

Conceptual combination: Conjunction and negation of natural concepts

JAMES A. HAMPTON
The City University, London, England

The operation of negation on combinations of natural categories was examined in two experiments. In the first, category membership ratings of lists of items were obtained for pairs of concepts considered individually and in two logical combinations: conjunctions (e.g., "Tools which are also Weapons") and negated conjunctions—forms of those conjunctions in which the modifier noun category was negated ("Tools which are *not* Weapons"). For conjunctions, results supported earlier findings of overextension and the geometric averaging of constituent membership values (Hampton, 1988b). Previous findings of concept dominance and noncommutativity within conjunctions were also replicated, both for typicality ratings and for probability of class membership. For negated conjunctions, the pattern of dominance was similar but interacted with order within the conjunction. Negated conjunctions were also overextended. The second experiment explored how the attributes of negated conjunctions were derived from those of the two component concepts. Frequency of generation of attributes expressed positively (*has wheels*) or negatively (*has no wheels*) followed rated frequency in the negated category. The distinctiveness of an attribute to distinguish the complement from the head-noun class was associated with the generation of attributes, particularly when there was relatively high overlap between the two categories.

The study of conceptual combination has recently come to assume considerable theoretical importance for psychological theories of concepts (Hampton, 1997a; Rips, 1995). The prototype theory of concepts (Rosch, 1978; Rosch & Mervis, 1975) postulated that the classification of objects in categories, such as, say, Furniture or Sports, is based on the overall similarity structure of the items composing the category. The theory proposed that the category was represented by a prototype, that is, an idealized representation of the set of attributes positively associated with category membership. Items are judged to belong in the category if they are sufficiently similar to the prototype (and dissimilar from the prototypes of contrasting categories). Against the prototype theory, it has been argued (Fodor, 1994; Osherson & Smith, 1981, 1982) that the lack of any clear set of rules for combining prototype concepts in logical combinations casts grave doubt on the general value of prototype theory as a theory of human concepts. Hence, research on how logical

functions operate on prototype concept categories is of considerable theoretical interest.

The arguments surrounding this issue have been widely aired (Cohen & Murphy, 1984; Hampton, 1987, 1988b, 1991, 1996, 1997a; Jones, 1982; Murphy, 1988; Murphy & Spalding, 1995; Osherson & Smith, 1981, 1982; Rips, 1995; Smith & Osherson, 1984; Thagard, 1983; Zadeh, 1982).¹ The problem addressed in this article is how similarity-based prototype concepts enter into logically constructed complex concepts. In particular, the focus here is on the logical operations of conjunction and negation. To examine these operations, in the studies reported below, a head noun plus relative clause construction was used as a means of expressing conjunction ("Tools which are also Weapons") and a similar construction was used with a negated modifier noun ("Tools which are not Weapons") as a means of studying the effect of negation. The questions raised were, first, to what extent is category membership in such classes predictable from category membership in each of the constituent classes, and second, to what extent are the attributes that are considered descriptive of the negated conjunctions also true of the constituent classes.

Much current theorizing on the formation of conjunctive concepts has arrived at the view that the most fruitful theoretical approach is an *intensional* one (Cohen & Murphy, 1984; Hampton, 1987; Smith & Osherson, 1984; Smith, Osherson, Rips, & Keane, 1988). This is to say that rather than modeling membership in the conjunction extensionally in terms of some function of degree of membership in the constituent classes (see, e.g., Zadeh, 1965, 1982, for the fuzzy logic approach, or Huttenlocher & Hedges, 1994, for the statistical approach), mod-

Part of this research was conducted while the author was on sabbatical leave at Stanford University, whose hospitality is gratefully acknowledged. It was made possible by the award of a fellowship from the Nuffield Foundation and a personal research award from the British Academy. The author also acknowledges support from the French Ministry for Higher Education and the Fulbright Commission Visiting Scholar program. The author's thanks are extended to the following for discussion and advice in the course of the research and in the preparation of this manuscript: Larry Barsalou, Herb Clark, Peter Fonagy, John Gardiner, Doug Medin, Greg Murphy, Greg Oden, and a number of anonymous reviewers. Thanks are also due Tony Heron and Antonia Gongorra for research assistance. Requests for reprints should be addressed to the author at: Psychology Department, City University, Northampton Square, London EC1V 0HB, England (e-mail: j.a.hampton@city.ac.uk).

els of conceptual conjunction should aim to define how the two prototypes (or schemas) representing the two concepts become combined into a modified or composite representation of the conjunctive class. (Dissenting accounts are offered by Chater, Lyon, & Myers, 1990, and Huttenlocher & Hedges, 1994.)

Two intensional models have been developed for concept conjunctions with some degree of detail. Smith et al. (1988) proposed a selective modification model for adjective–noun combinations such as “Red Apple.” In their model, a head noun such as “Apple” is represented by a frame (Minsky, 1975) composed of attributes such as COLOR, SIZE, or TASTE, each of which can take values, such as *red*, *large*, or *sweet*, respectively. Specifically, the representation of the head noun “Apple” would possess an attribute for COLOR which would normally take a range of values—*red*, *green*, *yellow*, and so on—each with an associated number of votes, reflecting its frequency of occurrence as the color of an apple. According to the model, this head-noun frame becomes selectively modified in the combination “Red Apple” by switching all the votes for COLOR to the value *red*, while at the same time increasing the overall weight of color in the determination of similarity to the concept schema.

The second intensional model was proposed by Hampton (1987, 1988b) in order to account for the way people understand the conjunction of two noun concepts in such phrases as “Sports that are also Games” or “Tools which are also Weapons.” Hampton’s composite prototype model for conjunctions proposed that each noun concept was represented by a prototype consisting of a list of attributes or properties. When the concepts are conjoined, then a new *composite prototype* is constructed by merging together the two sets of attributes defining the two constituent noun concept prototypes. This composite list of attributes is then subject to further modification in order to satisfy various constraints, such as Necessity (an attribute that is considered necessary for a constituent is also considered necessary for the conjunction), Impossibility (an attribute that is considered impossible for a constituent is also considered impossible for the conjunction), and Coherence (the composite prototype may not contain two incompatible attributes). A similar formal approach has been suggested by Thagard (1983, 1995).

Apart from being directed at different forms of conjunction, an important difference between the Smith et al. and Hampton models lies in their assumptions concerning the determination of set membership in the conjunctive class. Although Hampton (1987, 1988b) explicitly proposed that membership in the conjunction was determined by similarity of instances to the composite prototype, Smith et al. (1988) chose to limit their model to the determination of *typicality* or representativeness of instances in the conjunction. They recognized that, as constituted, their model failed to pick out a conjunctive concept category that would actually be the logical intersection of the two constituent sets (the same is true of Hampton’s model). Quite simply, a logical intersection requires that membership in the conjunction should depend on the level

of similarity to each constituent independently. Degree of Redness and degree of Appleness, for example, should form independent criteria, both of which need to be achieved for something to count as a Red Apple. However, if membership of the conjunction “Red Apple” is based on *overall similarity* to the conjunctively modified schema, then this independence of criteria will not be possible (see Ashby & Gott, 1988, for discussion of this issue).

For Hampton (1988b), this failure of intensional models to generate intersective conjunctions was taken as a virtue, inasmuch as his data apparently showed that people’s classification of instances in conjunctions was not, in fact, purely intersective, but showed the kind of interdependence predicted by the models. Smith et al. (1988), in contrast, argued that typicality and membership depended on two different types of semantic information. They suggested that concepts might have a core meaning—some central definitional component of the attribute structure associated with the concept—which was used in making class membership judgments (see also Miller & Johnson-Laird, 1976). If this core is defined as a necessary and sufficient set of common elements, then Boolean set logic can be used to determine how concepts combine. Effectively, an object is actually classified as a red apple only if it has the “core” features of both redness and applesness (no pun intended). Typicality judgments, however, would be based on similarity using the full range of prototype attributes, and so would not follow logical intersection. Their model for concept conjunction was therefore explicitly restricted in its scope to intuitions of typicality.

The primary aim of the present research was to replicate and extend the range of data considered by such models by exploring the use of *negation* in conceptual combinations. There has been very little research on how people interpret negated concepts. One obvious reason is that single negated terms have little meaning. People cannot sensibly rate items for their typicality as “not sports.” The category is infinite and indefinitely heterogeneous. Within a conjunctive phrase, however, the task is quite meaningful. Thus, “Games which are not Sports” is a concept for which subjects can sensibly judge the membership and typicality of items. The two experiments presented here explore both the extensional (category membership) and intensional (attribute listing) aspects of negated constituents in a conjunction. The aims of the research are primarily exploratory, as no previous work has looked at these issues within the framework of the current research tradition. Hypotheses can, however, be derived on the basis of earlier work within this paradigm. Hampton (1988b, Experiments 2, 3, and 4) asked people to make three categorization decisions about lists of items. In Stage 1 of the experiment, subjects made decisions first about whether activities on a list were (for example) “sports,” and then whether they were “games.” One week later, the subjects rated the same items again, this time for whether they were “Sports which are also Games” (or “Games which are also Sports”)—an apparently explicit conjunction of the two sets. The main finding was that

membership in the conjunction was not determined by Boolean set intersection. Rather, membership in the conjunction was generally a geometric average of the two constituent membership values. In addition, the criterion for membership in the conjunction was set quite low, with the result that the conjunctions were often overextended. Chess, for example, was judged by many subjects to be in the category "Game which is a Sport," even though they had said before that it was not a sport. The overextension effect occurred across a range of different conjunctions. These inconsistent class membership decisions were interpreted by Hampton (1987, 1988b) as evidence for a similarity-based categorization process in which similarity to the composite prototype was the basis for categorization in the conjunction. Hampton (1988b) used a combined typicality and membership rating scale, in which subjects first decided whether an item was a member of the category or not. If the answer was "yes," they then judged the item's typicality on a 3-point scale. If the answer was "no," they then judged its relatedness to the category on a 3-point scale. The two judgments were then combined into a 7-point scale from +3 (highly typical) to -3 (unrelated).² Hampton (1996) extended these results to visually presented category materials such as cartoon faces or ambiguous colored letter shapes.

Hampton (1988b) also found an asymmetry, or more properly a noncommutativity, in the conjunctions. Converse pairs of conjunctions like "Sports which are also Games" and "Games which are also Sports" had different graded structures in that the regression weight for a constituent concept predicting membership in the conjunction was higher when the concept was in the relative clause qualifier position than when it was the head noun. Over and above this positional effect, in several cases there was also an imbalance between the two constituent concepts. One concept (the dominant concept) tended to have a higher regression weight than the other, regardless of the order of the terms. This dominance effect has since been replicated in studies by Storms, De Boeck, van Mechelen, and Geeraerts (1993) (see also Storms, De Boeck, van Mechelen, & Ruts, 1996).

Consider then what may be expected when the second constituent is negated (as in the class "Sports which are *not* also Games"). On the basis of the previous results (Hampton, 1988b), similar patterns may be expected to occur. The regression weight of the second constituent should of course have a negative sign, since the better an activity is as an example of Games, the worse it should be as an example of "not Games." Otherwise, if the previous results are generalizable, then membership in the conjunctive category should not be predictable simply on the basis of set intersection or set complementation, there should be a greater regression weight for a category when it is in the relative clause qualifier position than in the head-noun position, and there should be dominance effects between the two constituent categories as before. Experiment 1 therefore had the aim of testing the degree

to which the pattern of results obtained with conjunctions could be generalized to negated conjunctions.

Experiment 1 had the secondary aim of testing two predictions of other models concerning conjunctive categorization through the introduction of two methodological changes. First, a between-subjects design was used so that the probability of categorization could be used to provide a more direct test of overextension of conjunctions. If categorization probability for an item is measured independently for two constituent categories and for their conjunction, then, if the conjunction is treated as an intersection of the two constituents, as predicted by classical models of concepts (Armstrong, Gleitman, & Gleitman, 1983; Osherson & Smith, 1981), the probability with which an item is judged to belong in the conjunction should lie within certain limits. Let the measured proportion of the population sampled who believe that the item is in category A be $p(A)$. Then, given $p(A)$ and $p(B)$ for a pair of constituent categories A and B, if the two beliefs are *uncorrelated* in the population, we would expect $p(A \& B)$, the probability that someone places the item in the conjunction "A which are also B," to be the product of the two probabilities $p(A)$ and $p(B)$. There are, however, reasons why the beliefs may be correlated. First, global individual differences in the breadth of category boundaries would give rise to a positive correlation between believing that a given item is in A and believing that it is in B. Second, variation in individual beliefs about the nature of the item being categorized could lead to either a positive or a negative correlation between $p(A)$ and $p(B)$, depending on how much the categories overlap or contrast.

If the two category membership beliefs are fully negatively associated, then the lower limit for $p(A \& B)$ should be zero or $[p(A) + p(B) - 1]$, whichever is the greater.³ Alternatively, if they are fully positively associated, then the upper limit for $p(A \& B)$ should be the *minimum* of the two constituent probabilities. For example, suppose that everyone holds the belief that all pets are animals. Then the probability that an item is classified as a "Pet which is an Animal" would be simply equal to the probability that it was classified as a Pet. Set logic applied to the incidence of beliefs in the population then predicts that $p(A \& B)$ should lie within the limits of zero or $[p(A) + p(B) - 1]$ and the minimum of the two constituent probabilities. To the extent that this constraint is broken, there will be evidence against the view that the psychological representations of conjunctions are based on the logic of set intersection, as proposed by Osherson and Smith (1981). Furthermore, if the binary view of the separation between typicality effects and set membership espoused by Smith and Osherson (1984) is correct, overextension, noncommutativity, and concept dominance effects should be found in the mean rated typicalities, but should be much reduced or even absent from the probability of categorization measure. Use of this design therefore helps to provide a test of Smith et al.'s (1988) the-

oretical position concerning the logical versus nonlogical nature of concept conjunction.

The between-subjects design also helps settle a question about the validity of the earlier studies (Hampton, 1988b), which used within-subjects designs. Hampton (1988b) obtained categorization judgments in a fixed order, presenting constituent categories first, followed by conjunctions. If this order confound was in some way responsible for the findings of overextension in conjunctions, then overextension should be absent with a between-subjects design.

The second methodological change was introduced to provide a test of a modified extensional account of concept conjunctions proposed by Jones (1982; see also Chater et al., 1990; Zadeh, 1982). The basis of this account is the idea that conjunctions are formed as a function of constituent class membership, but that conjunctive membership is *rescaled*, so that the best example of the set intersection becomes the prototype of the conjunctive set. Jones's (1982) conceptual stacks model proposed that people found the best fitting example of the two constituent sets and then based their conjunction on this exemplar. Thus, whatever instance happens to be the best example of a conjunction becomes maximally typical of the class. This proposal could also account for the emergent attributes and inheritance failure seen in Hampton (1987) if the conjunction's attributes are based on this maximal exemplar rather than on a combination of the two sets of constituent attributes. To test Jones's proposal, the materials for Experiment 1 deliberately included the best examples of the conjunctive concepts as generated by a pretest. If his proposal is correct, then at least some of these items should receive maximal typicality ratings for the conjunction (positive or negated) regardless of how typical they are in the constituent sets. Alternatively, if Hampton's intensional model is correct, then it may be possible to find conjunctive classes with "empty centers," where the composite prototype formed from the two constituent prototypes does not correspond closely to any known instance and the "best known example" of the conjunction is therefore still rated as relatively atypical.

EXPERIMENT 1

The primary aim of the first experiment was to explore the effects of negation in conceptual conjunctions. The following two questions were of particular importance concerning these negated conjunctions. First, how will membership in negated conjunctions be related to constituent membership? Will categorization probability for negated conjunctions [denoted as $p(A \& \neg B)$] follow the constraints of set logic? If not, will negated conjunctions be overextended like other conjunctions? On the one hand, given that a conjunction is overextended with respect to its constituents, the corresponding negated conjunction may

turn out to be underextended to a corresponding extent. On the other hand, if overextension is the result of some decision process applying equally to all kinds of conjunction, then negated conjunctions may be equally overextended. Second, when mean membership/typicality values are considered, will the negated conjunctions show effects of head versus modifier order and concept dominance parallel to those shown by their related conjunctions?

The secondary aim was to test the generality of the earlier results with a between-subjects design and with a more comprehensive selection of examples of the conjunctive categories. Two further hypotheses will then be tested: (1) that categorization probability differs from mean typicality in the degree of nonlogical effects observed, and (2) that the most common exemplar in a conjunction is automatically rated as highly typical.

Method

Subjects. The subjects were 120 students at Stanford University who participated for course credit. They were randomly assigned, 20 to each of the six main conditions. A further 16 students participated in the pretest, also for course credit.

Design and Materials. The six pairs of overlapping concepts were taken from categories used in earlier research on this topic (Hampton, 1987, 1988b). The use of the same categories allowed the generality of the previous dominance effects to be tested. The concept pairs were Birds–Pets, Buildings–Dwellings, Furniture–Household Appliances, Sports–Games, Tools–Weapons, and Vehicles–Machines. For each pair of categories, a new list of 20 items was created, containing exemplars of each concept, some in both concept categories, some in just one or the other, and others from the same general domain, but belonging to neither concept category. To ensure that the list contained the best examples of each conjunction, 16 subjects performed an item generation task as a pretest. For each of the six conjunctions, the subjects listed as many examples as possible with no time constraint. (Order within conjunctions was balanced between two subgroups of subjects.) Between 25 and 45 different examples were generated. The 6 most frequently produced examples were then always included in the list of 20 selected items for each pair of categories. Items were typed in a randomly ordered vertical list headed by a category title. There were two lists to a page, and page order was randomized for each subject. Six versions of the booklet were created by varying the category title at the head of the item lists. Each version was given to a different group of 20 subjects. Groups 1 and 2 each had one of the constituent concepts as category title (e.g., Bird for Group 1 and Pet for Group 2). Groups 3 and 4 each had a conjunction of the two concepts (e.g., Group 3 had "Birds which are also Pets," and Group 4 had "Pets which are also Birds"). Groups 5 and 6 had the negated conjunction sets corresponding to the sets of Groups 3 and 4 (e.g., Group 5 had "Birds which are not Pets," and Group 6 had "Pets which are not Birds").

Procedure. The booklet was administered in a group testing session. A cover sheet contained instructions on how the items were to be rated. The rating scale was the same as that used previously (Hampton, 1988a, 1988b). The subjects first had to decide whether an item belonged in the category, choosing a positive rating if it did, a negative rating if it did not, or a zero to indicate a borderline case. The positive and negative ratings each ranged from 1 to 3, corresponding either to the typicality of a category member (+1 = untypical, +3 = very typical) or to the relatedness of a category nonmember (−1 = related, −3 = completely unrelated). If an item

was unknown, the subjects could put a line through it. Those responses were then treated as missing data in the analyses.

Results

Two sets of analyses were performed. The first set used regression equations to analyze how mean rated typicality and the probability of a positive rating in the two kinds of complex concept—the conjunctions and the negated conjunctions—related to the respective values for the constituent concepts. The second set of analyses considered overextension of the same complex concepts and the question of whether the best example of a conjunctive concept automatically received a maximal typicality rating.

Regressions: Typicality. The first analysis used regression statistics to examine noncommutativity and concept dominance effects in membership ratings. Numerical scale values from +3 to -3 were averaged across subjects for each list item, for each of the subject groups.⁴ Four regression equations were calculated for each concept pair, predicting mean ratings for each of the four complex concepts from mean constituent ratings and an interaction term defined as their product. The regression statistics are shown in Table 1.⁵

For the conjunctions, the interaction term was significant in 6 of the 12 equations. Multiple *R* was between .892 and .992, corresponding, on average, to 93% of the variance. The average standard error of prediction was 0.40. For the negated conjunctions, the interaction term was significant in 10 of the 12 equations and multiple *R* ranged from .938 to .996. The average *R* squared was again .93, and average standard error of prediction was 0.45.

The noncommutativity of conjunctions was tested by comparing the absolute value of the regression coefficient for a concept when it was in the qualifier noun position (in the relative clause) with that when it was in the head-noun position. In all 12 positive conjunctions, concepts had higher regression weights in the qualifier position than in the head-noun position, thus replicating previous results (Hampton, 1988b). For the negated conjunction equations, however, there was no overall effect. Half of the concepts had higher regression weights in the qualifier, and half had higher regression weights in the head-noun position. The advantage of modifier over head position was confined to the conjunctions and was not seen in the negated conjunctions.

The second question concerned concept dominance—the observation that weights for particular pairs of con-

Table 1
Regression Statistics for Predicting Mean Rated Membership in Conjunctions and Negated Conjunctions From Membership in Their Constituents in Experiment 1

Head Noun A	Modifier B	Regression Weights				
		A	B	A × B	SE	R
Pets	Birds*	.30	.78	.10	.34	.99
Pets	Not Birds*	.32	-.75	-.11	.21	.99
Birds*	Pets	.65	.42	.14	.40	.98
Birds*	Not Pets	.72	-.36	-.14	.42	.97
Tools	Weapons*	.38	.72	.08	.34	.97
Tools	Not Weapons*	.18	-.76	-	.52	.95
Weapons*	Tools	.62	.50	.13	.48	.96
Weapons*	Not Tools	.78	-.39	-.20	.32	.99
Machines	Vehicles*	.46	.76	-	.42	.97
Machines	Not Vehicles*	.44	-.48	-.15	.37	.97
Vehicles*	Machines	.64	.65	-	.67	.94
Vehicles*	Not Machines	.73	-.63	-.24	.56	.95
Furniture	Household Appliance*	.12	.98	-	.35	.98
Furniture	Not Household Appliance*	.58	-.53	-.16	.34	.98
Household Appliance*	Furniture	.30	.56	-	.48	.89
Household Appliance*	Not Furniture	1.10	-.31	-	.66	.95
Games	Sports*	.38	.82	-	.29	.99
Games	Not Sports*	.85	-.82	-.21	.51	.99
Sports*	Games	.49	.66	.14	.42	.97
Sports*	Not Games	.87	-.93	-.33	.32	.97
Buildings	Dwellings*	.12	.98	-	.24	.99
Buildings	Not Dwellings*	.52	-.55	-.19	.43	.96
Dwellings*	Buildings	.42	.64	.05	.32	.99
Dwellings*	Not Buildings	.84	-.84	-.26	.71	.94

	Average Regression Weights			
	Conjunctions		Negated Conjunctions	
	Dominant	Nondominant	Dominant	Nondominant
Head noun	.521	.292	.839	.481
Qualifier	.841	.571	-.650	-.577

Note—Where a variable did not enter significantly, only a dash is shown. *Dominant concepts in Hampton’s (1988b) study.

cepts could be unequal. Defining the “dominant” concept of each conjunction as the concept with the higher mean regression weight in the earlier study (Hampton, 1988b), it was found that the same concepts were again dominant. (In Table 1, the dominant concepts are asterisked.) Averaging across the two orders of the conjunctions, in every case the dominant concepts had higher absolute regression weights [mean $|b| = .681$, compared with $.431$; $t(5) = 6.6, p < .01$]. For the negated conjunctions, the same was true, with one marginal exception [mean $|b| = .744$ vs. $.529$; $t(5) = 2.3, p < .05$]. The pattern of concept dominance was therefore replicated with new subjects, a between-subjects design, and new lists of items. This replication strengthens the case for the theoretical significance of the dominance effect (Hampton, 1987, 1988b; Storms et al., 1993; Storms et al., 1996), and suggests that its roots lie in the semantics of the particular concepts rather than in any artifact of the lists of items used. It was also shown that dominance was maintained when the qualifier concept was negated.

At the foot of Table 1, summary statistics are presented which show that, for conjunctions, the two effects of dominance and noncommutativity were nearly additive. On the other hand, for negated conjunctions they interacted. For negated conjunctions, the head-noun position showed a strong dominance effect, but in the (negated) modifier-noun position, regression weights for dominant and non-dominant concepts were approximately the same. The interaction was in the same direction in all six concept pairs. Analysis of variance (ANOVA) of absolute regression weights across the six conjunctions, with three repeated measures factors of negation (conjunction vs. negated conjunction), order (head vs. qualifier), and dominance (from Hampton, 1988b), showed significant effects of order [$F(1,5) = 14.13, p < .05$] and dominance [$F(1,5) = 15.82, p < .05$] and a significant three-way interaction [$F(1,5) = 10.96, p < .05$]. A breakdown of the interaction confirmed that, for negated conjunctions, order interacted with dominance [$F(1,5) = 10.33, p < .05$], while, for conjunctions, the interaction was absent ($F < 1$).

Regressions: Classification probability. According to predictions derived from Smith et al. (1988), classification probabilities should show few of the effects found for typicality/membership ratings. The probability of a positive rating was calculated for each of the six subject groups as the number of positive ratings given to an item, divided by the total number of nonzero ratings. (Instructions emphasized the use of zero on the scale to indicate the class-membership boundary. Relatively few zeros were given, and excluding them had no effect on the pattern of results reported.) Regression analyses were repeated using probabilities in place of mean rated typicalities. If there is a dissociation between typicality ratings and category membership, as proposed by Smith and Osherson (1984), then the previous effects such as dominance and noncommutativity should be greatly reduced, or even eliminated, when considering probability of categorization alone.

Regression analysis was run on each pair of categories to predict probability of categorization for each of the four complex concepts, from $p(A)$ and $p(B)$ for the two constituents. An interaction term, defined as $[p(A) - .5] \times [p(B) - .5]$ was also included. The regression statistics are summarized in Table 2, where it may be seen that the pattern of results was indeed very similar to that for typicalities in Table 1. The noncommutativity effect was again seen only in the conjunctions. For negated conjunctions, the effect was not significant. ANOVA of the regression weights confirmed the main points of the previous analysis of typicalities. For conjunctions, the two main effects of dominance and noncommutativity were approximately additive [for the two-way interaction, $F(1,5) = 2.12, p > .2$], whereas for negated conjunctions, the dominance effect was again much stronger for head-noun concepts than for the negated qualifiers [for the two-way interaction, $F(1,5) = 11.22, p < .05$]. The three-way interaction was highly significant [$F(1,5) = 17.09, p < .01$]. In all major respects, although the dominance effects were rather less reliable, the analysis of categorization probabilities produced the same pattern of results as the analysis of mean typicality. The finding that the “nonlogical” effects of concept conjunction, such as noncommutativity and dominance, were present not only in the typicality/membership ratings but also in the probability of categorization is strong evidence against the view proposed by Smith and colleagues (Smith & Osherson, 1984; Smith et al., 1988) that membership in a conjunction should be determined by set intersection applied to the constituent memberships. In their selective modification model, they proposed that typicality in a conjunction would be determined by similarity to a semantic representation of the combined concepts. Yet they also proposed that categorization would not be based on this representation but, rather, would be a logical function of the “core” meaning of each constituent. The finding here that typicality and categorization probability produce parallel results is clearly contrary to this proposal, and instead provides justification for assuming a common basis for typicality and category membership (Hampton, 1988b). Furthermore, there is no obvious way in which models that assume classical core definitions for determining category membership could account for the different systematic effects observed here.

Conjunctive typicality and conceptual stacks. Mean rated typicality was also analyzed to examine Jones’s (1982) conceptual stacks proposal that the best example of a conjunction should receive a maximal typicality rating. To test this effect, the present experiment deliberately included the most commonly generated items for each of the six conjunctions. For five of the conjunctions, the best exemplars of the conjunction were indeed rated as highly typical, but as these exemplars were also very typical members of one or another constituent, they did not provide a clear test of the hypothesis. For example, the three most typical Pet Birds were very typical of the conjunction but were also highly typical as Birds.

Table 2
Regression Statistics for Equations Predicting Probability of a Positive Rating
for Conjunctions and Negated Conjunctions in Experiment 1

Head Noun A	Modifier B	Regression Weights				R
		A	B	A × B	SE	
Pets	Birds*	.24	.72	.43	.08	.98
Pets	Not Birds*	.22	-.79	-.37	.04	.99
Birds*	Pets	.55	.37	.55	.12	.95
Birds*	Not Pets	.57	-.40	-.78	.14	.91
Tools	Weapons*	.56	.73	-	.07	.97
Tools	Not Weapons*	.24	-.67	-	.12	.93
Weapons*	Tools	.70	.55	-	.13	.90
Weapons*	Not Tools	.64	-.40	-1.05	.11	.95
Machines	Vehicles*	.51	.59	-	.11	.94
Machines	Not Vehicles*	.39	-.43	-.88	.13	.90
Vehicles*	Machines	.39	.66	.66	.11	.96
Vehicles*	Not Machines	.52	-.66	-.95	.08	.98
Furniture	Household Appliance*	.19	.92	-	.07	.98
Furniture	Not Household Appliance*	.54	-.53	-	.09	.95
Household Appliance*	Furniture	.25	.60	.57	.09	.89
Household Appliance*	Not Furniture	1.06	-.22	-	.12	.96
Games	Sports*	.29	.64	.65	.08	.97
Games	Not Sports*	.65	-.57	-1.64	.10	.98
Sports*	Games	.34	.57	1.26	.09	.97
Sports*	Not Games	.74	-.93	-1.39	.10	.93
Buildings	Dwellings*	.18	.83	-	.06	.99
Buildings	Not Dwellings*	.38	-.53	-.77	.10	.93
Dwellings*	Buildings	.42	.65	.29	.07	.98
Dwellings*	Not Buildings	.60	-.70	-.94	.22	.84

	Average Regression Weights			
	Conjunctions		Negated Conjunctions	
	Dominant	Nondominant	Dominant	Nondominant
Head noun	.441	.330	.688	.403
Qualifier	.739	.566	-.588	-.552

Note—Where a variable did not enter significantly, only a dash is shown. *Dominant concepts in Hampton's (1988b) study.

The conjunction of Furniture and Household Appliances provided the only good test of the effect, since the best rated members (TV, Hifi, Desk Lamp) were not typical members of either constituent. These exemplars were rated between 0.5 and 1.5 on the scale for the conjunction, so that even the best members of the conjunction were no more than moderately typical of the conjunctive class. However, even excluding those subjects who gave negative ratings for the conjunction, the mean conjunctive typicality ratings for these exemplars were only: TV = 2.3, Hifi = 1.8, Desk Lamp = 2.2. In effect, the conjunctive category had an “empty center.” The possibility of a category with no typical members has important theoretical implications and deserves further study. In the present case, it rules out the conceptual stacks notion (Jones, 1982) by which the most typical exemplar in any category should automatically be seen as maximally typical.⁶

For negated conjunctions, the analysis was more difficult because of the lack of direct ratings for the negated constituent concepts (it is necessary to infer membership in the negated constituent “Not B” from membership in the unnegated constituent “B”). Furthermore, no attempt was made to ensure that the most typical members of the negated conjunction classes were included in the experiment. However, the results suggest that negated con-

junctions may show a much greater “rescaling” effect than the conjunctions. For example, *Dog sled* was rated as a Vehicle, 1.45; as a Machine, -1.55; and as a “Vehicle which is not a Machine,” 2.50. And *Tent* was rated as a Dwelling, 0.75; as a Building, -0.85; and as a “Dwelling which is not a Building,” 2.90. High typicality could be achieved in the negated conjunction with only moderate typicality in the head-noun class and moderate unrelatedness to the modifier.

Overextension. A critical issue arising from the earlier work by Hampton (1988b, 1996) is the overextension of conjunctive categories. Hampton (1988b, 1996) examined overextension by considering how each individual subject categorized each item across the three category decisions—constituent A, constituent B, and the conjunction “A which are B.” Overextension (and underextension) could then be defined at the individual level in terms of the proportion of *inconsistent* response triples made by each subject in categorizing an item in each of the three categories. The between-subjects design adopted here required a different method for assessing overextension. The dependent variable used was categorization probability for each item in the three categories, as estimated from the different groups of subjects. The following issues were then addressed for both conjunctions and negated

conjunctions: First, were the conjunctions overextended? When freed from the constraints of making three decisions about the same item, are people less likely to overextend the conjunction? Second, how will people categorize items with respect to the negated conjunction class—will it also be overextended or will it be correspondingly underextended? If people are classifying according to a logical combination of their beliefs, then the effects of negation on categorization probability should follow set theory as applied to the population of people holding different beliefs. For example, the sum of the categorization probabilities for the two complex concepts [$p(A \& B)$ and $p(A \& \neg B)$] should be no greater than that for the head-noun constituent [$p(A)$] on the grounds that a person cannot consistently hold the belief that an item is simultaneously both in the conjunction and in its negated form.

To assess overextension, expected values of categorization probability for the four complex concepts were derived from categorization probability for the constituents, according to two different functions. The first, the intersection rule (corresponding to Zadeh's, 1982, function for interactive conjunction) defined $p(A \& B)$ as the product of the two constituent values $p(A)$ and $p(B)$. If someone judges an item to belong in a conjunction only if they hold the beliefs that the item belongs in both constituents, and if the incidence of the two beliefs is statistically independent across individuals, then the probability of being in the conjunction should be the product of each constituent probability. The second function was the *minimum* of the two constituent probability values. Even allowing for maximum positive correlation between the incidence of the two constituent beliefs, the expected value of $p(A \& B)$ should never exceed the *minimum* of the two constituent values. For example, if 60% consider item I to belong in category A, then no more than 60% should consider I to be in the conjunction of A

with B. Note that many previous demonstrations of the failure of the minimum rule (Smith & Osherson, 1984) have been confined to typicality judgments. Here we are considering whether or not the rule constrains class membership of conjunctions. Smith and Osherson (1984) and Smith et al. (1988) would predict that the minimum rule *should* apply to categorization judgments, since these decisions depend on core meaning rather than on diagnostic prototype information. In calculating the minimum constraint for negated conjunctions, the probability of an item's *not* being in a class was defined as the proportion of nonzero ratings that were negative for that category [$p(B) = 1 - p(\neg B)$].

Table 3 shows the predicted values for the two functions compared with observed values for conjunctions [$p(A \& B)$] and for negated conjunctions [$p(A \& \neg B)$]. Both functions strongly underestimated the categorization probability to both conjunctions and negated conjunctions. Using the minimum criterion, mean probability was overextended on average by .105 for conjunctions and by .143 for negated conjunctions. The overall root mean square deviation for $p(A \& B)$ was .203 for the product rule and .195 for the minimum rule. For $p(A \& \neg B)$, it was, respectively, .214 and .200.

The data were also analyzed to see whether there were any items where the probabilities for complex concepts broke the lower limit derived from probability theory—that is, where $p(A \& B)$ was less than [$p(A) + p(B) - 1$]. There were only minor cases in which this lower limit was broken. There was therefore no evidence for any corresponding *underextension* of category items.

Subjects in the groups judging the complex concepts overextended the categories considerably with respect to the subjects in the simple concept groups. The items showing the strongest overextension are shown in the Appendix, where it can be seen that both head-noun and

Table 3
Observed Probabilities of Positive Categorization and Probabilities Expected From the Intersection and the Minimum Rules for the Complex Concepts for Each Pair of Categories in Experiment 1

Category A	Category B	Conjunctions					
		A&B			B&A		
		Obs	INT	MIN	Obs	INT	MIN
Pets	Birds	.62	.47	.48	.55	.47	.48
Tools	Weapons	.66	.46	.47	.66	.46	.47
Machines	Vehicles	.76	.59	.60	.66	.59	.60
Furniture	Household Appliances	.42	.30	.35	.43	.30	.35
Sports	Games	.69	.67	.68	.73	.67	.68
Buildings	Dwellings	.74	.59	.60	.70	.59	.60
		Negated Conjunctions					
		A&¬B			B&¬A		
		Obs	INT	MIN	Obs	INT	MIN
Pets	Birds	.23	.18	.19	.40	.26	.27
Tools	Weapons	.39	.27	.28	.28	.22	.23
Machines	Vehicles	.25	.15	.18	.26	.17	.19
Furniture	Household Appliances	.59	.34	.38	.45	.21	.25
Sports	Games	.32	.11	.13	.28	.18	.19
Buildings	Dwellings	.33	.15	.17	.34	.15	.17

Note—Obs, observed; INT, intersection rule; MIN, minimum rule.

qualifier-noun categories could be overextended in both conjunctions and negated conjunctions. There were 79 items (29%) in conjunctions and 166 items (69%) in negated conjunctions that exceeded the minimum constraint. Categorization probability for the conjunction most often lay between the two constituent probabilities, while for the negated conjunction it most often lay between $p(A)$ and $1 - p(B)$.

Overextension of categorization probability for both constituents at the same time was largely confined to the negated conjunctions. For conjunctions, no items had $p(A\&B)$ exceeding that of both constituents by more than .025. For negated conjunctions, however, there were 18 (8%) such items. Most of these items had $p(A)$ in the region of .7 for the head-noun constituent and $p(A\&\neg B)$ of around .8 for the negated conjunction. For example, Tents and Tree Houses were both poor Dwellings (.75 and .737) and differed widely as Buildings (.25 and .8, respectively), but both had a high probability of being rated as "Dwellings which are not Buildings" (.89 and 1.0). Note the striking overextension whereby 80% of subjects in one group considered Tree Houses to be Buildings, yet 100% of subjects in another group considered them to be "Dwellings which are not Buildings."

"Double overextension" of negated conjunctions can be interpreted as demonstrating the viability of the idea that some items may be reasonably considered to be both B and Not B. For instance, a Tent could be both a Building (it has to be erected) but also Not a Building (it is temporary, portable, and so forth). It is interesting to note that this overlap between A and Not A is one of the points that Osherson and Smith (1981) found particularly counterintuitive, in their first attack on the adequacy of prototype theory. It is, however, quite consistent with Tversky's (1977) finding that similarity judgments and dissimilarity judgments may be made on the basis of different semantic information (see also Medin, Goldstone, & Gentner, 1993). Across all items, there were, in fact, 28 cases (12%) for which categorization probability was greater than .5 both for a conjunction and for its negated form.

This paradoxical overlap between a conjunction and its negated form was examined further by comparing the values of $p(A\&B)$ and $p(A\&\neg B)$ across items. When added together, these values, in principle, should never be greater than $p(A)$ and certainly should always be less than or equal to 1.0. The observed sum, S , of the two probabilities in fact ranged from .053 to 1.725, with a mean of .98 and standard deviation of .345. The finding of sums greater than 1 is clear evidence against core definitions and the classical view of categorization. If items fall into categories if and only if they meet the core-defining criteria for membership, then it is not possible that they should both meet the criteria and also not meet the criteria. There is no way, therefore, in which a classical model of concepts could allow that a majority of people believe that something is in the class "A which are B" and that a majority of people (in a different random sample from the same population) believe that it is in the class "A which are not B."

Since the set A is composed of two parts "A which are B" and "A which are not B," the statistic S , being the sum of these two parts, should correlate with the probability of being in Set A, $p(A)$. As predicted, the correlation between $p(A)$ and S was high, at .845. The residual deviation in S once variance due to $p(A)$ was removed was not correlated with $p(B)$ but was significantly correlated [$r(237) = .30, p < .001$] with closeness to the borderline for category B, as measured by $.5 - |[p(B) - .5]|$. Thus, for a given value of $p(A)$, the closer an item was to the borderline of B, the greater was the excess in the sum S (and hence the overlap of the conjunction and negated conjunction concepts). As expected by the prototype view, overextension of each concept was associated with the fuzziness of category membership at the class boundary. According to the prototype model, overextension can be interpreted as being the result of people setting a broad decision criterion for judging category membership, as a result of which the boundaries for both conjunctions and negated conjunctions appear overextended and items near the borderline of B are judged to belong in both complex categories.

Discussion

For conjunctions, the phenomena of overextension, noncommutativity, and concept dominance reported in previous research (Hampton, 1988b) were replicated in a between-subjects design with new items, and appeared equally in analyses of mean typicality ratings and of categorization probability. The binary model suggested by Smith and Osherson (1984) in which typicality and categorization are based on different sources of semantic information received no support in that, even when considering the likelihood of a positive categorization, there was a strong tendency for the value for a conjunction to be some average of the two constituent values.

Negated conjunctions also appeared to be defined as an interactive combination of the two constituent concepts. For both complex concepts, the criterion for membership was drawn more widely than would be predicted by extensional set criteria. Thus it was shown that conjunctions overlapped with their negated conjunctions and that the amount of overlap depended on the closeness of the item to the boundary of the qualifier noun. All the data presented here are completely consistent with the existence of a procedure to form a composite prototype to represent complex concepts (including conjunctions with negated qualifiers), so that membership in the complex concepts depends not on membership in the constituent sets but, rather, on similarity to the composite prototype. The data also rule out the possibility that overextension in earlier studies was simply an artifact of the fixed order of obtaining the categorization decisions in a within-subjects design. In a similar study using a between-subjects design, Storms et al. (1996) collected membership judgments for 30 items in each of 50 different concept conjunctions and their constituents. Analysis of their data (kindly provided by the authors) revealed a strong degree of overextension, with 60% of

items and 82% of the category pairs showing net overextension, compared with 15% of items and only 4% of categories showing net underextension.

The regression analyses confirmed earlier findings of noncommutativity and dominance effects in “straight” conjunctions. Dominance was still effective when the qualifier concept was negated, but noncommutativity was not. In fact, negated conjunctions showed an interaction between the two effects. Dominance was seen in the head nouns of negated conjunctions but was not found when the same nouns acted as negated qualifiers. To explain this effect within a composite prototype model would require a specification of how negation operates on the intension of a concept, a process that is little understood (but see Smith et al., 1988). For example, if negation has an effect on the relative number of salient attributes that the qualifier concept contributes to the composite, this might explain why the dominance effect is less apparent for negated qualifiers (see Hampton, 1987). The aim of Experiment 2 was to provide some preliminary data on this issue.

A further interesting difference between conjunctions and negated conjunctions was that double overextension was found only for negated conjunctions. The best exemplars of conjunctions were either equally typical of one or the other set or else were not rated as maximally typical of the conjunction. Sets such as “Furniture which is also a Household Appliance” appeared to be nonempty but yet to have empty centers (at least among possible exemplars known to the subjects). The existence of empty-centered conjunctions is inconsistent with Jones’s (1982) extensional conceptual stacks model.⁷ For the negated conjunctions, a different pattern was observed. Best exemplars of the negated conjunctions were rated as highly typical, even when they were not typical of the head constituent and were relatively highly related to the modifier. In terms of probability of categorization, double overextension (breaking the minimum constraint for both constituents at once) was also found only for negated conjunctions. There were more people willing to classify an item like *tent* in the category “Dwellings which are not Buildings” than there were people who would classify it as a Dwelling or people who believed it was not a Building.

Why should double overextension and high typicality for best exemplars be more common for negated conjunctions? While it would be wise to study a wider range of concepts before reaching any firm conclusions, one tentative answer may lie in previous theoretical discussion of interactive processes in concept combination (Hampton, 1987, 1988b, 1991) where it was argued that an item’s superior typicality in a conjunction depends on the interaction of intensional information when a pair of concepts are conjoined. For example, pets and fish have different and incompatible values for attributes such as habitat and warm/cold blooded. The conjunction pet fish takes its habitat from pets and its blood temperature from fish, thus excluding some attributes of each concept (pet fish do not live in the ocean, lakes, or rivers like other fish and

they are not warm blooded like most other pets). Hampton (1987) suggested that it is this mutual exclusion in the inheritance of particular attributes which may account for double overextension effects, since it permits an item (such as a guppy) to be more similar to the composite representation of pet fish than it is to either of the constituent concept prototypes.

To explain the stronger double-overextension effect found for negated conjunctions in this model, one would have to assume that the resulting complex concept bears *less* relation to its constituents when the modifier concept is negated than when the modifier is not negated. Why should this assumption be correct? One argument in favor of the assumption is that, as will be argued below, the process by which a negated conjunction inherits attributes from the negated-modifier constituent is considerably more complex than it is for straight conjunctions. Some of the attributes of the negated modifier may in fact be inherited in negated form (e.g., “Dwellings which are not Buildings” are *not permanent*). Others may be inherited as positive attributes in spite of the negation (e.g., “Dwellings which are not Buildings” *serve some specific purpose, are places*, etc.). Yet others may not be inherited at all. In broad terms, then, one could argue that negation introduces a greater difference between the attributes of a complex concept and its constituents, and hence produces more extreme overextension effects. This explanation may also serve to explain why concept dominance is no longer apparent for the negated-modifier constituents.

To understand the effects of negation, it is clear that a considerable amount of further research will be required. In particular, there is a need to look at the attribute inheritance pattern for negated conjunctions. This was according to the aim of Experiment 2.

EXPERIMENT 2

In order to study attribute inheritance in conjunctions, Hampton (1987) had subjects first generate attributes of component concepts, and of their conjunctions, and then rate them for their definingness for each concept. For example, “Is competitive,” which was generated as an attribute of Sports, was rated for its importance as an attribute of Sports, of Games, and of “Sports which are also Games.” It was found that the pattern of inheritance of attributes by conjunctions from their constituents was generally regular, with conjunctive importance of an attribute predictable from an average of its importance for the constituents. The average rule, however, was subject to the constraint that attributes that were either necessary or impossible for constituents were also necessary or impossible for their conjunctions with other sets. Other considerations (such as a dominance effect for attribute inheritance matching that for category membership judgments) led Hampton (1987) to propose that categorization in a conjunction is actually made by some similarity computation between the conceptual representation of an item being categorized and the composite prototype rep-

resentation of the conjunction resulting from the process of attribute inheritance. (For evidence that the process may work in the opposite direction, see Storms et al., 1993.) Discovering the nature of the attributes involved in representing negated conjunctions could then provide an account of how people make category judgments with respect to these concepts.

When a person forms a complex concept such as “Tools which are not Weapons,” how do they identify the intensional attribute information for such a concept? According to the classical model of concepts, a “Tool which is not a Weapon” should possess all of the core-defining attributes of Tools and should fail to possess *at least one* of the core-defining attributes of Weapons. Logically, this missing defining attribute (or these attributes) of Weapons cannot also be a defining attribute of Tools. The classical model, however, has little or nothing to say about other nondefining attributes. Indeed, according to Fodor (1994), it is in principle impossible for a theory of concepts to determine which nonnecessary attributes would be inherited.

How would the composite prototype model attempt to answer the question? The attributes of Tools would be inherited as a part of the composite concept, subject to a consistency check with the attributes inherited from the Not Weapons constituent. For the inheritance of the attributes of Weapons, there are, then, three possibilities. Some attributes may be inherited by the composite prototype in spite of the effect of negation. These attributes would include general aspects, such as “Is used by people.” It is, in principle, possible for the negated conjunction to inherit any one or more of the attributes that are generally true of the negated modifier Weapons, provided that it does not inherit too many of them (and hence provide too good a similarity fit to items that are, in fact, Weapons). Good candidates for attributes that are inherited unchanged from the negated modifier will be those that are also important attributes of the head-noun Tool concept.

A second possibility is that attributes that tend to be more specific to Weapons, such as “Used in war,” may be omitted from the composite prototype. They are simply not inherited as part of the complex concept representation (even though tools that are not weapons most probably are used in war). Finally, a third possibility is for attributes to be inherited, but in a *negated* form. Thus, “Can be used to kill people” might be inherited as its opposite “Cannot be used to kill people.” Attributes that distinguish Weapons from Tools are the most likely candidates to fall in this third class, because negating the Weapons attribute should not involve negating an attribute of Tools.

Further constraints on attribute inheritance under negation can be derived from pragmatic considerations. For example, the phrase “someone who is not a bachelor” is more likely to refer to a married adult male than to a baby girl. Marital status is relevant only if a person is of a marriageable age, and so takes a foreground role in

the concept’s definition. We should therefore expect that this presuppositional structure, together with contextual relevance, may also determine the scope of negation within the attributes defining a concept. Information that is “foregrounded” in the negated constituent should be the most likely to be inherited in negated form, while “background” attributes would be unaffected by the negation.

One way that a model for negation could be developed would be as an extension of Smith et al.’s (1988) selective modification model. Assuming a frame representation for a noun concept, then, negation of an attribute could be modeled by increasing the salience of the slot and decreasing the votes for the specified value to zero. Thus, an Apple that is not red would have the slot [COLOR] given increased salience and the slot value *red* given zero votes. As a result, typicality would then be increased for any object, to the extent that it was some common color of apples other than red.

Extending the general idea to the case of relative clause conjunctions and composite prototypes, the problem is then to determine which slots will be negated and which will remain unaffected. Consider an example such as a “Building which is not a Dwelling.” The analysis requires that we start with a semantic representation of Building. We then have to look at the foreground information in Dwelling (the most salient attributes, however salience might be defined), which in the case of Dwelling (as for many artifact categories) might be a [FUNCTION] slot

Dwelling: [FUNCTION {human habitation}].

Negation then denies this possible value for the same slot in the frame representation of Buildings. The intriguing question is then: how far down the line does this process go? Does negation apply to all the Dwelling attributes? This is clearly unnecessary. As discussed above, if there are central or necessary features in the concept definition, then only one such feature needs to be missing to render some item out of the class, so, provided the foregrounded attribute is a necessary one, there is no reason to negate any further attributes (like roofs, or doors, or heating). On the other hand, the consistency requirement for the composite concept is liable to propagate changes through the concept—no habitation function is liable to imply no bedrooms, no curtains in the windows, since the building will be empty at night, and so on. We may therefore expect *some but not all* attributes of the negated concept to be also negated.

The extent to which the head-noun attributes are changed by the modifying relative clause should also depend on the degree of overlap between the two categories. For example, the category “Birds which are not Aeroplanes” is unproblematically just the same concept as Birds (ignoring for now metaphorical interpretations, such as that the bird is not huge and noisy, or that it does not fly by gliding with outstretched wings). It is primarily only to the extent that two concepts overlap extensionally that they will interact. Thus, negated conjunction categories formed from strongly overlapping sets like “Dwell-

ings which are not Buildings” should involve greater modification of the head noun.

Theoretical analysis of this kind depends on purely intuitive judgments to understand the process of negation. There is clearly a need for empirical data to provide some harder evidence of the phenomenon. To investigate how composite prototypes are constructed for negated conjunctions, attributes for these concepts were elicited from subjects in Experiment 2, and judgments were then collected of the frequency of occurrence of the attributes in the constituent and complex categories.

Method

Subjects. The subjects were 60 students at City University London who volunteered for the experiment and received £2 for their participation.

Materials. The same six pairs of concepts as in Experiment 1 were used to generate 12 negated-conjunction classes. For example, the pair Dwellings–Buildings generated two negated conjunctions: “Dwellings which are not Buildings” and “Buildings which are not Dwellings.” Subjects in Phase 1 generated the attribute materials to be used in Phase 2.

Procedure. In Phase 1, two groups of 10 subjects each were asked to generate attributes to describe and define a set of six negated conjunctions. One of each pair of concepts was randomly assigned to be the A category; the other was the B category. One group of subjects listed attributes for six negated conjunctions created by instantiating the phrase “A which are not B” with the six category pairs. The second group listed attributes for the converse sets (“B which are not A”). Thus, for example, where Group 1 was given “Vehicles which are not Machines,” Group 2 was given “Machines which are not Vehicles.”

All of the attributes listed by each group were then collated as follows to produce a master list for each category pair. The two lists of attributes for each pair of categories were first collapsed into a single list. This list contained two kinds of attributes: those that were listed in a direct positive form (e.g., *Has windows, Provides comfort*) and those that were listed in a negative form (e.g., *Has no roof, Does not use fuel*). The lists were edited so that each negatively expressed attribute was reexpressed in a positive form (e.g., *Has a roof, Uses fuel*) in order to avoid the potential confusion involved in judging the degree to which the *lack* of an attribute was present or absent in a class. A pilot study had suggested that subjects were easily muddled when called on to make negative judgments about negative properties.

To summarize, the master list was composed of attributes that were generated in either positive or negative form, either for “A which are not B” or for “B which are not A.” The master list was then rated in Phase 2 by a further four groups of 10 subjects each. Two groups judged the list with respect to each of the constituent categories A and B, respectively, and two other groups judged the list with respect to each of the negated conjunctions, “A which are not B” and “B which are not A,” respectively. Their task, in every case, was to work down the list of attributes, arranged in alphabetical order, in two columns per page for each concept, and, by giving an estimate of the percentage of category members that would have that attribute, to judge how likely it was that a member of the category would have each property. Previous research on attributes (Hampton, 1987) used a more qualitative scale involving “importance” and “necessity” of attributes for concept definitions. Use of the percentage scale was intended to allow issues of the distribution of attributes in concept classes to be addressed more directly. It also provided a concrete interpretation for the otherwise rather vague notion of “attribute importance.”

Results and Discussion

Generation of attributes.⁸ Preliminary analysis of individual conjunctions suggested the advisability of performing separate analyses of conjunctions with highly overlapping constituent categories (e.g., Vehicles and Machines) and those with little overlap (e.g., Furniture and Household Appliances). It was expected that where two categories had weak overlap, there would be little need to generate many negatively expressed attributes, since “A which are not B” would be very similar to A. For highly overlapping categories, however, more negatively expressed attributes (and emergent attributes) were expected.

To separate the conjunctions on the basis of overlap, a group of 10 graduate students, otherwise unfamiliar with the aims of the research, were asked to judge for each pair of categories what proportion of examples of A also belonged in B, using a 5-point scale with numbers labeled 1 = *almost all*, 2 = *more than half*, 3 = *about half*, 4 = *less than half*, and 5 = *very few*. They also provided numerical estimates of the percentage of examples of A which were also in B. Each judge rated the six pairs of categories twice, once with one assignment of categories to A and B (e.g., Pets = A, Birds = B) and once with the reverse assignment (e.g., Birds = A, Pets = B). The results provided a clear division of the 12 negated conjunctions into three groups. Concepts that had high overlap were Sports–Games, Dwellings–Buildings, and Vehicles–Machines (mean ratings 1.1–1.7, mean percentage estimate 80%–82%); those with medium overlap were Weapons–Tools, Tools–Weapons, Buildings–Dwellings, and Games–Sports (2.2–3.5, 35%–48%); and those with low overlap were Birds–Pets, Pets–Birds, Machines–Vehicles, Furniture–Household Appliances, and Household Appliances–Furniture (3.9–4.8, 6%–26%).⁹

The numbers of positively and negatively expressed attributes generated for each set are shown in Table 4. Recall that a negatively expressed attribute is an attribute that was generated expressed in a negative form (e.g., “has no wheels”) and that high-overlap categories are those where a high proportion of examples in the head-noun category are also in the unnegated-modifier category (e.g., “Sports which are not Games”). There were, on average, around 24 different positively expressed and 8 different negatively expressed attributes produced to each negated conjunction. Across the three levels of overlap, the expected trend was seen, with the high-overlap negated conjunctions having more negatively expressed

Table 4
Mean Number of Attribute Types and Tokens
Generated Positively or Negatively per
Negated Conjunction in Experiment 2

	Low Overlap		Medium Overlap		High Overlap	
	Types	Tokens	Types	Tokens	Types	Tokens
Positive	23.0	36.8	27.0	40.2	23.7	34.3
Negative	7.6	15.4	7.7	12.7	11.7	23.7

attributes generated than the other negated conjunctions (mean of 11.7 types and 23.7 tokens, as compared with 7.6 types and 14.0 tokens for the other two sets of negated conjunctions). However, the interaction did not reach significance in an ANOVA of generation frequency across categories. The lack of significance reflected large differences within category pairs at the same level of overlap. For example, within the high-overlap group, “Dwellings which are not Buildings” had 19 different negatively expressed attributes, but “Sports which are not Games” had only 5.

Where do the positively and negatively expressed attributes come from? A simple hypothesis would be that positively expressed attributes are chosen from the head noun and negatively expressed attributes are chosen from the modifier. Table 5 shows the number of positive attributes generated to the negated conjunctions in Phase 1 of the experiment as a function of the median rated frequency of the attribute for each constituent. (In all tables and analyses where frequency judgments were collapsed across subjects, the median was always used as the measure of central tendency.) For this analysis, ratings of frequency of occurrence for categories A and B alone were collapsed into five equal intervals corresponding to the percentage ranges of 0–20, 21–40, 41–60, 61–80, and 81–100.

The expected trend can clearly be seen. Positive attributes tended to be rated as high frequency for the head-noun category A (81%–100% row in Table 5) and as low frequency for the modifier category B (the 0%–20% column in Table 5). These two intervals together accounted for 245 (83%) of the 294 positive attributes generated. Analysis of the frequency ratings confirmed that positively expressed attributes were rated as more frequent for the head noun than for the modifier [mean ratings of 53% for A, 31% for B; $t(293) = 11.8, p < .001$].

This result should not be taken at face value, however, since the mean is not a good summary of either scale. The two most common groups of positively expressed attributes generated were those that were uncommon for both A and B (76 attributes) and those that were rated as frequent in both categories (46 attributes). Both marginal distributions showed bimodality. The uncommon attributes could be termed *emergent attributes* (Hampton, 1997b; Kunda, Miller, & Clare, 1990) in that they are attributes that would not normally be part of the represen-

tation of either concept alone but would emerge in the complex concept representation. Hampton (1987, 1997b) suggested that one source of emergent attributes, which he termed extensional feedback, is information linked to particular instantiations or exemplars of the complex category. Subjects may generate distinctive attributes of particular exemplars of the complex category that are not generally common in either constituent class. Inspection of the data showed that this was often the case, with many attributes naming exemplars (e.g., *Is a tent, Is a burrow* for “Dwellings which are not Buildings”) or describing attributes of those exemplars (e.g., *Has flaps, Has holes*).

The positive attributes rated as common for *both* constituents tended to be generally true of the domain. For instance, Birds and Pets are both alive, animals, need to eat, and need to breathe. Such attributes are also true of “Birds which are not Pets,” and “Pets which are not Birds,” and were indeed rated as such. These attributes were not generated as negative attributes, even though they were common attributes of the negated concept, because they were also central to the definition of the head-noun concept.

Table 6 shows the equivalent breakdown for the negatively expressed attributes. (Recall that the frequency ratings were made to the attributes in their unnegated form. The table shows the rated frequency of the *positively expressed* form of these attributes. For example, *is not permanent* was generated as an attribute of “Dwellings which are not Buildings.” The corresponding positively expressed attribute *is permanent* was judged as true of 90% of dwellings and 80% of buildings.) When choosing attributes to be produced in negated form, the preference was for attributes with high frequency in category B (55 in the 81%–100% column) and those with low frequency for category A (45 in the 0%–20% row). Together, these two bands accounted for 83 (80%) of the 104 negatively expressed attributes generated. Analysis of ratings confirmed that negatively expressed attributes were rated as less frequent for the head noun than for the modifier [mean ratings of 41.8 for A, 73.0 for B; $t(103) = 8.56, p < .001$]. Comparing Tables 5 and 6, an ANOVA of the rated frequencies confirmed the significant crossover interaction between ratings for head versus modifier nouns, and whether attributes were expressed positively or negatively [$F(1,396) = 196.5, p < .0001$].

Interestingly, there were a fairly high number of negatively expressed attributes with high-rated frequencies

Table 5
Number of Positively Expressed Attributes Generated to Negated Conjunctions, Tabulated as a Function of Their Judged Frequency in Each Constituent Concept in Experiment 2

Rated Frequency in Category A	Rated Frequency in Category B					Total
	0%–20%	21%–40%	41%–60%	61%–80%	81%–100%	
0%–20%	76	3	2	2	0	83
21%–40%	24	4	1	1	1	31
41%–60%	20	6	6	2	2	36
61%–80%	20	4	4	4	7	39
81%–100%	28	11	13	7	46	105
Total	168	28	26	16	56	294

Table 6
Number of Negatively Produced Attributes Tabulated as a
Function of Their Judged Frequency in Each Constituent Concept

Rated Frequency in Category A	Rated Frequency in Category B					Total
	0%–20%	21%–40%	41%–60%	61%–80%	81%–100%	
0%–20%	4	7	6	11	17	45
21%–40%	1	0	2	1	4	8
41%–60%	0	1	2	4	10	17
61%–80%	0	0	2	3	1	6
81%–100%	0	1	3	1	23	28
Total	5	9	15	20	55	104

for both A and B. Comparing Tables 5 and 6, there were 46 positively expressed and 23 negatively expressed attributes generated in this cell. In addition to the domain general (background) attributes that did not get negated, there were also, therefore, attributes which were generally true of both A and B but which *did* get negated. For instance, Vehicles and Machines both typically use fuel, yet an attribute generated for “Vehicles which are not Machines” was *does not use fuel*. There was no consistency between the two negated conjunctions formed from the same pair of categories in whether an attribute that was common to both categories would be generated as positive or negative. There were 18 attributes generated (either positively or negatively) to both “A which are not B” and “B which are not A,” which were also rated as over 80% frequency in both A and B categories. Of these 18, 9 were generated positively to both negated conjunctions (for example, *is an activity* was positive for both “Sports which are not Games” and “Games which are not Sports”), and 1 was generated negatively to both negated conjunctions (*is competitive* for the same pair of negated conjunctions). The remaining 8 attributes were generated positively for one negated conjunction but negatively for the other (for example, *has moving parts* was generated positively for “Machines which are not Vehicles” and negatively for “Vehicles which are not Machines”).

To see how these data differed as a function of the relatedness of the combined categories, Table 7 shows the same data broken down into high-, medium-, and low-overlap negated conjunctions on the basis of the previously described overlap measure. To compensate for low cell frequencies, the five frequency bands used in Tables 5 and 6 were collapsed into just two, using a 40% cutoff

point. The table shows the mean number of attributes (positive and negative) generated per conjunction, as a function of a binary division of attributes that were rated as high or low frequency for each constituent.

The upper half of Table 7 confirms that positively expressed attributes had mainly low frequencies for the constituent B and high frequencies for constituent A. However, this was not true across all types of category. In particular, it was not true for the category pairs with high overlap. When the two concepts, A and B, had medium or low overlap, one would expect there to be a larger number of distinctive attributes, true of one concept and not true of the other. Subjects were apparently able to access this pool of attributes as positive descriptions of the class “A which are not B,” producing, on average, around 9 such attributes per conjunction. For the high-overlap categories, however, there were only 2 attributes generated per conjunction which were high frequency for A and low frequency for B. For these strongly overlapping categories, attributes that are false of B will also tend to be false of A. Hence, the positive attributes generated for the negated conjunction tended to be mostly low frequency for both A and B, and hence generally true of neither set (on average 14.3 such attributes were generated). As expected, a negated conjunction created from such closely overlapping sets as “Vehicles which are not Machines” involves the construction of a prototype with emergent attributes that are not usually true of the head noun.

Negatively expressed attributes showed the converse pattern. For all conjunctions, negatively generated attributes tended to be high frequency for Set B. For low-overlap categories, being common in B would entail being

Table 7
Mean Numbers of Positively and Negatively Expressed Attributes Generated Per
Conjunction as a Function of Constituent Frequency Greater or Less Than
40% and as a Function of Constituent Category Overlap

Rated Frequency for Head Constituent A	Rated Frequency for Negated Constituent B					
	High Overlap		Medium Overlap		Low Overlap	
	High	Low	High	Low	High	Low
Positively Expressed Attributes						
High	7.0	2.0	9.0	9.0	6.8	9.4
Low	0.3	14.3	0.2	8.7	1.0	5.8
Negatively Expressed Attributes						
High	9.7	0.7	3.2	0.0	1.4	0.0
Low	0.7	0.7	4.2	0.2	4.4	1.8

uncommon in Set A. Hence, attributes generated as negative for low-overlap category pairs tended to be mostly high frequency for B and low frequency for A (mean number generated = 4.4). For high-overlap categories, being common in B tends to entail also being common in A. Consequently, for these sets, the negative attributes appeared in the high-high cell (mean = 9.7). The medium overlap sets fell in between these two extremes.

Effectively, then, positive attributes were of two kinds—those that were high frequency for both constituents (91 attributes) and those that were low frequency for B (196 attributes). Together, these two sources accounted for 98% of positively expressed attributes. Whether an attribute that was low frequency for B would be high or low frequency for A depended simply on the overlap between the two categories. Conversely, the main criterion for generating a negatively expressed attribute was that it was high frequency for category B (90 attributes, accounting for 87% of attributes). Whether or not the attribute was also common for category A depended again simply on the overlap between the two categories. This rather complex pattern of data can be summarized by the rule that apart from the domain of general attributes, it was frequency in category B which largely determined whether an attribute was generated (expressed either positively or negatively).

Estimated percentages in the negated conjunction.

In addition to considering the *number* of attributes produced as a function of constituent percentage occurrence, the average median rated percentage occurrence of each attribute for the negated conjunction was calculated in a second analysis. The measure showed how frequent each positively or negatively expressed attribute was judged to be for the complex category. Table 8 shows the data for positively expressed attributes. From the marginal weighted averages, it is clear that for positive attributes the strongest influence on the frequency of an attribute in the negated-conjunction category was its frequency in the head-noun category A. Estimates ranged from 19% in the 0%–20% band for A up to 87% for the 81%–100% band. In fact, the mean values for each row lay within the range of each band for the A category, so that $\%(A-B)$ was approximately equal to $\%(A)$. Looking across the columns, there was no apparent systematic effect of the rated frequency in B on the mean data in the table. In fact, when regressions were run on each of the 12 negated conjunctions, predicting $\%(A-B)$ from $\%(A)$ and $\%(B)$ for positively expressed attributes, the head-noun

rated frequency $\%(A)$ entered significantly in every case, with a mean regression weight of 0.83, while the modifier-noun rated frequency $\%(B)$ entered significantly in only 3 of the 12 equations, with a mean regression weight across all equations of -0.01 . (Root mean squared R for the regressions was 0.90.) This result was curious in that one would predict that attributes that were more common in category B should have been considered less common in the conjunction “A which are not B.” However, there are also the domain-general attributes which were judged 100% true of both A and B and of “A which are not B,” and these clearly counteracted the tendency of frequency in B to correlate negatively with frequency in the negated conjunction.

Recall that 76 positive attributes were generated with rated percentages of less than 20% for both A and B (see Table 5). The average rated percentage in the negated conjunction for these attributes was only 18%. In other words, these attributes were also uncommon in the negated conjunction class. When examined more closely, many of these attributes (45, or 59%) turned out to refer to particular instantiations of the negated-conjunction categories, either by naming them directly or by describing very specific attributes of an instantiation. For example, for the Dwellings–Buildings category pair, there were 27 positive attributes with low frequency for both categories. Twelve of these named specific instantiations—*is a burrow, is a cave, is a tunnel* (for “Dwellings which are not Buildings”) or *is a place of worship, is a monument* (for “Buildings which are not Dwellings”). A further 7 were very specific attributes derived from an instantiation (e.g., *has flaps, is for entertainment, is for storage*). The remaining 8 attributes were more general, as in *is natural, is portable, is cold*, and so could be seen as being generated more from the composite prototype representation than from specific instantiations. In terms of Hampton’s (1987) analysis, the instantiation-based attributes are examples of emergent attributes derived from extensional feedback—the strategy of identifying specific exemplars in the complex category and basing the attribute description on those exemplars. The more general attributes, of which there were 31 (41%), would be emergent attributes resulting from the attempt to form a coherent category from the composite prototype representation.

Table 9 shows the equivalent mean estimated percentages for negated attributes. One would expect that the frequency of negatively expressed attributes in the negated conjunction “A which are not B” should be gener-

Table 8
Mean Estimated Frequency (%) for Positively Expressed Attributes
as a Function of Estimated Frequency for Each Constituent Category

Rated Frequency in Category A	Rated Frequency in Category B					Weighted Mean
	0%–20%	21%–40%	41%–60%	61%–80%	81%–100%	
0%–20%	18	39	–	–	–	19
21%–40%	37	24	–	–	–	33
41%–60%	49	44	58	–	–	50
61%–80%	63	80	76	72	68	68
81%–100%	83	88	91	74	91	87
Weighted Mean	41	63	73	66	85	

ally lower than that for positively expressed attributes, and the table confirms this expectation. Although the smaller number of negative attributes left many of the cells in Table 8 empty, the data show that, as in the case of the positive attributes, the degree of occurrence in the head-noun category A was the most influential factor. The more common an attribute was in category A, the more common it was in the class “A which are not B.” However, the slope of this relationship was much less steep than before, so that attributes with high frequency in A averaged only 44% for “A which are not B.” Because there were many fewer negatively expressed attributes, only two regression analyses were conducted, combining the six pairs of categories together for each order of the constituents. Frequency in the head noun was a significant predictor of frequency in the negated conjunction (regression weights = .37 and .54), while frequency in the modifier was not a significant predictor in either case.

Distinctiveness. A plausible explanation of why subjects should generate positively or negatively expressed attributes might appeal to the distinctiveness of the attributes for distinguishing between the head-noun class A and the negated conjunction class “A which are not B.” Pragmatically, one would expect a predicate to be relevant and informative if it helped the hearer to distinguish the negated conjunction class from its superordinate category A. Hence, positively expressed attributes should have a higher rated frequency in the negated conjunction than in the head-noun class alone, whereas negatively expressed attributes should have lower rated frequency in the negated conjunction than in the head-noun class. Table 10 shows a breakdown of numbers of attributes generated per conjunction as a function of whether they were positively or negatively expressed, and of whether they were rated as more common in the negated conjunction than in the head noun or as less common.

For the concept pairs with high overlap, the distinctiveness hypothesis was very clearly confirmed. Of 63 positively expressed attributes, 45 (71%) were rated as more frequent in the negated conjunction class than in the head-noun class, whereas 32 (94%) of the 34 negatively expressed attributes were judged as less frequent in the negated-conjunction class than in the head-noun class [$\chi^2(1) = 38.0, p < .0001$]. For the medium- and low-overlap categories, the distinctiveness hypothesis was also supported [$\chi^2(1) = 8.4$ and 6.1 , respectively, $p < .02$], but the effect was less strong, with considerable numbers of attributes going against the hypothesis.

Emergent attributes and inheritance failure. What evidence was there for emergent attributes—that is, those produced as describing “A which are not B” which were considered generally untrue of A?¹⁰

To detect emergent attributes, a strict definition of emergent attributes was used by defining them as those attributes with a median rated frequency of 0% for A but greater than 0% for “A which are not B.”¹¹ Of the 51 attributes with median rated proportion for A as 0%, 24 also had 0% ratings for “A which are not B.” Eleven had median ratings of 5% or more, and only 6 attributes had median ratings for “A which are not B” of 10% or more. There were therefore very few examples of strictly emergent attributes here, in the sense of attributes that break the logical consistency constraint of being judged to be entirely absent from the head-noun class, and yet present to some degree in the negated conjunction. Table 11 shows the 6 most extreme cases. The strongest example numerically—birds which are not pets having noses—seems to depend on a shift in the meaning of the concept of a nose; half of the subjects apparently considered beaks to be noses when judging birds that were not pets (all subjects except 1 gave a rating of either 0% or 100% to the question), but none of the subjects considered beaks to be noses when just considering Birds alone. This shift is probably best explained as an example of metonymy—the concept *nose* is a typical example of the concept *proboscis*. By metonymy, a typical case can also be used as a name to refer to the class as a whole (Lakoff, 1987), so that “nose” may be given a narrow sense and a more general sense. Why there should have been an interaction with the complexity of the category in interpreting “nose” broadly or narrowly remains an intriguing question. It could be that mention of the category Pets (even in the negative) leads to consideration of a broader range of animals, and so to a broader interpretation of the part named by the word “nose.”

The dwellings example involves a shift from dwelling as permanent residence (the default meaning of dwelling) to dwelling as shelter, as in a tent used on a camping trip. The vehicle examples show a shift from thinking of a vehicle as a mechanical means of transport to thinking of a broader category which would include animals such as horses, camels, and donkeys. The latter two cases involve broadening of the concept definitions, which would suggest that some default attributes from the prototype concept have been dropped. Such a result would be compatible with forming a composite prototype, since in

Table 9
Mean Estimated Frequency (%) for Negatively Expressed Attributes
as a Function of Estimated Frequency for Each Constituent Category

Rated Frequency in Category A	Rated Frequency in Category B					Weighted Mean
	0%–20%	21%–40%	41%–60%	61%–80%	81%–100%	
0%–20%	10	12	6	4	4	6
21%–40%	–	–	–	–	21	20
41%–60%	–	–	–	24	35	29
61%–80%	–	–	–	28	–	25
81%–100%	–	–	47	–	43	44
Weighted Mean	12	20	21	13	27	

Table 10
Mean Number of Attributes Generated per Conjunction With Median Rated Frequency Greater in the Negated Conjunction Than in the Head-Noun Class or Vice Versa, as a Function of Constituent Category Overlap

	Low Overlap		Medium Overlap		High Overlap	
	Positive	Negative	Positive	Negative	Positive	Negative
$p(A \text{ not } B) > p(A)$	10.8	2.4	12.7	1.7	15.0	0.7
$p(A \text{ not } B) < p(A)$	8.4	5.0	11	5.7	6.0	10.7

both of these cases the concepts are strongly overlapping (most vehicles are machines and most dwellings are buildings), and hence conflict between inherited attributes should be strongest. The broadening of concept definitions is also consistent with the extensional class membership data from Experiment 1, where subjects overextended categories when they were placed in conjunctions and negated conjunctions.

The converse of emergent attributes is the case of inheritance failure—those attributes rated as true of 100% of class A but as true of less than 100% for “A which are not B.”¹² Of 54 attributes with median ratings for A of 100%, 32 had median ratings of 100% for “A which are not B” and a further 11 had ratings of 95% or more. There were only 10 attributes with median ratings of 90% or less for the negated conjunction. Table 12 lists these cases of inheritance failure of apparently “necessary” attributes.

Several of these examples (dwellings having doors and roofs, being occupied, etc.) show an interesting ambiguity in the interpretation of the task in that they seem to involve a move from an intensional answer to an extensional one. When considering the category on its own, people answered in terms of the general prototype picture, even though the task asked for a rating of the proportion of exemplars, which is an explicitly extensional task. Thus, they may have considered that *all* buildings are used, otherwise why would they have been built? The negated-conjunction class “Buildings which are not dwellings” may, however, have encouraged subjects to think of actual exemplars, and of course there are many buildings standing empty and unused. Making the concept more specific through modification may have led subjects to pay closer attention to exemplar information. Alternatively, removing the most typical instances from the class may also have led subjects to consider a wider range of buildings (such as monuments) and so to change their frequency estimates.

Other cases are further examples of the stretching of the concept categories when conjoined with others with which

they share little overlap. Dwellings are usually buildings. When considering those that are not, the concept became extended to include dwellings for nonhuman animals and temporary dwellings. Similarly, the category of vehicles was extended (as discussed previously) to include animals.

Finally, the case of “Birds which are not Pets” having no names, is a case of ambiguity in the sense of the word “name.” Birds all have names, in the sense of a species or generic name, whereas pets of course have names as individuals. In the context of pets, the concept of name is more likely to be interpreted in the latter sense.

GENERAL DISCUSSION

To summarize the findings of the two experiments presented here, the first experiment found that the rated membership of items in conjunctions and negated-conjunction classes was highly predictable from a knowledge of an item’s membership in the two constituent categories. However, there were interesting aspects of these data which suggested an intensional account of the data. Subjects overextended both conjunctions and negated conjunctions with respect to the membership of constituent classes, and there was evidence for systematic variation in the regression weight of different categories, depending on whether they were in head-noun or qualifier-noun position, and depending on the “dominance” of the categories established in earlier research (Hampton, 1987, 1988b; Storms et al., 1993; Storms et al., 1996).

The dominance effect (whereby one category had a greater effect in the regression equation predicting either typicality or the probability of positive categorization in the complex concepts) was found to interact with whether or not a conjunction was negated. In positive conjunctions, dominance and the head/modifier effects were additive, whereas for negated conjunctions, dominance was seen in the head-noun position but not in the negated-qualifier position. An explanation for this interaction

Table 11
Examples of Attributes Estimated to Have Zero Frequency in the Head-Noun Category, But to Have At Least 10% Frequency in the Negated Conjunction, Together With Estimated Frequency for the Negated Conjunction

Negated Conjunction	Attribute	Frequency
Dwellings which are not Buildings	Are for relaxation out of the home	10%
Household Appliances which are not Furniture	Are for play	20%
Vehicles which are not Machines	Are natural	35%
	Are self-motivating	13%
	Are rafts	10%
	Have a nose	53%
Birds which are not Pets		

Table 12
Examples of Attributes Estimated to Have 100% Frequency in the Head-Noun Category,
But to Have 90% or Less Frequency in the Negated Conjunction,
Together With Estimated Frequency for the Negated Conjunction

Negated Conjunction	Attribute	Frequency
Dwellings which are not Buildings	Have a door	26%
	Have a roof	65%
	Are structures	83%
	Are somewhere to keep dry and warm	85%
	May be occupied	95%
Buildings which are not Dwellings	Are used	88%
	Are places	90%
Vehicles which are not Machines	Are inorganic	70%
	Accommodate people	80%
Birds which are not Pets	Have names	75%

was offered in terms of attribute inheritance by the complex concept representations. When a concept is negated, its attributes are divided into those that are inherited unchanged as positive attributes and those that are inherited as negated attributes. It was argued by Hampton (1987) that the dominant concepts were those that had a larger number of attributes considered as important to their definition. When in the negated-qualifier position, the dominance effect might be reduced, given that some attributes are inherited positively and others negatively. Experiment 2 confirmed this prediction. Whereas the rated frequency of an attribute in the negated conjunction was positively correlated with its rated frequency in the head-noun category, correlation with frequency in the negated-modifier category was near zero.

Another interesting result from Experiment 1 was the finding that, for conjunctions, it is possible that conjunctions of categories with low overlap (like “Household Appliances which are Furniture”) may have empty centers. The best examples of such categories may still be considered *atypical* of them. This result is counter to the “guppy effect” introduced by Osherson and Smith (1981). According to their intuitive account, a guppy may be considered a better example of a pet fish than it is an example of either a pet or a fish. In the data presented in Experiment 1, the guppy effect was largely absent from the conjunctive concepts. Even though an attempt was made through pretesting to include the best examples of each conjunction in the set of items, the results showed that these best examples were not always rated as highly typical of the conjunctive category. It was very rare that an item was rated more typical in the conjunction than in *both* of the constituent categories. Interestingly, the same result did not hold true for negated conjunctions. The best examples of negated conjunctions like “Dwellings which are not Buildings” were, in fact, rated as maximally typical in the class, although they had not been rated as typical in either category individually. This effect was found even though the materials had not been expressly designed to include the best examples of the negated-conjunction sets. A tentative explanation of this effect may also be provided by the results of Experiment 2. The inheritance of attributes was far more complex from the negated-qualifier concept than from simple unnegated

qualifiers (Hampton, 1987). Some of the qualifier attributes were inherited in negative form, reflecting the negation, but others were inherited in positive form, regardless of the fact that the qualifier was negated. Hampton (1987) argued that the greater the difference in representation between a complex concept and each of its constituents, the greater would be the likelihood of observing overextensions in the conjunction. The fact that negated conjunctions show a more complex inheritance pattern of attributes is therefore consistent with the greater occurrence of overextension and double overextension in Experiment 1.

The attributes generated in Experiment 2 were analyzed with respect to their rated frequency in each constituent category. To show the data most clearly, the attributes were divided into those that were generated as positive predicates and those that were generated as negative. The pairs of categories were also divided into those with high, medium, or low overlap as categories. As expected, there was a proportion of attributes that were generated positively and were also considered to be generally true of both constituent classes. These were domain-general attributes that were true of the superordinate category to which the two constituents both belonged, and hence were also true of both negated conjunctions. Aside from these positive attributes, the main influence on the generation of both positive and negative attributes was frequency for the modifier category B. The large majority of positive attributes generated were rated as low frequency for category B, while the majority of negative attributes generated were rated as high frequency for category B. Frequency in category A had an apparent effect on numbers of attributes generated only as a result of the relative overlap between A and B. Thus, when overlap was high and A and B shared many of the same attributes, the positive attributes generated were also low frequency for A and the negative attributes generated were also high frequency for A. When overlap was low and A and B shared few attributes, positive attributes generated were high frequency for A and negative attributes generated were low frequency for A.

Although the likelihood of generation as a positive or negative attribute depended solely on category B, the rated frequency of the attributes in the negated conjunc-

tion itself was apparently largely controlled by their frequency in class A. The mean rated frequency for "A which are not B" for positively generated attributes closely followed the mean rated frequency for the A category alone. For negatively generated attributes, rated frequency of the positive form of the attribute in "A which are not B" was also greater as a function of the rated frequency in category A, although mean frequencies in the negated conjunction were understandably lower than they were for the positive attributes.

On the basis of informativeness, it was predicted that positive attributes should be judged to be more true of the negated conjunction than of the head noun, whereas negative attributes (in their positive form) should be judged to be more true of the head noun than of the negated conjunction. This prediction was generally supported, although more strongly so in the categories with high overlap (where most members of the head-noun category A are also in category B). Low overlap, of course, means that forming a negated conjunction for such categories requires less cognitive effort, since the majority of each class also falls in the negated conjunction. For example, since most household appliances are not furniture, "Household Appliance which is not Furniture" may initially be considered as being very close to the concept Household Appliance on its own. It is perhaps for this reason that the distinctiveness effect was less pronounced for low-overlap category pairs.

Hampton (1987) found that attribute inheritance in conjunctions was largely constrained by a *necessity constraint* and an *impossibility constraint*. Where an attribute was judged as necessary (or impossible) for a constituent class, then it also tended to be judged as necessary (or impossible) for the conjunction. The equivalent analysis on the present data showed relatively few cases where these constraints were broken. Where cases did violate these two constraints, this was attributed to the effects of the flexibility of concepts (Hampton, 1997b) and context sensitivity in the meaning of both the category terms and the attributes.

Although no very systematic analysis was possible, it appeared that subjects could change the meaning of category terms and attributes between narrow and broad interpretations. A Vehicle could be thought of as necessarily mechanical, but if asked to consider vehicles that were not machines, a more abstract sense of vehicle as being any means of transport replaced the more narrow conception. Attributes could also change the scope of their truth conditions in subtle ways. Thus *is lived in* would be a commonly considered necessary attribute of Dwellings. But, of course, when individual dwellings are considered, there are many which are not currently lived in. There is thus a subtle shift from considering the prototype intensional concept of a Dwelling as an ideal type of a place where people live to the consideration of the *extensional* set of actual objects which have been constructed as dwellings. This shift between intensional and extensional modes of thinking about categories is of fundamental importance in theories of concept representation

(Hampton, 1993, in press) and other cognitive skills (e.g., probabilistic reasoning; Tversky & Kahneman, 1983).

What has been learned from these studies about the action of negation on concept representations? Two accounts remain as viable (and not necessarily incompatible) explanations of the data. Intensional models (Hampton, 1987, 1988b; Smith et al., 1988) propose that the formation of a complex concept involves the merging of attribute information into a composite representation and the subsequent modification of that information in order to render the representation coherent as a concept. Often this may involve the creation of novel information in order to improve the coherence of the complex concept. Negation appears to affect the creation of this composite through inheritance by the composite of negated attributes of the negated constituent concept. Not all attributes are negated, however, and the choice of which are inherited as positive and which as negative is not a simple matter of the frequency of occurrence of such attributes in the two constituent classes. An appeal could be made to background causal theories (Murphy, 1988; Rips, 1995) to account for why some attributes are negated and others are not. The results, particularly from Experiment 1, however, are not consistent with Smith et al.'s (1988) proposal that logical functions such as conjunction and negation as applied to natural concepts should follow the prescriptive rules of set logic. It is probably true that most individuals can exercise logical thought, given sufficient motivation, time, and (perhaps) education. What is apparent from studies such as those presented here is that when interpreting phrases in natural language, such as "Tools which are Weapons" or "Vehicles which are not Machines," logical set theory provides a very poor account of people's categorization behavior.

Aside from the composite prototype account, there is also considerable evidence in the data presented here for a model of concept conjunctions, based on the identification of *exemplars* that fall in the extension of the complex concept. Particularly for the negated conjunction concepts, the results show that people may well be inclined to use this more extensional strategy. For example, in Experiment 1, negated conjunctions (but not conjunctions) showed the rescaling effect proposed by Jones (1982), whereby the most typical exemplars of the negated conjunction class were judged to be highly typical of that class. Negated conjunctions (but not conjunctions) also showed double overextension effects where an item was more likely to be judged in the complex class than in either constituent. In Experiment 2, there was also evidence for exemplar-based effects in the considerable number of instantiations generated as attributes of the negated conjunctions. Support for a more extensional account of concept conjunction is also presented by Storms et al. (1993). It is noteworthy that there was more evidence for mediation by exemplars for the negated conjunctions than for the others. Perhaps as linguistic phrases become more complex in form, exemplar retrieval becomes a preferred strategy for handling their interpretation.

In conclusion, this article has presented new data on the operation of negation within the process of constructing relative clause conjunctions. Typicality and class membership in negated conjunctions appeared to follow broadly the same pattern of results as for unnegated conjunctions. Attribute inheritance, however, was clearly more complex, and the data presented here represent just a preliminary exploration of the processes whereby people generate complex concepts using negation.

REFERENCES

- ACHEN, C. H. (1982). *Interpreting and using regression*. Beverly Hills, CA: Sage.
- ARMSTRONG, S. L., GLEITMAN, L. R., & GLEITMAN, H. (1983). What some concepts might not be. *Cognition*, **13**, 263-308.
- ASHBY, F. G., & GOTT, R. E. (1988). Decision rules in the perception and categorization of multidimensional stimuli. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **14**, 33-53.
- BARSALOU, L. W. (1983). Ad hoc categories. *Memory & Cognition*, **11**, 211-227.
- CHATER, N., LYON, K., & MYERS, T. (1990). Why are conjunctive concepts overextended? *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **16**, 497-508.
- COHEN, B., & MURPHY, G. L. (1984). Models of concepts. *Cognitive Science*, **8**, 27-58.
- FODOR, J. A. (1994). Concepts—a pot-boiler. *Cognition*, **50**, 95-113.
- HAMPTON, J. A. (1987). Inheritance of attributes in natural concept conjunctions. *Memory & Cognition*, **15**, 55-71.
- HAMPTON, J. A. (1988a). Disjunction of natural concepts. *Memory & Cognition*, **16**, 579-591.
- HAMPTON, J. A. (1988b). Overextension of conjunctive concepts: Evidence for a unitary model of concept typicality and class inclusion. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **14**, 12-32.
- HAMPTON, J. A. (1991). The combination of prototype concepts. In P. Schwanenflugel (Ed.), *The psychology of word meanings* (pp. 91-116). Hillsdale, NJ: Erlbaum.
- HAMPTON, J. A. (1993). Prototype models of concept representation. In I. van Mechelen, J. A. Hampton, R. S. Michalski, & P. Theuns (Eds.), *Categories and concepts: Theoretical views and inductive data analysis* (pp. 67-95). London: Academic.
- HAMPTON, J. A. (1996). Conjunctions of visually based categories: Overextension and compensation. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **22**, 378-396.
- HAMPTON, J. A. (1997a). Conceptual combination. In K. Lamberts & D. Shanks (Eds.), *Knowledge, concepts and categories* (pp. 135-162). Hove, U.K.: Psychology Press.
- HAMPTON, J. A. (1997b). 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.
- HAMPTON, J. A. (in press). Psychological representation of concepts. In M. A. Conway & S. E. Gathercole (Eds.), *Cognitive models of memory*. Hove, U.K.: Psychology Press.
- HUTTENLOCHER, J., & HEDGES, L. V. (1994). Combining graded categories: Membership and typicality. *Psychological Review*, **101**, 157-165.
- JONES, G. V. (1982). Stacks not fuzzy sets: An ordinal basis for prototype theory of concepts. *Cognition*, **12**, 281-290.
- KATZ, J. J., & FODOR, J. A. (1963). The structure of a semantic theory. *Language*, **39**, 170-210.
- KUNDA, Z., MILLER, D. T., & CLARE, T. (1990). Combining social concepts: The role of causal reasoning. *Cognitive Science*, **14**, 551-578.
- LAKOFF, G. (1987). *Women, fire and dangerous things*. Chicago: University of Chicago Press.
- MEDIN, D. L., GOLDSTONE, R. L., & GENTNER, D. (1993). Respects for similarity. *Psychological Review*, **100**, 254-278.
- MILLER, G. A., & JOHNSON-LAIRD, P. N. (1976). *Language and perception*. Cambridge, MA: Harvard University Press.
- MINSKY, M. (1975). A framework for representing knowledge. In P. H. Winston (Ed.), *The psychology of computer vision* (pp. 211-277). New York: McGraw-Hill.
- MURPHY, G. L. (1988). Comprehending complex concepts. *Cognitive Science*, **12**, 529-562.
- MURPHY, G. L., & SPALDING, T. (1995). Knowledge, similarity, and concept formation. *Psychologica Belgica*, **35**, 127-144.
- 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. (1982). Gradedness and conceptual conjunction. *Cognition*, **12**, 299-318.
- RIPS, L. J. (1995). The current status of research on concept combination. *Mind & Language*, **10**, 72-104.
- ROSCH, E. (1978). Principles of categorization. In E. Rosch & B. Lloyd (Eds.), *Cognition and categorization* (pp. 27-48). Hillsdale, NJ: Erlbaum.
- ROSCH, E., & MERVIS, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, **7**, 573-605.
- SMITH, E. E., & OSHERSON, D. N. (1984). Conceptual combination with prototype concepts. *Cognitive Science*, **8**, 337-361.
- SMITH, E. E., OSHERSON, D. N., RIPS, L. J., & KEANE, M. (1988). Combining prototypes: A selective modification model. *Cognitive Science*, **12**, 485-527.
- STORMS, G., DE BOECK, P., VAN MECHELEN, I., & GEERAERTS, D. (1993). Dominance and noncommutativity effects in concept conjunctions: Extensional or intensional basis? *Memory & Cognition*, **21**, 752-762.
- STORMS, G., DE BOECK, P., VAN MECHELEN, I., & RUTS, W. (1996). The dominance effect in concept conjunctions: Generality and interaction aspects. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **22**, 1266-1280.
- THAGARD, P. (1983, June). *Conceptual combination: A frame-based theory*. Paper presented at the meeting of the Society for Philosophy and Psychology, Wellesley, MA.
- THAGARD, P. (1995, May). *Conceptual combination, coherence and creativity*. Paper presented at the Creative Concepts Conference, Texas A&M.
- TVERSKY, A. (1977). Features of similarity. *Psychological Review*, **84**, 327-352.
- TVERSKY, A., & KAHNEMAN, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, **90**, 293-315.
- ZADEH, L. (1965). Fuzzy sets. *Information & Control*, **8**, 338-353.
- ZADEH, L. (1982). A note on prototype theory and fuzzy sets. *Cognition*, **12**, 291-297.

NOTES

1. It is important to distinguish between two rather different forms of conceptual combination. Osherson and Smith's (1981) critique was concerned with the *logical* functions of conjunction, disjunction, negation, and class inclusion as applied to prototype concepts. More recently, a new line of research has emerged which examines noun-noun combinations as conceptual combinations (for example, *apartment dog* or *expert repair*; see Murphy, 1988). The interpretation of such noun phrases is beyond the scope of simple logical connectives, and so will not be the focus of the present paper.
2. Chater et al. (1990) replicated the Hampton (1988b) study using simple yes/no category decisions, and found the same pattern of results as Hampton had found using the combined membership/typicality scale.
3. For example, if 50% believe a tomato to be a Fruit and 60% classify it as a Vegetable, there must be at least 10% who believe it is both, to go with 40% who think it is just a Fruit and 50% who think it is just a Vegetable.
4. Although treating the rating scale as an interval scale carries with it strong scaling assumptions (not least of which is that a unit of typicality is equivalent to a unit of relatedness), the technique was considered useful for the following reasons. First, the theoretical position

adopted assumes that people base categorization on a continuous underlying scale of similarity to the concept prototype. If this is so, then the constructed scale should be monotonically related to that underlying similarity. Second, previous use of the technique (Hampton, 1988a, 1988b, 1996; Storms et al., 1993; Storms et al., 1996) has provided a useful indicator of phenomena such as dominance and noncommutativity, and the purpose of the current experiment was to examine the generality of these phenomena.

5. It is arguable whether standardized (beta weights) or unstandardized regression weights are more appropriate measures of dominance in the equations (Achen, 1982, p. 68). Since the two independent variables are defined on the same scale, the unstandardized weight is probably more valid, as the beta weight is confounded with sampling differences in the variance of the variables within the (nonrandom) sample of items. The major conclusions to be drawn do not critically depend on which measure is used.

6. While subjects could not apparently generate any highly typical members of the class of "Furniture which is a Household Appliance," it is of course possible that they could do so if they were asked to imagine novel classes of object that possessed many of the attributes of each class. The concepts are thus not necessarily empty centered in all possible worlds only in the current world of known exemplars (see Hampton, 1997b).

7. Analysis of data provided by Storms et al. (1996) confirmed the low degree of double overextension for conjunctions. Only 6% of their items showed overextension relative to both constituents.

8. The full set of attributes is available upon request from the author.

9. It is notable that the measure of *relative* overlap between concepts tallied very closely with their relative dominance in Experiment 1.

Dominant concepts tended to be those that had a higher proportion of their members included within the other category. This measure has been shown to predict dominance quite generally across a large sample of concept pairs (Storms et al., 1996).

10. Note that it is very hard to consider the converse case of attributes that are generally untrue of *not* B—presumably they could be attributes that are true of B. However, the problem is that whether an attribute is true of *all* B, *some* B, or *no* B allows no inference at all about the attribute's frequency for A that are not B. Knowing that x% of Pets have attribute F, we can make no sensible prediction of the proportion of "Animals which are not Pets" which may have F.

11. It is necessary for an attribute to be actually *absent* from the head class in order for it to count as emergent. There is no inconsistency in an attribute's having a frequency in a subset that is much higher than its frequency in the general class. Hence, a strict criterion was adopted of requiring the majority of subjects to give a 0% rating.

12. Since attributes were not generated to class A alone in this experiment, there may be more of these attributes than those included here—only those listed by subjects negatively for "B which are not A" or positively for "A which are not B" are included. Examination of the attributes generated for the same categories in the study by Hampton (1987) suggests that there could have been several other disappearing attributes. For example, for "Vehicles which are not Machines" the following may have been considered universally true of Vehicles and yet not been inherited: *is made of metal, needs maintenance, is technological, and replaces people*. Rates of incidence of inheritance failure of necessary attributes may therefore be underestimated here.

APPENDIX
Items With Categorization Probability in a Conjunction at Least .25 Greater Than
Minimum Constituent Categorization Probability in Experiment 1

Category A	Category B	Item	Difference			
			A	B	C	
Conjunctions						
PETS	BIRDS	Peacock	.35	1.0	.64	.29
		Lark	.21	1.0	.58	.37
		Eagle	.15	.95	.42	.27
		Cuckoo	.41	.85	.80	.39
TOOLS	WEAPONS	Chisel	1.0	.55	.85	.30
		Drill	1.0	.35	.66	.31
		Saw	1.0	.35	.66	.31
		Club	.35	1.0	.74	.39
		Spear	.26	1.0	.67	.41
		Whip	.25	1.0	.51	.26
		Arrow	.05	1.0	.56	.51
		Razor	.5	.9	.82	.32
MACHINES	VEHICLES	Lawn mower	1.0	.15	.67	.52
		Elevator	.9	.53	.82	.29
		Skateboard	.15	.8	.42	.27
FURNITURE	HOUSEHOLD APPLIANCES	Can opener	.1	.95	.51	.41
BUILDINGS	DWELLINGS	Trailer	.35	1.0	.66	.31
		Cave	.05	.94	.59	.54
Negated Conjunctions						
			A	1-B	C	
PETS	NOT BIRDS	Elephant	.1	1.0	.7	.6
BIRDS	NOT PETS	Robin	1.0	.7	.95	.25
		Toucan	1.0	.26	.71	.45
		Mynah bird	1.0	.12	.47	.35
		Homing pigeon	1.0	.05	.47	.42
		Hammer	1.0	.2	.5	.3
TOOLS	NOT WEAPONS	Screwdriver	1.0	.15	.6	.45
		Axe	1.0	.05	.3	.25

APPENDIX (Continued)

Category A	Category B	Item	Difference					
			A	1-B	C			
WEAPONS	NOT TOOLS	Club	1.0	.65	.95	.3		
		Knife	1.0	.1	.37	.27		
MACHINES	NOT	Road roller	1.0	.16	.59	.43		
	VEHICLES	Bulldozer	1.0	.1	.4	.3		
		Dog sled	.2	.05	.3	.25*		
FURNITURE	NOT	Chair	1.0	.7	.95	.25		
	HOUSEHOLD	Heated waterbed	1.0	.58	.85	.27		
		APPLIANCES	Desk lamp	.8	.32	.7	.38	
	Television		.75	.1	.45	.35		
	Cooking stove		.55	.0	.35	.35		
	Painting		.4	1.0	.73	.33		
	Closet		.1	.9	.42	.32		
	Floor mat		.35	.75	.67	.32		
HOUSEHOLD APPLIANCES	NOT	Television	.9	.25	.89	.64		
	FURNITURE	Cooking stove	1.0	.45	.95	.5		
		Refrigerator	1.0	.45	.95	.5		
		Hifi	.7	.3	.88	.58*		
		Desk lamp	.68	.2	.74	.54*		
		Clothes washer	1.0	.6	1.0	.4		
		Vacuum cleaner	1.0	.74	1.0	.26		
		SPORTS	NOT GAMES	Canoeing	1.0	.4	.85	.45
Rowing	1.0			.25	.58	.33		
Boxing	1.0			.2	.45	.25		
Mountaineering	.95			.55	.85	.3		
Judo	.9			.3	.7	.4		
Stamp collecting	.05			.63	.3	.25		
GAMES	NOT SPORTS			Bridge	1.0	.75	1.0	.25
				Chess	.95	.6	.89	.29
		Darts	.95	.37	.74	.37		
		Dominoes	.9	.75	1.0	.25*		
		Synagogue	1.0	.5	.75	.25		
BUILDINGS	NOT DWELLINGS	Barn	1.0	.15	.5	.35		
		Castle	1.0	.0	.26	.26		
		Igloo	.75	.0	.25	.25		
		Phone box	.11	.95	.6	.49		
		DWELLINGS	NOT	Hut	1.0	.35	.79	.44
BUILDINGS	Igloo		1.0	.25	.84	.59		
	Tent		.75	.75	1.0	.25*		
	Tree house		.74	.2	.89	.69*		
	Phone-box		.05	.89	.31	.26		

Note—A = $p(A)$ for constituent category A. B = $p(B)$ for constituent category B. C = $p(A\&B)$ or $p(A\&\neg B)$ for negated conjunctions. Difference = C minus the minimum of A and B. For negated conjunctions, 1-B is $1-p(B)$. Asterisks indicate double overextension.

(Manuscript received October 31, 1995;
revision accepted for publication July 15, 1996.)