participants, they may nevertheless have differed in their familiarity in two ways. First, it is possible that just by chance the randomly generated expressions contained more noun–noun compounds than phrasal nominals that participants had encountered in their past. Second, background knowledge plays an important role in forming new concepts (Hoffman, Harris, & Murphy, 2008; Murphy & Allopena, 1994), and thus it is possible that the two conditions could have differed in the extent to which participants’ background knowledge led them to think that the likely referents of the two types of expressions actually exist. If either of these scenarios turned out to be true, a familiarity-based explanation of participants’ reluctance to endorse kind specification in comparison with class inclusion statements for phrasal nominals may be possible.

We ran a follow-up to Experiment 3 to investigate the plausibility of this alternative explanation. In the follow-up experiment, we asked a different set of participants to make judgments regarding the actual existence of instances of the expressions (e.g., Do magic tables really exist?), and if they had ever seen the expression before (e.g., Have you ever seen this expression before: *magic table*). The experiment consisted of two blocks, one for each of the judgment tasks, with each block containing all terms from Experiment 3. The two blocks were counterbalanced for order between two versions. Within each block the ordering of the experimental questions was randomized. Seventeen participants took part in the follow-up experiment.

The mean ratings for the two judgments for each type of expression is given in Table 1. Statistical analyses showed that novel noun–noun compounds and the things they refer to

Table 1

Ratings of (i) the likelihood that instances of lexical nominals, phrasal nominals, and noun–noun compounds used in Experiment 3 actually exist, and (ii) the familiarity of the three types of expressions

	Lexical Nominals	Phrasal Nominals	Noun–Noun Compounds
Existence of instances (i)	6.634	4.841	2.947
Familiarity with expression (ii)	6.458	3.598	1.929

were judged as being *less* familiar than phrasal nominals and the things they refer to. These results argue strongly against a familiarity-based explanation of participants' reluctance to endorse kind specification in comparison with class inclusion statements for phrasal nominals.

In sum, the results of Experiments 2 and 3 suggest that phrasal nominals map onto class representations and thus support class inclusion relations, but not kind specification relations. In contrast, lexical nominals may be mapped onto kind representations and thus support both class inclusion and kind specification relations.

7. Experiment 4

Experiment 4 investigated *how* the hierarchical relations supported by the class and kind representations underlying phrasal and lexical nominals are represented. It was hypothesized that phrasal nominals are mapped onto classes that have criteria of application in addition to those of the kind specified by the head noun of the phrase. On the other hand, lexical nominals are mapped onto kinds that do not specify criteria of application in addition to those specified by their superordinate concept. Instead, they specify one way to meet the criteria specified by the more general concept.

Experiment 4 tested these predictions by asking participants to evaluate the truth of statements such as (16)–(19), which characterized being a member of a less general category as either being something in addition to being a member of a more general category (16)–(17) or as a specific way of being a member of the more general category (18)–(19).

(16) #To be a polar bear is to be something in addition to being a bear.

(17) To be a white bear is to be something in addition to being a bear.

(18) To be a polar bear is a specific way to be a bear.

(19) #To be a white bear is a specific way to be a bear.

If phrasal nominals, but not lexical nominals, map onto representations that specify criteria in addition to the criteria specified by the higher level category, participants should endorse statements such as (17), but not (16). On the other hand, if lexical nominals, but not phrasal nominals, may be mapped onto representations that are interpreted as specifying one way to be a member of a higher level category, participants should endorse (18), but not (19).

7.1. Methods

7.1.1. Participants

Forty native English speaking students from an introductory psychology course at Hunter College participated in exchange for course credit.

7.1.2. Stimuli

The stimuli for this experiment were the same as those used in Experiments 1 and 2.

7.1.3. Procedure

Stimuli were presented in a coupled manner such that participants rated both the additional criteria statements (16)–(17) and the specific way statements (18)–(19) for each item one right after the other. Order of presentation was controlled across four versions. Each participant received one of the four versions. A total of 400 sentences were presented one at a time on a computer screen. The sentence pairs were presented in a randomly generated order using PsyScope software. Participants were given directions asking them to rate the “extent to which the statement strikes you as being true” on a 7-point scale ranging from “definitely false” to “definitely true.” The participants were then presented with 24 warm-up trials to complete before moving on to the target items in the experiment.

7.1.4. Results

Mean ratings are given in Fig. 4. A two-way repeated measures ANOVA with Expression (lexical nominal/phrasal nominal) and Interpretation (“specific way”/“something in addition”) as the factors and participants’ ratings as the dependent variable was run. As predicted, there was a significant interaction between Expression and Interpretation, $F_1(1, 39) = 12.0, p < .01$; $F_2(1, 197) = 414.1, p < .001$. The interaction resulted from participants rating the “in addition” interpretation as better for the phrasal nominals than the lexical nominals, $F_1(1, 39) = 9.3, p < .01$; $F_2(1, 197) = 257.0, p < .001$, while preferring the “specific way” interpretation for the lexical nominals, $F_1(1, 39) = 11.2, p < .01$; $F_2(1, 197) = 256.7, p < .001$. The main effects of Expression and Interpretation were not significant in the subject analysis, but the effect of Interpretation was significant in the item analysis, $F_2(1, 197) = 228.1, p < .001$.

To test whether differences in complexity of linguistic form were driving the results, we performed a 2×2 ANOVA using data only from the lexical nominals that were compounds that repeated the basic-level noun (e.g., *monarch butterfly*) and their matched phrasal

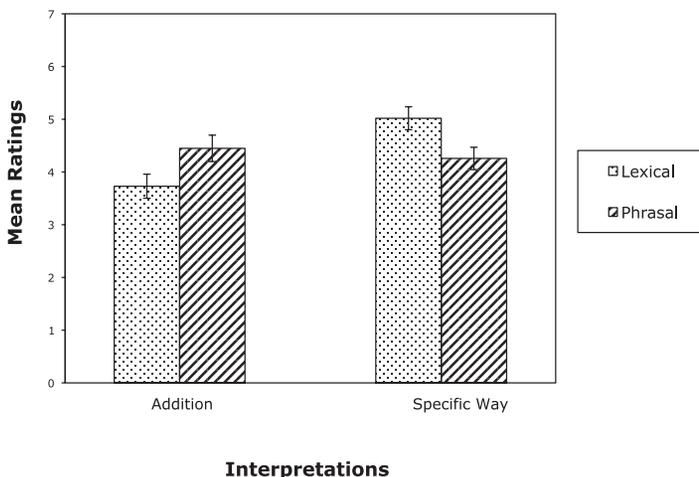


Fig. 4. Results for Experiment 4. Error bars show standard errors.

nominals (*unfriendly butterfly*), with Expression and Interpretation as the factors, and participants' ratings as the dependent variable. As predicted, there was a significant interaction between Expression and Interpretation, $F_1(1, 39) = 10.9, p < .01$; $F_2(1, 197) = 157.5, p < .001$. The interaction was due to participants rating the "in addition" interpretation as better for the phrasal nominals than the lexical nominals, $F_1(1, 39) = 7.5, p < .01$; $F_2(1, 104) = 90.1, p < .001$, while preferring the "specific way" interpretation for the lexical nominals, $F_1(1, 39) = 11.4, p < .01$; $F_2(1, 104) = 110.2, p < .001$. The main effects of Expression and Interpretation were not significant in the subject analysis, but the effect of Interpretation was significant in the item analysis, $F_2(1, 104) = 123.6, p < .001$. This pattern parallels the results of the full data set and thus rules out the possibility that the differences between lexical nominals and phrasal nominals are due to differences in complexity.

7.1.5. Discussion

The results of the experiment provide evidence that our conceptual systems distinguish two ways of representing hierarchical relations. One way of representing hierarchical relations is by representing additional conditions of application for the more specific concept. This way of representing hierarchical relations was preferred for the class representations expressed via phrasal nominals. The other way of representing hierarchical relations is to represent the more general and specific category as supplying the same conditions of application but understanding the specific concept as providing a specific way of realizing the conditions specified by the more general concept. This way of representing hierarchical relations was preferred for the kind representations expressed via lexical nominals.

8. General discussion

The experiments in this paper provide evidence that phrasal nominals are mapped onto class representations, whereas lexical nominals may be mapped onto kind representations. Class and kind representations are formally distinct and provide distinct perspectives from which to represent and reason about categories. Experiment 1 showed that phrasal nominals are mapped onto class representations which may be constituted by things of many kinds, whereas lexical nominals may be mapped onto kind representations which are necessarily constituted by instances of a single kind. Experiments 2 and 3 showed that kind representations are given voice via lexical nominals which support both class inclusion and kind specification relations, whereas phrasal nominals are mapped onto class representations which only support class inclusion relations. Experiment 4 showed that phrasal nominals are mapped onto class representations that have criteria of membership in addition to those of the kind specified by the head noun of the phrase. On the other hand, lexical nominals may be mapped onto kind representations that do not specify criteria of membership in addition to those specified by their superordinate concept. Instead, they specify one way to meet the criteria specified by the more general concept. Together, the experiments present clear evidence that lexical and phrasal nominals are mapped onto formally distinct representations of categories and that these categories support distinct types of hierarchical relations. These

data provide strong evidence that our conceptual systems provide two ways of representing categories. The conceptual distinction between kinds and classes underlies systematic differences in the ways we think and talk about categories and their relations in the context of linguistic tasks of the sort employed in the present research. Studies of the effects of lexicalization suggest that the distinction between kind and class representations extends to contexts in which participants are making inferences about the stability of properties and providing explanations of properties (e.g., Cimpian & Markman, 2011; Gelman & Heyman, 1999). The extent to which the distinction between kind and class representations may be relevant in purely non-linguistic tasks is unknown; however, it is not immediately obvious how the category type manipulation may be made in the absence of language.

In what follows, we consider the implications of the present research for a number of areas, including (i) the interface of the linguistic and conceptual system, (ii) hierarchical representation in commonsense conception, (iii) conceptual combination, (iv) the representation of kinds and classes, and (v) the nature of lexicalization.

8.1. *The linguistic–conceptual representation interface*

The results of the present research suggest a relatively systematic mapping between linguistic form (lexical versus phrasal nominals) and conceptual form (kind versus class representations), but only an indirect mapping between linguistic form and conceptual content (arbitrariness of category structure). It must be noted, however, that the mapping cannot be entirely systematic, as adjectives such as *corporate* and *civil* combine with nouns to name kinds (e.g., corporate lawyers are a kind of lawyer; civil engineers are a kind of engineer) (Prasada, 1992). These adjectives, however, cannot appear as predicates with the same meaning (e.g., *#the lawyer is corporate*; *#the engineer is civil*) and have a number of other properties that distinguish them from other adjectives (Levi, 1978). As such, there is a purely linguistic basis for distinguishing phrases involving these non-predicating adjectives, and thus a basis for phrases involving them to have a different mapping to conceptual representation.

It remains for future research to determine the extent to which the interface between linguistic and conceptual representation can generally be characterized as involving systematic relations between linguistic form and conceptual form, but only indirect relations between linguistic form and conceptual content. Finally, it is important to note that though the mapping between class representations and phrasal nominals is bi-directional, the mapping between kind representations and lexical nominals is uni-directional such that names for kinds are lexical nominals, but not all lexical nominals name kinds.

8.2. *Hierarchical representations in commonsense conception*

The results of Experiments 2 through 4 provide evidence for Shipley's (1989, 1993) proposal that our conceptual systems distinguish class inclusion and kind hierarchies. The experiments also add to Shipley's proposal by addressing the question of whether kinds and classes support distinct hierarchical relations. That is, the question of whether *is a* is

conceptually different from *kind of*. Experiment 4 suggests that class inclusion and kind specification really are distinct hierarchical relations. Class inclusion relations assume different criteria of category membership for hierarchically related classes with the criteria of application for the less general class including the criteria of the more general class along with additional criteria. On the other hand, kind specification relations assume that hierarchically related kinds have the same criteria of membership with the less general kind specifying one way to realize the criteria provided by the more general concept.

To our knowledge, this is the first empirical demonstration of these formally distinct hierarchical relations. Our impression is that despite many substantive differences in how hierarchical representations are represented or computed, most of the literature, with some notable exceptions (e.g., Wierzbicka, 1984; Winston, Chaffin, & Herrmann, 1987), has assumed or criticized and investigated the class inclusion relation rather than the kind specification relation (e.g., Collins & Quillian, 1969; Hampton, 1982; McCloskey & Glucksberg, 1979; O'Connor, Cree, & McRae, 2009; Rips, Shoben, & Smith, 1973; Sloman, 1998). The developmental literature often uses the word *kind* in experiments involving hierarchical relations, and sometimes the kind specification locution *kind of* (e.g., Callanan, 1989; Diesendruck & Shatz, 2001; Markman, Horton, & McLanahan, 1980; Waxman et al., 1991); however, the studies have generally not sought to distinguish class inclusion relations from kind specification relations (but see Shipley, 1989). It will be important for future research to investigate how both of these relations are acquired, represented, and used. Potential differences in the way kind specifications may differ when characterizing the relation between subordinate and basic-level concepts as opposed to the relation between basic- and superordinate-level concepts (Markman & Wisniewski, 1997) also need to be investigated.

8.3. *Implications for theories of conceptual combination*

The finding that phrasal nominals and lexical nominals are mapped onto formally distinct representations also has implications for research on conceptual combination. Current theories do not generally consider the question of whether the result of a conceptual combination is a kind or class representation (e.g., Connolly, Fodor, Gleitman, & Gleitman, 2007; Estes & Glucksberg, 2000; Gagne & Shoben, 1997; Hampton, 1997; Jonsson & Hampton, 2007; Machery & Lederer, 2012; Medin & Shoben, 1988; Murphy, 1988, 1990; Osherson & Smith, 1981; Smith, Osherson, Rips, & Keane, 1988; Wisniewski, 1997; Wisniewski & Love, 1998). The results of the experiments in this paper suggest that whereas the concepts underlying simple nouns and noun–noun compounds are kind representations, the concepts underlying adjective–noun phrases are class representations. Though knowledge-based processes determine the content of conceptual combinations underlying both adjective–noun phrases and noun–noun compounds (Murphy, 1988), it appears that the resulting combinations differ in their formal properties. Theories of conceptual combination must incorporate this new finding because kind and class representations support different ways of thinking, reasoning, and talking about categories and their properties and relations.

It is also likely that the emergent features of class and kind representations are represented, reasoned, and spoken about in different ways. Emergent features are features of a

conceptual combination that are not features of the concepts that enter the conceptual combination (Hampton, 1987; Johnson & Keil, 2000; Murphy, 1988). For example, “eats seals” may be an emergent feature of the concept *WHITE BEAR*. Emergent features of the kind concepts picked out by noun–noun compounds are predicted to be more likely to be represented as aspects of being the relevant kind of thing than the emergent features of the class concepts picked out by adjective noun phrases. Thus, “eats seals” is more likely to be understood to be an aspect of being an ice bear than an aspect of being a white bear. As such, “eats seals” would be represented as having a principled connection to the kind ice bear, but only a statistical connection to the class white bear (Prasada & Dillingham, 2009). This has both linguistic and nonlinguistic consequences (Prasada, 2010; Prasada & Dillingham, 2006, 2009).

8.4. *Representation of kinds and classes*

The phenomena investigated in the present study suggest that rather than simply differing in their content, kind and class representations provide fundamentally different, but complementary, ways of representing categories.

Another fundamental difference between kind and class representations is that kinds are understood to be intrinsically general and are constituted by indefinitely many instances, whereas classes may be finite. Thus, the prime numbers less than 36 constitute a perfectly good class, as do the books that are on my bookshelf at this moment. Thinking of these classes does not require us to also think about non-actual but potential members of the class. Kinds, on the other hand, are intrinsically general and are constituted by their actually existing members as well as indefinitely many potentially existing members. As such, kind representations require a modal distinction between actual and potential existence, whereas no such distinction is implicated by class representations.

In representing a class, we construct a description that specifies the conditions that members of the class must meet in order to belong to the class. In doing so, we make no assumptions as to whether anything exists that meets the description (i.e., how many members the class has, if any). Nor do we posit the existence of the class itself. On the other hand, in representing a kind, we posit the existence of a kind constituted by indefinitely many instances, some of which may actually exist, whereas indefinitely many others exist potentially. Further, kind representations do not specify conditions that must be met in order for something to be an instance of the kind. Instead, they specify certain properties (*k*-properties) that have a principled connection to the kind and thus are understood to be aspects of being that kind of thing (Prasada & Dillingham, 2006, 2009). These properties are expected to generally be true of instances of the kind, are understood to be true of instances of the kind because they are the kind of things they are, and they have normative significance (i.e., instances that lack *k*-properties are understood to be defective or incomplete rather than simply atypical).

As such, kind and class representations embody very different perspectives from which we can represent and reason about multiplicities and their properties. Kind representations posit certain forms of regularity, intelligibility, and normativity, and thus provide relatively fixed and stable perspectives from which to think and reason about things. Kind representations

also support the positing of essences that may be understood to be causally responsible for many properties of instances of the relevant kind (Gelman, 2003; Medin & Ortony, 1989). On the other hand, class representations do not provide a basis for positing essences, since, as Experiment 1 documents, members of a class need not be of the same kind and thus need not share an essence and may differ from one another in fundamental ways. As such, class representations do not underwrite assumptions of uniformity but, instead, provide the means for flexibly classifying things according to the dimensions and criteria of interest.

8.5. *The nature of lexicalization*

As has often been noted (e.g., Birnbaum, Deeb, Segal, Ben-Eliyahu, & Diesendruck, 2010; Cimpian & Markman, 2011; Gelman & Heyman, 1999; Hall & Moore, 1997; Markman & Smith cited in Markman (1989); Wierzbicka, 1986), the lexicalization of a phrase such as *black bird* into *blackbird* is not a merely an act of proceduralization that increases the efficiency with which information may be accessed and used (Anderson, 1982). Instead, a number of conceptual and linguistic changes accompany lexicalization. The results of the present experiments and the ideas considered above suggest that conceptually, lexicalization involves a shift from a class representation to a kind representation. These representations are formally distinct and differ radically in the perspectives they provide for reasoning about multiplicities and their members.

It is interesting to note that the linguistic expressions that map onto kind representations are simpler than those that are mapped onto class representations. This suggests the intriguing possibility that kind representations may also be conceptually and developmentally primitive in comparison to class representations. It is beyond the scope of the present paper to consider these possibilities; however, future research should consider these possibilities.

8.6. *Conclusions and open questions*

Research on concepts has implicitly assumed that we have a single way of representing categories. The experiments reported in this paper suggest that, contrary to this assumption, our conceptual systems formally distinguish two ways of representing categories—as kinds and as classes. These representations are mapped onto the linguistic distinction between lexical nominals and phrasal nominals in English and provide radically different perspectives from which we can reason about categories and their members. Recently, there have been a number of calls for pluralism of one or another sort in the study of conceptual representation (Dove, 2009; Machery, 2009; Weiskopf, 2009). The results of the present paper support representational pluralism with respect to the conceptual representation of categories. An important question for future research concerns the formal and developmental relations of kind and class representations to the set representations that underlie number concepts (Barner, Wood, Hauser, & Carey, 2008; Carey, 2009; Feigenson & Halberda, 2004).

Work on the formal dimension of conceptual representation complements the large amount of work that focuses on conceptual content. By identifying the formal characteristics

of kind representations, the present work, along with related work (Prasada, 2010; Prasada & Dillingham, 2006, 2009; S. Prasada, S. Khemlani, S. Glucksberg & S. Leslie, unpublished data), helps to characterize domain-general principles of conceptual representation. The formal dimension of conceptual representations identifies different perspectives or ways of thinking about things. These perspectives or ways of thinking are formally defined and thus applicable to things from any content domain that may be usefully represented as displaying the relevant formal characteristics. An important question for future research concerns how domain-specific principles of conceptual representation interact with the formal dimension of conceptual representation in the acquisition and use of conceptual knowledge.

Acknowledgments

This work was supported by PSC-CUNY 38, GRTI, and startup grants to the first author. We received infrastructure support from RCMI grant RR03037 from the National Center for Research Resources (NIH) to the Gene Center at Hunter College. Experiments 2 and 3 were part of Laura Hennefield's MA thesis. She is currently in the PhD program in psychology at Washington University, St. Louis, MO.

References

- Anderson, J. R. (1982). Acquisition of cognitive skill. *Psychological Review*, *89*, 396–406.
- Barner, D., Wood, J., Hauser, M., & Carey, S. (2008). Evidence for a non-linguistic distinction between singular and plural sets in rhesus monkeys. *Cognition*, *107*, 603–622.
- Birnbaum, D., Deeb, I., Segal, G., Ben-Eliyahu, A., & Diesendruck, G. (2010). The development of social essentialism: The case of Israeli children's inferences about Jews and Arabs. *Child Development*, *81*, 757–777.
- Callanan, M. A. (1989). Development of object categories and inclusion relations: Preschoolers' hypotheses about word meanings. *Developmental Psychology*, *25*, 207–216.
- Carey, S. (2009). *The Origin of Concepts*. New York: Oxford University Press.
- Chaffin, R., & Herrmann, D.J. (1988). The nature of semantic relations: A comparison of two approaches. In M.E. Evens (Ed.), *Relational models of the lexicon* (pp. 289–334). New York: Cambridge University Press.
- Cimpian, A., & Markman, E. M. (2011). The generic/non-generic distinction influences how children interpret new information about social others. *Child Development*, *82*, 471–492.
- Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior*, *8*, 240–247.
- Connolly, A. C., Fodor, J. A., Gleitman, L. R., & Gleitman, H. (2007). Why stereotypes don't even make good defaults. *Cognition*, *103*, 1–22.
- Diesendruck, G., & Shatz, M. (2001). Two-year-olds' recognition of hierarchies: Evidence from their interpretation of the semantic relation between object labels. *Cognitive Development*, *16*, 577–594.
- Dove, G. (2009). Beyond perceptual symbols: A call for representational pluralism. *Cognition*, *110*, 412–431.
- Estes, Z., & Glucksberg, S. (2000). Interactive property attribution in concept combination. *Memory & Cognition*, *28*, 28–34.
- Feigenson, L., & Halberda, J. (2004). Infants chunk object arrays into sets of individuals. *Cognition*, *91*, 173–190.
- Gagne, C. L., & Shoben, E. J. (1997). Influence of thematic relations on the comprehension of modifier-noun combinations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*, 71–87.

- Gelman, S. A. (2003). *The essential child: Origins of essentialism in everyday thought*. London: Oxford University Press.
- Gelman, S. A., & Heyman, G. D. (1999). Carrot-eaters and creature-believers: The effects of lexicalization on children's inferences about social categories. *Psychological Science*, *10*, 489–493.
- Gleitman, L. R., & Gleitman, H. (1970). *Phrase and paraphrase: Some innovative uses of language*. New York: W. W. Norton & Co.
- Hall, D. G., & Moore, C. E. (1997). Red bluebirds and black greenflies: Children's understanding of the semantics of adjectives and count nouns. *Journal of Experimental Child Psychology*, *67*, 236–267.
- 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. (1997). Conceptual combination. In K. Lamberts & D. R. Shanks (Eds.), *Knowledge, concepts and categories* (pp. 135–162). Hove, UK: Psychology Press.
- Hoffman, A. B., Harris, H. D., & Murphy, G. L. (2008). Prior knowledge enhances the category dimensionality effect. *Memory & Cognition*, *36*, 256–270.
- Johnson, C., & Keil, F. C. (2000). Explanatory knowledge and conceptual combination. In F.C. Keil & R.A. Wilson (Eds.), *Explanation and cognition* (pp. 327–359). Cambridge, MA: MIT Press.
- Jonsson, M. L., & Hampton, J. A. (2007). On prototypes as defaults. *Cognition*, *106*, 913–923.
- Keil, F.C. (1989). *Concepts, kinds, and conceptual development*. Cambridge, MA: MIT Press.
- Levi, J. (1978). *The syntax and semantics of complex nominals*. New York: Academic Press.
- Machery, E. (2009). *Doing without concepts*. Oxford, UK: Oxford University Press.
- Machery, E., & Lederer, L. G. (2012). Simple heuristics for concept composition. In M Werning, W. Hinzen, & E. Machery (Eds.), *The Oxford handbook of compositionality* (pp. 454–472). Oxford, UK: Oxford University Press.
- Macnamara, J. (1982). *Names for things: A study of human learning*. Cambridge, MA: MIT Press.
- Macnamara, J. (1999). *Through the rearview mirror: Historical reflections on psychology*. Cambridge, MA: The MIT Press.
- Markman, E. M. (1989). *Categorization and naming in children*. Cambridge, MA: MIT Press.
- Markman, E. M., Horton, M. S., & McLanahan, A. G. (1980). Classes and collections: Principles of organization in the learning of hierarchical relations. *Cognition*, *8*, 227–241.
- McCloskey, M. E., & Glucksberg, S. (1979). Natural categories: Well defined or fuzzy sets? *Memory & Cognition*, *6*, 462–472.
- Medin, D., & Ortony, A. (1989). Psychological essentialism. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning*. Cambridge, UK: Cambridge University 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. (1990). Noun phrase interpretation and conceptual combination. *Journal of Memory and Language*, *29*, 259–288.
- Murphy, G. L., & Allopena, P. D. (1994). The locus of knowledge effects in concept learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *20*, 904–919.
- Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, *92*, 289–316.
- O'Connor, C. M., Cree, G. S., & McRae, K. (2009). Conceptual hierarchies in a flat attractor network: Dynamics of learning and computations. *Cognitive Science*, *33*, 665–708.
- Osherson, D. N., & Smith, E. E. (1981). On the adequacy of prototype theory as a theory of concepts. *Cognition*, *11*, 35–58.
- Prasada, S. (1992). *Acquisition of adjective meanings: A lexical semantic approach*. Unpublished doctoral dissertation, Massachusetts Institute of Technology, Cambridge, MA.

- Prasada, S. (1999). Names for things and stuff: An Aristotelian perspective. In R. Jackendoff, K. Wynn, & P. Bloom (Eds.), *Language, logic, and conceptual representation: Essays in honor of John Macnamara* (pp. 119–146). Cambridge, MA: MIT Press.
- Prasada, S. (2010). Conceptual representations and some forms of genericity. In F. J. Pelletier (Ed.), *Kinds, things and stuff* (pp. 36–59). New York: Oxford University Press.
- Prasada, S., & Dillingham, E. M. (2006). Principled and statistical connections in common sense conception. *Cognition*, 99, 73–112.
- Prasada, S., & Dillingham, E. M. (2009). Representation of principled connections: A window onto the formal aspect of common sense conception. *Cognitive Science*, 33, 401–448.
- Rips, L. J., Shoben, E. J., & Smith, E. E. (1973). Semantic distance and the verification of semantic relations. *Journal of Verbal Learning and Verbal Behavior*, 12, 1–20.
- Shipley, E. F. (1989). Two kinds of hierarchies: Class inclusion hierarchies and kind hierarchies. *The Genetic Epistemologist*, 17, 31–39.
- Shipley, E. F. (1993). Categories, hierarchies, and induction. *Psychology of Learning and Motivation*, 30, 265–301.
- Sloman, S. A. (1998). Categorical inference is not a tree: The myth of inheritance hierarchies. *Cognitive Psychology*, 35, 1–33.
- Smith, E. E., Osherson, D. N., Rips, L.J., & Keane, M. (1988). Combining concepts: A selective modification model. *Cognitive Science*, 12, 485–527.
- Waxman, S. R., Lynch, E. B., Casey, K. L., & Baer, L. (1997). Setters and samoyeds: The emergence of subordinate level categories as a basis for inductive inference. *Developmental Psychology*, 33, 1074–1090.
- Waxman, S. R., Shipley, S., & Shepperson, B. (1991). Establishing new subcategories: The role of category labels and existing knowledge. *Child Development*, 62, 127–138.
- Weiskopf, D. A. (2009). The plurality of concepts. *Synthese*, 169, 145–173.
- Wierzbicka, A. (1984). Apples are not a ‘kind of fruit’: The semantics of human categorization. *American Ethnologist*, 11, 313–328.
- Wierzbicka, A. (1986). What’s in a noun? (Or: How do nouns differ in meaning from adjectives?). *Studies in Language*, 10, 353–389.
- Winston, M.E., Chaffin, R., & Herrmann, D. (1987). A taxonomy of part-whole relations. *Cognitive Science*, 11, 417–444.
- Wisniewski, E. J. (1997). When concepts combine. *Psychonomic Bulletin & Review*, 4, 167–183.
- Wisniewski, E. J., & Love, B. C. (1998). Relations versus properties in conceptual combination. *Journal of Memory and Language*, 38, 177–202.

Appendix

Lexical Nominals	Phrasal Nominals	Noun–Noun Compounds
(*used in Exp. 3)	(*used in Exp. 3)	(Exp. 3)
iguana*	serious lizard*	activity lizard
gecko*	american lizard*	american lizard
dalmatian*	gothic dog*	anthology dog
greyhound*	uncomfortable dog*	arch dog
bulldog*	alexandrian dog*	atheism dog
daisy*	pristine flower*	bib flower
sunfish*	athletic fish*	candle fish
oak tree*	everlasting tree*	catholicism tree

Appendix (Continued)

Lexical Nominals	Phrasal Nominals	Noun–Noun Compounds
pickup truck*	filthy truck*	chimney truck
bullfrog*	dull frog*	christmas frog
cobra*	indonesian snake*	coliseum snake
polar bear*	spotty bear*	collaborator bear
garbage truck*	pungent truck*	confederate truck
daffodil*	scarlet flower*	consolation flower
hammerheads*	gigantic shark*	container shark
jet*	magnificent airplane*	counsel airplane
rattlesnake*	omnivorous snake*	dart snake
jeans*	posh pants*	decapitation pants
honeybee*	unpolished bee*	decoy bee
pine tree*	albanian tree*	endorsement tree
tree frog*	busy frog*	exercise frog
great white shark*	petite shark*	fancier shark
rottweiler*	unkind dog*	footnote dog
chihuahua*	observable dog*	georgian dog
black widow*	good spider*	good spider
desk chair*	public chair*	hand chair
rocking chair*	japanese chair*	handle chair
sailboat*	returnable boat*	honeycomb boat
mallard*	unclothed duck*	honeydew duck
basketball*	genuine ball*	interview ball
folding chair*	indoor chair*	jumble ball
box turtle*	spellbound turtle*	kosher turtle
chardonnay*	fabulous wine*	lash wine
beluga*	mad whale*	legislature whale
earthworm*	untouched worm*	locker worm
steel*	unexceptional metal*	manipulator metal
goldfish*	southern fish*	member fish
catfish*	unimpaired fish*	navel fish
chainsaw*	cantonese saw*	notification saw
golden retriever*	idealistic dog*	onlooker dog
needlenose pliers*	complementary pliers*	outrage pliers
tugboat*	uninsured boat*	parallelogram boat
owl*	gassy bird*	parasite bird
rose*	undistinctive flower*	penguin flower
poodle*	opaque dog*	pinch dog
sunflower*	desert flower*	planet flower
blue jay*	unofficial bird*	potassium bird
hummingbird*	freakish bird*	prussian bird
picnic table*	unfinished table*	punctuation table
tuna*	wiggly fish*	racer fish
chameleon*	unbeaten lizard*	rescuer lizard
anaconda*	comical snake*	scallop snake
cement truck*	nondescript truck*	scoop truck
barracuda*	cuban fish*	servant fish

Appendix (Continued)

Lexical Nominals	Phrasal Nominals	Noun–Noun Compounds
silver*	splendid metal*	sheriff metal
snapping turtle*	beastly turtle*	sleuth turtle
carnation*	unstudied flower*	snowflake flower
soccer ball*	unsafe ball*	stimulant ball
grizzly bear*	haitian bear*	substructure bear
tow truck*	retrograde truck*	tablet truck
canoe*	traditional boat*	tone boat
maple tree*	prehistoric tree*	turnover tree
tulip*	decorative flower*	umbrella flower
white-tailed deer*	shapely deer*	vagabond deer
clownfish*	gaunt fish*	wallet fish
tarantula*	inedible spider*	waterglass spider
humpbacks	gaelic whale	
school bus	blue bus	
palm tree	urban tree	
penguin	dizzy bird	
ostrich	deaf bird	
dump truck	new truck	
trout	unpaired fish	
yacht	egyptian boat	
taxi	dingy car	
salmon	unreal fish	
blue whale	infectious whale	
mini-van	delicate van	
tennis ball	elizabethan ball	
collie	awkward dog	
semi truck	buyable truck	
minnow	slippery fish	
beach ball	normal ball	
dining room table	magic table	
sparrow	impotent bird	
python	trusty snake	
german shepherd	unavoidable dog	
lily	intact flower	
violet	armenian flower	
coffee table	describable table	
sedan	dampish car	
electric drill	used drill	
bottlenose dolphin	whimsical dolphin	
lilac	bristly flower	
monarch butterfly	unfriendly butterfly	
siamese cat	venomous cat	
kitchen table	roundish table	
chicken	mythical bird	
robin	wicked bird	
flamingo	solitary bird	