Journal of Experimental Psychology:
Learning, Memory, and Cognition
2004, Vol. 30, No. 3, 573–582

0278-7393/04/$12.00 DOI: 10.1037/0278-7393.30.3.573

Typical Versus Atypical Unpacking and Superadditive Probability Judgment

Steven Sloman
Brown University

Yuval Rottenstreich
University of Chicago Graduate School of Business

Edward Wisniewski
University of North Carolina at Greensboro

Constantinos Hadjichristidis
University of Plymouth

Craig R. Fox
University of California, Los Angeles

Probability judgments for packed descriptions of events (e.g., *the probability that a businessman does business with a European country*) are compared with judgments for unpacked descriptions of the same events (e.g., *the probability that a businessman does business with England, France, or some other European country*). The prediction that unpacking can decrease probability judgments, derived from the hypothesis that category descriptions are interpreted narrowly in terms of typical instances, is contrasted to the prediction of support theory that unpacking will generally increase judged probabilities (A. Tversky & D. J. Koehler, 1994). The authors varied the typicality of unpacked instances and found no effect of unpacking with typical instances (*additivity*) and a negative effect with atypical instances (*superadditivity*). Support theory cannot account for these findings in its current formulation.

A basic normative requirement for likelihood judgment is the notion of *description invariance:* The judged probability of an event should not depend on how the event is described. Empirically, however, alternative descriptions of the same event can produce systematically different judgments. For instance, Rottenstreich and Tversky (1997) asked respondents to judge the probability that a randomly selected death was a murder and obtained higher probabilities when this event was described as "homicide by an acquaintance or stranger" than when it was merely described as "homicide." When distinct exemplars of homicide were made salient, judged probability increased.

Support theory (Rottenstreich & Tversky, 1997; Tversky & Koehler, 1994) addresses how the description of a category affects judged probability. One of the theory's foundational assumptions is that the perceived probability of an event generally increases when the event's description is unpacked into a disjunction of component events (as in the homicide example). Such an increase is referred to as *implicit subadditivity* (as is a decrease caused by unpacking a contrasting event). In this article, we observe that the evidence for implicit subadditivity is mixed, and we present several examples of the opposite pattern, *implicit superadditivity,* observed when unpacking an event decreases its perceived probability. We show that an analysis of the way that people represent category structure helps one to characterize the circumstances that lead to implicit subadditivity, additivity, or implicit superadditivity. In particular, we find that the probability assigned to an unpacked description is proportional to the typicality of instances used to unpack the category, that is, to the degree that they are good representatives of the category being judged.

We begin by describing the major features of support theory and reviewing extant evidence for implicit subadditivity. Next, we present a novel interpretation of unpacking that we call the *narrow interpretation conjecture,* and we provide evidence for it, mainly in the form of implicit superadditivity. We conclude with a discussion of the implications for support theory and for probability judgment in general.

Support Theory

The Support Scale Mediates Judged Probability

Support theory attaches subjective probability not to events, as do most theories of subjective probability, but rather to descriptions of events, called *hypotheses.* The theory associates each hypothesis A with a nonnegative support value, $s(A)$, that corresponds to the perceived strength of evidence for that hypothesis.

Steven Sloman, Cognitive and Linguistic Sciences, Brown University; Yuval Rottenstreich, Center for Decision Research, University of Chicago Graduate School of Business; Edward Wisniewski, Department of Psychology, University of North Carolina at Greensboro; Constantinos Hadjichristidis, Department of Psychology, University of Plymouth; Craig R. Fox, The Anderson School, University of California, Los Angeles.

This work was funded by National Aeronautics and Space Administration Grant NCC2-1217, an American Philosophical Society Sabbatical Fellowship to Steven Sloman, and National Science Foundation Grant BCS-9975198 to Ed Wisniewski. We thank Marianne Harrison, Erica Middleton, and Kelly Trindel for assistance in conducting the experiments, and Louis Narens for comments on an earlier version of this article.

Correspondence concerning this article should be addressed to Steven Sloman, Cognitive and Linguistic Sciences, Brown University, Box 1978, Providence, RI 02912. E-mail: steven_sloman@brown.edu

The probability that hypothesis A holds rather than hypothesis B, assuming exactly one of them obtains, is given by the following:

$$P(A, B) = \frac{s(A)}{s(A) + s(B)}. \quad (1)$$

Thus, the theory posits the construct of support as an intermediate component of judgment (i.e., hypotheses → support → probability). Judged probability is interpreted as an expression of the proportion of total evidence favoring the focal hypothesis A rather than the alternative B.

Subadditivity of Support

Support theory imposes specific constraints on the support scale. In particular, the support of a hypothesis A (e.g., homicide) is assumed to be less than the sum of the support of (exclusive and exhaustive) constituent hypotheses A_1 and A_2 (e.g., homicide by an acquaintance, homicide by a stranger):

$$s(A) \leq s(A_1) + s(A_2). \quad (2)$$

This pattern, known as *subadditivity*, implies a corresponding condition on judged probabilities: $P(A, \neg A) \leq P(A_1, \neg A_1) + P(A_2, \neg A_2)$, where $\neg$ indicates the negation of its associated hypothesis. That is, the judged probability of an event is generally less than the sum of probabilities of its constituents. Furthermore, the theory decomposes subadditivity into two parts: (a) the effect of separately describing constituent hypotheses and (b) the effect of separately evaluating these constituents. First, the theory assumes that unpacking a hypothesis A into a disjunction of exclusive constituents, $A_1 \vee A_2$, generally increases support. That is, when (A_1, A_2) is recognized as a partition of A,

$$s(A) \leq s(A_1 \vee A_2). \quad (2A)$$

This relation implies *implicit subadditivity*, the corresponding pattern for judged probability: Unpacking the focal hypothesis increases judged probability, $P(A, \neg A) \leq P(A_1 \vee A_2, \neg A)$, whereas unpacking the alternative hypothesis decreases judged probability, $P(\neg A, A) \geq P(\neg A, A_1 \vee A_2)$. Second, the theory assumes that the evaluation of an unpacked hypothesis yields less support than does separate evaluation of its components. That is,

$$s(A_1 \vee A_2) \leq s(A_1) + s(A_2). \quad (2B)$$

This pattern implies *explicit subadditivity* (Rottenstreich & Tversky, 1997), a corresponding condition on judged probabilities: $P(A_1 \vee A_2, \neg A) \leq P(A_1, A_2 \vee \neg A) + (A_2, A_1 \vee \neg A)$.

Experimental Evidence for Subadditivity

Subadditivity (Equation 2) and explicit subadditivity (Equation 2B) appear to be sizable and robust phenomena. They have been observed in numerous studies reviewed by Tversky and Koehler (1994) and Fox and Tversky (1998; see also Brenner & Rottenstreich, 1999). Subadditivity, for example, has been documented in studies of medical doctors making diagnoses (Redelmeier, Koehler, Liberman, & Tversky, 1995), professional options traders forecasting the future prices of specific stocks (Fox, Rogers, & Tversky, 1996), expert sports fans predicting game outcomes (Fox, 1999), and attorneys forecasting trial outcomes (Fox & Birke, 2002). Subadditivity has also been observed in willingness to pay for hypothetical insurance policies (Johnson, Hershey, Meszaros, & Kunreuther, 1993).[1]

The evidence for implicit subadditivity is less consistent. In their first experiment, Rottenstreich and Tversky (1997) obtained implicit subadditivity among participants forecasting a trial outcome in which the hypothesis *not a guilty verdict* was unpacked into *a not guilty verdict or a hung jury.* However, these authors failed to register a significant effect when unpacking the hypothesis *the winner of the next presidential election will not be a Democrat* into *the winner of the next presidential election will be a Republican or an Independent rather than a Democrat.* In their second experiment, Rottenstreich and Tversky found implicit subadditivity when the hypothesis that a randomly selected death is due to *homicide* was unpacked into *homicide by an acquaintance or homicide by a stranger*, a description designed to promote attention to different causes of homicide. However, they registered no implicit subadditivity when the hypothesis *homicide* was unpacked into *daytime homicide or nighttime homicide,* a description designed to provide no such advantage.

Fox and Tversky (1998) reported that judged probabilities increased when the hypothesis *an Eastern Conference team wins the NBA championship* was unpacked into a list that explicitly mentioned the four strongest teams in the conference. However, they found no implicit subadditivity when the Western Conference was similarly decomposed. Fox and Birke (2002) found that judgments by experienced attorneys increased when the hypothesis that the Microsoft antitrust case would *go directly to the Supreme Court* was unpacked into *go directly to the Supreme Court and be affirmed, reversed, or modified.* However, these authors reported only a nonsignificant tendency toward implicit subadditivity in two other experiments. Fox and See (2003) reported strong implicit subadditivity in judgments of the probability that various categories of teams would win a conference basketball championship (e.g., *A school from North Carolina* vs. the unpacked *A public or private school from North Carolina*). These authors also found significant implicit subadditivity when specific scenarios (e.g., *Duke beats Virginia*) were unpacked into elaborated scenarios through conjunction with a second event (e.g., *Duke beats Virginia and beats Clemson or beats Virginia and loses to Clemson*). However, these authors failed to observe implicit subadditivity when scenarios (e.g., *Duke beats UNC*) were unpacked into disjunctions of dimensional subhypotheses (e.g., “Duke beats UNC by 1–15 points or more than 15 points”). Likewise, Fox and Clemen (2003) reported that business students did not show implicit subadditivity when events were unpacked into obvious constituents (e.g., when “a school other than Wharton” was unpacked into “Chicago, Harvard, Kellogg, Stanford, or another school other than Wharton” when judging the top-rated business school in the next *Business Week* survey), though participants did exhibit strong

[1] An exception under a special condition is reported by Macchi, Osherson, and Krantz (1999). They found that unpacking low-support instances resulted in violations of binary complementarity. For example, the judged probability that *the 1995 birthrate in Burma was greater than that of Thailand* was less than .5 as was that of its complement, a form of explicit superadditivity.

explicit subadditivity when each of these possibilities was evaluated separately and their probabilities summed.

Koehler, Brenner, and Tversky (1997) observed implicit subadditivity in an experiment that unpacked the alternative hypothesis. To illustrate, judged probabilities were higher for the hypothesis that an individual "majors in Economics rather than another social science" than for the hypothesis that the same individual "majors in Economics rather than Political Science, Psychology, or Sociology."

Some older evidence can also be interpreted as support for implicit subadditivity. Tversky and Kahneman (1983) obtained higher ratings of likelihood for "Linda is a bank teller whether or not she is active in the feminist movement" than for "Linda is a bank teller." Assuming the first description would be treated as equivalent to "Linda is a feminist bank teller or Linda is a nonfeminist bank teller," it serves as an unpacking of the second description.

Thus, a review of the literature reveals strong and consistent evidence for generic subadditivity (Equation 2) and explicit subadditivity (Equation 2B) but inconsistent evidence for implicit subadditivity (Equation 2A). Hadjichristidis, Sloman, and Wisniewski (2001) and Hadjichristidis, Stibel, Sloman, Over, and Stevenson (1999) provided preliminary evidence that selectively unpacking hypotheses into components that enjoy low levels of support results in implicit superadditivity, wherein the judged probability of the packed hypothesis is greater than that of the unpacked one. To illustrate, students gave higher probability estimates for the packed hypothesis "death from a natural cause" than for its coextensional unpacked counterpart "death from asthma, the flu, or some other natural cause." These studies, however, do not provide for a unique explanation of implicit superadditivity. For one, they fail to distinguish whether the effect is cognitive, resulting from the form of mental representations of categories, or pragmatic, resulting from an inference about the intended referent of the category label. The purpose of this article is to distinguish these alternative accounts in the course of clarifying the relation between descriptions and probability judgments. In the process, we extend the range of conditions for observing implicit superadditivity.

Accounting for Unpacking Effects: The Narrow Interpretation Conjecture

In light of the paucity of evidence for implicit subadditivity, why has support theory taken its existence for granted? Consider Tversky and Koehler's (1994) original account of subadditivity:

> When people assess their degree of belief in an implicit disjunction . . .they tend to form a global impression that is based primarily on the most representative or available cases. Because this mode of judgment is selective rather than exhaustive, unpacking tends to increase support. . . .Unpacking a category. . .might remind people of possibilities that would not have been considered otherwise. Moreover the explicit mention of an outcome tends to enhance its salience and hence its support. (p. 549)

Tversky and Koehler note that unpacking a category into a list of exemplars might remind people of possibilities they would not have otherwise considered. On the assumption that increasing the availability of possibilities increases support, this reminding-based analysis implies that unpacking will always result in either implicit additivity or implicit subadditivity but never in implicit superadditivity.

However, the appeal to salience reduces the motivation for the subadditivity prediction in some cases. If the unpacked instances have low support, then increased attention to them at the cost of decreased attention to other instances should reduce probability judgment and might even lead to superadditivity.

We agree with Tversky and Koehler (1994) that judgments of support for a category are based primarily on the most representative category instances (see also Koehler, Brenner, & Tversky, 1997), and because categories are usually represented in a manner consistent with their most typical exemplars (see Murphy, 2003; Rosch & Mervis, 1975). For instance, when asked to evaluate the probability that a randomly selected death is due to "a disease," one might assess the relative frequency of this event by the ease with which instances come to mind (Tversky & Kahneman, 1973). Because deaths are normally attributed to specific maladies, the judge will spontaneously unpack the generic hypothesis "disease" into a limited set of typical exemplars, instances (e.g., heart attack, cancer, and stroke) that are frequent and that manifest the more common features of the category.

Typical instances tend to be the best representatives of a category by virtue of being relatively common and the most similar to other category members (Hampton, 1998). Because frequency is also an important source of assessments of support and because instances that are similar to other instances are easier to elicit from memory, typicality is correlated with support: Good category representatives tend to be judged more probable than bad ones. But exceptions do exist. For example, cardinals are typical birds in the western United States despite their rarity, whereas the Canada Goose is atypical despite being common (in winter). More generally, the primary determinant of support is typicality or representativeness, but it is not the only one. Note that typicality gradients exhibit as much reliability for ad hoc categories like "things to take out of a burning house" as they do for categories that are well established in memory, though their determinants can differ (Barsalou, 1985).

We also agree with Tversky and Koehler (1994) that retrieved instances are used to generate judgments of support. However, increasing the availability of instances in memory is not a sufficient condition to increase probability judgment. Hadjichristidis, Sloman, and Wisniewski (2001) showed that merely presenting instances immediately prior to probability judgment—without using them to unpack the judged category—had no effect on judgments. As suggested by Tversky and Koehler, unpacking a hypothesis into a list of explicitly mentioned components can focus attention on the listed components. This presumably explains subadditivity in cases in which the sample space involves a small number of specified outcomes that are unlikely to be forgotten (e.g., Brenner & Koehler, 1999; Koehler et al., 1997; Koehler, White, & Grondin, 2003). But focus can also inhibit the judge's capacity to consider components that are not listed, especially for large categories with many potential elements to retrieve. The focusing of attention may occur because explicitly listed constituent events tax working memory so that the attentional resources for further unpacking are unavailable. Moreover, explicitly naming a subset of elements from a category may inhibit subsequent recall of other elements drawn from that category, an episodic memory effect known as *part-set cuing* (Slamecka, 1968) that has also been shown in semantic memory (Nickerson, 1984; Sloman, 1991). Whatever the underlying mechanism, the effect of focusing depends on the nature of the listed components. The judged probability of an unpacked category should be proportional to the support of the unpacked instances.

We refer to the combination of suppositions that categories are represented via typical instances and that unpacked instances serve as the focus of judgment as the *narrow interpretation* conjecture. It yields several new predictions. First, unpacking a category into a small number of typical exemplars should generally yield additivity. The judged probability of the packed category should approximately equal the judged probability of a typical unpacking because the packed category is already assumed to be interpreted in terms of typical exemplars. Second, unpacking a category into atypical exemplars with weak support should generally yield superadditivity. These exemplars will capture attention from more typical exemplars and thus lower support and judged probability. Finally, when a hypothesis is unpacked into atypical exemplars that have stronger support than typical exemplars, we expect implicit subadditivity. In the following experiments, we tested these three predictions.

Experiment 1: Implicit Superadditivity and Additivity (Diseases)

In our first experiment we tested our superadditivity prediction in a situation in which we presumed that participants recruit support through availability of instances. When relying on the availability heuristic, people assess support by evaluating the ease with which instances come to mind (Tversky & Kahneman, 1973). We asked University of Chicago undergraduates ($N = 137$) to complete the following written item:

> Consider all the people that will die in the U.S. next year. Suppose that we select one of these people at random. Please estimate the *probability* that this person's death will be attributed to the following causes.

The packed group ($n = 49$) judged the probability of the hypothesis "disease." The typical group ($n = 53$) judged the probability of a hypothesis mentioning the three most common causes of death and a residual category: *heart disease, cancer, stroke, or any other disease.* Finally, the atypical group ($n = 35$) judged the probability of a hypothesis listing relatively less common causes of death and a residual category: *pneumonia, diabetes, cirrhosis, or any other disease.*

Consistent with our predictions, the median judged probability was not significantly higher in the typical group (.60) than in the packed group (.55). Moreover, the median packed judgment was significantly higher than the median atypical judgment (.40; $p = .05$ by Mann–Whitney). That is, unpacking into typical causes of death yielded additivity, whereas unpacking into atypical causes of death with weak support (plus a catch-all category) yielded implicit superadditivity. Including low support, atypical instances in the judged category gave them substantially more weight in the judgment process.

Experiment 2: Implicit Superadditivity and Additivity While Controlling for Pragmatic Bias

In the following experiment we examined the possibility of implicit additivity and superadditivity using larger samples of stimuli covering wider spectra of categories. To review, Tversky and Koehler's (1994) reminding-based account of unpacking predicts universal implicit subadditivity. The narrow interpretation conjecture suggests that additivity generally obtains for typical unpackings but superadditivity obtains for unpackings into atypical components with weak support.

Unpacking a category may not only change how people think about the category being judged; it can also change the participant's interpretation of the experimenter's intended meaning of the category label. For example, unpacking the category *article of clothing* into *necklace or other article of clothing* may bias people to reinterpret articles of clothing to include jewelry. If the question asked about the likelihood that a husband would buy an "article of clothing" for his wife as a wedding gift, the unpacked version might produce subadditivity for pragmatic reasons (jewelry is more representative of a wedding gift than articles of clothing). If the question asked about the likelihood that a father would buy an "article of clothing" for his son as a birthday gift, an unpacking using "necklace" might produce superadditivity (jewelry is more unrepresentative of a birthday gift to a son than articles of clothing). This experiment evaluated the pragmatic bias account by using well-defined, ad hoc categories with unambiguous extensions.

As before, typicality was varied. This time, we provided an independent measure of it. We also varied the number of instances into which categories were unpacked. Unpacking into two instances brings more instances to mind and increases the salience of more instances than unpacking into a single instance. Therefore, the reminding hypothesis predicts more subadditivity with two rather than one unpacked instance, irrespective of the typicality value of those instances. The narrow interpretation conjecture implies that the number of unpacked exemplars is unimportant relative to how typical they are. Thus, all else being equal, the conjecture implies that there should be no difference between unpackings that include one versus two elements.

Method

Participants. The participants were 124 undergraduates from the University of North Carolina at Greensboro who participated in the experiment for course credit. Each participant completed both a probability judgment task and an assessment of typicality.

Materials and procedure. The stimuli consisted of nine well-defined categories and four members of each category (see Table 1). We selected category members such that two were typical of the category and two atypical, according to our intuitions. Each category was well defined in the sense that membership in the category was clear and unambiguous. For example, one category was *seven letter words that begin with a consonant.*

For the probability judgment task, participants were asked to report a percentage ranging from 0% to 100% for nine hypotheses drawn from one of the following five conditions: packed, single typical unpacked, double typical unpacked, single atypical unpacked, double atypical unpacked. Table 2 shows an example of one question from each of the five conditions. In each condition, half of the participants were given the nine questions in one ordering, whereas the other half of the participants were given the questions in the opposite order.

Table 1
Ad Hoc Categories in Experiment 2 and Their Typical and Atypical Instances

Category	Instance: Typical	Instance: Atypical
1. Times that classes normally start before 8:45 am	8:15, 8:30	8:07, 8:23
2. Heights of NBA players that are less than 6′ 8″	6′ 3″, 6′ 4″	5′ 10″, 5′ 11″
3. Weights of spaghetti packs that are less than 600 grams	250 grams, 500 grams	129 grams, 317 grams
4. Island vacation spots for New Yorkers	Jamaica, Hawaii	Japan, Ireland
5. Amount of money that an undergraduate has on them that is less than 40 dollars	5 dollars, 10 dollars	50 cents, 1 dollar
6. Seven letter words beginning with a consonant	*m, b*	*z, q*
7. Geometric shapes of the top views of buildings	Square, parallelogram	hexagon, octagon
8. European countries that US businessmen do business with	England, France	Hungary, Finland
9. Number of cookies in a box that contains some number of cookies divisible by 5	15, 20	55, 60

Immediately after completing the probability judgment task, participants answered questions about how good an example an instance was of the category. They were instructed that they would be presented with lists of examples of categories and, using their best judgment, they were to rank how good those examples are of the categories. The task was illustrated with four examples of birds: "ostrich," "robin," "sparrow," and "chicken." Participants were told that they might rank "robin" as the best example of the birds listed, "sparrow" as the second best example, "chicken" as the third best example, and "ostrich" as the fourth best or poorest example. They were also told that they might not agree with these rankings and that these rankings were used only to illustrate the task. Participants were to make their rankings by writing the number 1, 2, 3, or 4 next to each example. Participants filled out the booklet at their own pace.

Table 2
Examples of Questions From the Probability Judgment Task in Experiment 2

Packed condition
How likely do you think it is that a randomly selected word from the dictionary is a seven letter word beginning with a consonant (as opposed to a seven letter word beginning with a vowel)? ______
Single-component typical unpacking
How likely do you think it is that a randomly selected word from the dictionary is a seven letter word beginning with m or some other consonant (as opposed to a seven letter word beginning with a vowel)? ______
Double-component typical unpacking
How likely do you think it is that a randomly selected word from the dictionary is a seven letter word beginning with m, b, or some other consonant (as opposed to a seven letter word beginning with a vowel)? ______
Single-component atypical unpacking
How likely do you think it is that a randomly selected word from the dictionary is a seven letter word beginning with z or some other consonant (as opposed to a seven letter word beginning with a vowel)? ______
Double-component atypical unpacking
How likely do you think it is that a randomly selected word from the dictionary is a seven letter word beginning with z, q, or some other consonant (as opposed to a seven letter word beginning with a vowel)? ______

Results and Discussion

Participants' typicality rankings substantially agreed with our intuitions. Participants' average rankings placed 16 of the 18 typical instances (two from each category) higher than both atypical instances of the category.

Table 3 displays the mean probability ratings for the packed categories and the corresponding unpacked categories. Almost all of the findings conformed to our predictions. Superadditivity obtained for judgments of atypically unpacked categories and additivity for typical ones. Single- and double-component unpackings did not differ. These results are supported by two 2 (instances) × 2 (typicality) analyses of variance, one by participants (F_1) and one by items (F_2). There was a statistically reliable main effect for typicality by participants that was marginally reliable by items, $F_1(1, 97) = 4.78$, $p < .04$; $F_2(1, 8) = 3.64$, $p < .09$. No other effects were reliable (all Fs < 1). Two-tailed t tests compared each Instances × Typicality condition with the packed condition. The only differences that were statistically reliable were those between the single-atypical-unpacked and packed conditions: participants, $t(48) = 2.03$, $p < .05$; items, $t(8) = 2.81$, $p < .02$, and the

Table 3
Mean Probability Judgments for Packed and Unpacked-Typical and Unpacked-Atypical Categories in Experiment 2

No. of instances	Typical (%)	Atypical (%)
1	63.8	58.9
2	63.5	58.9

Note. Mean probability judgment for packed condition was 65%.

Table 4
Categories and Their Typical and Atypical Instances

Category	Instance: Typical	Instance: Atypical
1. Bottles made of glass	baby bottle	shampoo bottle
2. Seats that a randomly chosen American would be sitting on	car seat	bean bag chair
3. Belts that a randomly chosen American would be wearing	seat belt	white belt
4. Kinds of not guilty verdicts	hung jury	not guilty by reason of insanity
5. Strings that would be found in the living room of a typical house	guitar string	hemp string
6. Chairs that a randomly chosen American in an airport would be sitting in	wheelchair	lawn chair
7. Guns that you buy at a hardware store	staple gun	antique gun
8. Cards that on any given day you will receive in the mail	credit card	postcard
9. Mammals that can hold their breath	whale	weasel

double-atypical-unpacked and packed conditions: participants, $t(48) = 2.42$, $p < .02$; items, $t(8) = 2.71$, $p < .03$.[2]

In sum, the experiment replicated the superadditivity and additivity findings of Experiment 1 in a way that cannot be attributed to pragmatic bias because the categories were all well defined. Moreover, the experiment failed to confirm the hypothesis of greater subadditivity with more unpacked instances.

Experiment 3: Superadditivity and Additivity With Fuzzy Categories

Natural categories, unlike the events in the clearly defined sample spaces used in probability textbooks or Experiment 2, are characterized by lack of enumerability. Natural categories are fuzzy; they have a graded structure: Not only do instances vary in typicality but category membership may be unknown (e.g., is coral a plant or an animal?) or undefined (e.g., in what sense is a tomato a vegetable?). In Experiment 3 we attempted to replicate the superadditivity and additivity effects of Experiment 2 using fuzzy categories.

Method

Participants. The participants were 169 undergraduates from the University of North Carolina at Greensboro who participated as part of a course requirement. Specifically, 129 students participated in the probability judgment task and 40 in the typicality judgment task.

Materials. The stimuli consisted of nine fuzzy categories and two members from each category, one typical and one atypical. For example, one category was “mammals that can hold their breath for more than two minutes” with the typical member “whale” and atypical member “weasel.” Table 4 lists the categories and instances. Item 4, the hung jury item, comes from Rottenstreich and Tversky’s (1997) study.

For the probability judgment task, participants were presented with booklets containing nine questions asking them to make likelihood judgments about the packed categories, the unpacked categories with a typical instance, or the unpacked categories with an atypical instance. Table 5 lists an example of each type of question. A second form of each of the three booklets was created by reversing the order of the questions.

For the typicality judgment task, booklets were presented that contained questions about how good an example an instance was of the category. The booklet contained 18 questions, two instances of each of the nine categories. A second form of the booklet was created by reversing the order of the questions.

Procedure. In the probability judgment task, instructions asked participants to report a percentage ranging from 0% to 100% to describe how probable they believed a particular event to be, based on the information given. The task was illustrated with a simple example: “Consider a randomly selected bear. How likely do you think the bear is dark colored rather than light colored?” The instructions suggested that the answer might be fairly high (perhaps around 80%) as many bears are brown or black, though some, like polar bears, are light colored. Participants wrote responses after each question.

Participants were randomly assigned to the packed, unpacked-typical, or unpacked-atypical condition. They completed the corresponding probability judgment booklet at their own pace. Approximately equal numbers of participants were in each condition and filled out each of two forms of the corresponding booklet.

In the typicality judgment task, instructions informed participants that they would be making a series of goodness-of-example judgments. They were to describe how good an example a given item was of its category on a scale ranging from 0% (*not a good example*) to 100% (*the best example or prototype*) on the basis of the information given. The task was illustrated with a simple example involving the instances “robin” and “ostrich” for the category *bird.* Participants were told that because they may think that a robin is very similar to a prototypical bird, they may assign it a very high value close to 100% and that because they may think that an ostrich is very dissimilar to a prototypical bird, they may assign it a very low value close to 0%.

Participants filled out the booklet at their own pace. Approximately equal numbers of participants filled out each form of the booklet. Each task took about 15 min to complete.

Results and Discussion

Unlike in earlier studies (Study 2 in this article; Hadjichristidis et al., 1999, 2001), our intuitions about which category instances were typical versus atypical did not agree with participants’ judgments. Therefore, we performed a median split; we rank ordered

[2] For two of the nine categories (Items 2 and 5; see Table 1) one of our unrepresentative instances was ranked higher on representativeness than one of our representative instances. Analyses were conducted discarding these two items. Exactly the same pattern of results was obtained.

the typicality ratings of the 18 category instances from lowest to highest, classifying the 9 instances with the lowest ratings as atypical of their respective category and the 9 instances with the highest ratings as typical. Thus, instances that we had thought were atypical were reclassified as typical and vice versa. As a result of this reclassification, participants in each of the unpacked conditions ended up having made probability judgments about unpacked categories with both typical and atypical instances.[3]

Almost all of the findings conformed to our predictions. Participants exhibited superadditivity for probability judgments of unpacked categories with atypical instances. Specifically, the mean probability judgment for the unpacked-atypical categories was 46.4% compared with 64.1% for the corresponding packed categories. This difference was statistically reliable by an item analysis, $t(8) = 6.36$, $p < .01$. A small amount of subadditivity can be seen in the unpacked-typical versus packed conditions. The mean probability judgment for the unpacked typical categories was 65% compared with 56.7% for the corresponding packed categories. However, the effect was not statistically significant across items, $t(8) = 1.31$, $p < .23$. Table 6 shows the mean probability ratings for the packed descriptions (averaged over all categories, both those with typical and those with atypical instances in the unpacked conditions) and the unpacked categories. The relation between typicality and probability judgment is apparent even without using the median split to collapse typicality judgments. The correlation between mean typicality ratings and mean unpacked ratings across the 18 items was .58.

We also performed analyses across participants. We collapsed the two groups of participants who provided probability judgments for the unpacked categories and calculated the mean judgment for typical and atypical unpackings for each participant.[4] We compared the unpacked-atypical mean (46.8%) with the packed group's mean for the same items (65%). As in the item analysis, we found a highly significant superadditivity effect, $t(167) = 8.62$, $p < .01$. We also compared the unpacked-typical mean (62.2%) with the mean probability judgments of the packed group for that set of items (56.7%). This time we found a significant subadditivity effect, $t(167) = 2.38$, $p < .02$. However, the entire effect was carried by one item, Item 7, "Guns that you buy at a hardware store" with the typical instance of "staple gun." Excluding this item, typical unpackings proved additive, $t(167) < 1$, *ns.* We speculate that this item produced subadditivity because participants did not think of staple guns when considering "guns that you buy at a hardware store." Staple guns may be typical of such guns, but they are not accessible in memory, and they were likely not understood as an intended referent in the packed condition. Participants may have been judging a larger category in the unpacked-typical condition.

Table 5
Examples of a Probability Judgment From Each Condition in Experiment 3

Typical condition
How likely is it that you can buy a staple gun or some other type of gun in a hardware store?
Atypical condition
How likely is it that you can buy an antique gun or some other type of gun in a hardware store?
Packed condition
How likely is it that you can buy a gun in a hardware store?

Table 6
Mean Probability Judgments for Packed and Unpacked Categories in Experiment 3

Category	*M* (%)
Unpacked with typical member	65.0
Packed (all categories)	60.4
Unpacked with atypical member	46.4

Experiment 4: Superadditivity via Representativeness, Subadditivity via Availability

The results of Experiments 1–3 have provided evidence of implicit additivity when implicit hypotheses are unpacked into typical component hypotheses and superadditivity when they are unpacked into atypical components. However, the reminding-based hypothesis implies that unpacking a category using a component with strong support that participants would not otherwise have thought of will increase judged probabilities and therefore give rise to implicit subadditivity. To validate this claim, in the following experiment we attempted to show how different background features can lead to either superadditivity or subadditivity for the same hypotheses.

In the case of superadditivity, we used the following item based on Tversky and Kahneman's (1983) study:

> Linda is 31 years old, single, outspoken and very bright. She majored in philosophy. As a student she was deeply concerned with issues of discrimination and social justice, and also participated in anti-apartheid demonstrations.
>
> What do you think is Linda's current profession? Please rank the following professions from 1 (*most likely*) to 7 (*least likely*). That is give 1 to the profession that you think has the highest probability and 7 to the profession that has the lowest probability. Please use each rank only once.
>
> ____Linda is a teacher in elementary school, junior high, or high school.
> ____Linda works in a bookstore.
> ____Linda is a doctor or nurse.
> ____Linda is a politician.
> ____Linda is a lawyer.
> ____Linda is a psychiatric social worker.
> ____Linda is a bank teller.

[3] This reclassification resulted in three categories having two typical instances each but no atypical instances, and three categories having two atypical instances but no typical instances.

[4] As a result of the median split on typicality, for this analysis one group provided judgments on six typical instances and three atypical instances and vice versa for the other group.

The question asks for a ranking of conditional probabilities, the probabilities that Linda is in various categories given her description. We compared rankings given the list shown to rankings in a condition in which the category *lawyer* was unpacked into *tax, corporate, patent, or other kind of lawyer.* These unpacked subcategories are all atypical (i.e., unrepresentative) of Linda and therefore we expect superadditivity in this condition (higher rankings for the unpacked than for the packed description).

To reverse this effect and obtain implicit subadditivity, we replaced the description of Linda with another description adopted from Tversky and Kahneman's (1983) study:

> Bill is 34 years old. He is intelligent, but unimaginative, compulsive, and generally lifeless. In school he was strong in mathematics but weak in social studies and humanities.

Because Bill is technically oriented, the subcategories *tax, corporate, or patent lawyer* are typical of him. Moreover, the description of Bill does not refer to lawyers and does not bring any kind of lawyer to mind spontaneously. Therefore, the description *tax, corporate, patent, or other lawyer* should make available possibilities that have relatively strong support but that participants would not have otherwise considered. The unpacked description should increase probability judgments relative to the packed one; we expected to find subadditivity.

To test our predictions, we gave 46 University of Chicago master of business administration students the Linda problem (half in each condition) and 29 Duke University undergraduates the Bill problem (14 and 15 in the unpacked and packed conditions, respectively). Some of these students were randomly selected to receive a $20 prize in exchange for participation. The Duke students participated in exchange for a charitable donation. Participants completed a short questionnaire with either the packed or unpacked description of lawyer.

The results are displayed in the first two columns of Table 7 and are as we predicted. In the Linda problem, the packed description of lawyer received a median rank of 2, whereas the unpacked one received a median rank of 5 ($p < .01$ by Mann–Whitney). In the Bill problem, the probability rankings reversed revealing implicit subadditivity: The median rank of "Bill is a tax, corporate, patent, or other lawyer" was 2, whereas the median rank of "Bill is a lawyer" was 4 ($p < .01$).

These results suggest that one must be aware of the process by which people recruit support in order to predict whether subadditivity or superadditivity will arise. To emphasize this point, we asked a second group drawn from the same sample of Duke undergraduates to judge the probability that "a randomly chosen American" holds one of the professions in question. We surmised that lacking a specific target against which to evaluate each profession, participants would assess relative frequency by the relative accessibility of instances (i.e., the availability heuristic). Because tax, corporate, and patent attorneys are fairly common but not representative, we again predicted implicit subadditivity. Indeed, participants' rankings indicate higher probability judgments for the unpacked version of lawyer (*Mdn* rank $= 1$, $n = 15$) than for the packed version (*Mdn* rank $= 2.5$, $n = 14$). This result approached statistical significance ($p = .16$ by Mann–Whitney, see Random column of Table 7). Unpacking lawyer into a list of common but easily overlooked types of lawyers increased the availability of the explicitly mentioned types and thus increased support.

Table 7
Median Ranks of Lawyer Among 7 Professions for Unpacked and Packed Conditions in Experiment 4

Profession	Condition		
	Linda–Atypical[a]	Bill–Typical[b]	Random[b]
Lawyer	2.0	4.0	2.5
Tax, corporate, patent, or other type of lawyer	5.5	2.0	1.0

[a] These data were derived from the University of Chicago MBA student sample.
[b] These data were derived from the Duke University undergraduate sample.

General Discussion

The present experiments investigated the hypothesis that people interpret category descriptions narrowly, in terms of typical instances. As predicted by the narrow interpretation hypothesis, we found that atypical unpacking generally yielded implicit superadditivity, a result not expected by support theory (Rottenstreich & Tversky, 1997; Tversky & Koehler, 1994). We also found that unpacking with typical instances generally yielded additivity. Experiment 2 showed that these results are not due to the pragmatics of the question asked but rather to how people think about categorical information. Experiment 3 showed that when categories were more naturalistic and fuzzy, the effect of superadditivity was enhanced. Experiment 4 confirmed another prediction associated with narrow interpretation: Considerations of typicality versus availability allowed us to accurately predict cases of implicit super- versus subadditivity.

Overall, the present experiments support the narrow interpretation conjecture and thus elaborate the reminding-based explanation proffered by Tversky and Koehler (1994). The narrow interpretation hypothesis explains our results regarding unpacking as follows:

1. Typical unpackings yielded additivity because participants made probability judgments on the basis of typical instances and neglected atypical ones. Thus, the objects of judgment were very similar when categories were packed or unpacked with typical instances. In contrast, atypical unpackings inhibited spontaneous unpacking into typical exemplars, thereby lowering support and engendering superadditivity.

2. Implicit subadditivity required specific conditions: unpacking into exemplars that would not come to mind spontaneously but were of higher support than the typical exemplars that do come to mind. In general, available exemplars are also the exemplars with highest support; only when this correlation is broken can implicit subadditivity obtain.

Point 2 hints at an interpretation of implicit subadditivity for those studies that have found it (Fox & Birke, 2002; Fox & See, 2003; Koehler, Brenner, & Tversky, 1997; Rottenstreich & Tversky, 1997). Consider Rottenstreich and Tversky's demonstration of implicit subadditivity comparing the packed category *not guilty verdict* with the unpacked category *hung jury or a not guilty verdict* (our own subadditivity effect was small and not significant for this example). "Hung jury" is unlikely to come to mind spontaneously in the context of a "not guilty verdict." The fact that "hung jury" is also unlikely to have high support may explain why it failed to show subadditivity in Experiment 3.

Consistent with the part-list cuing effect in studies of memory, we have suggested that when instances are unpacked, other relevant instances tend to be neglected. But they may not be completely ignored. Probability judgments may be mediated by an anchoring and adjustment process such that people anchor on unpacked instances and then adjust insufficiently for other instances. When typical instances are unpacked, people will usually adjust down in the direction of atypical instances. They may do the same in the packed condition. When atypical instances are unpacked, there may be some adjustment upwards.

Our demonstrations of implicit superadditivity have been limited to unpacked categories that include a packed residual (*some other member of the category*). Partitioning a category into atypical instances plus a packed residual could also potentially yield explicit superadditivity. The current absence of reports of explicit superadditivity (though see Footnote 1) suggests either that the right experiment has not been done or that explicit subadditivity arises in part because of processes that operate independently of the choice of unpacked description, perhaps having to do with use of the probability response scale.

Implications for Support Theory

We do not question support theory's key contribution: that probability judgments are mediated by judgments of evidence or support. Moreover, the data collected to date are largely consistent with the assumptions of generic and explicit subadditivity (Equations 2 and 2A). However, the present data do contradict support theory's assumption that unpacking does not decrease subjective probability judgments, at least when sample spaces are large. Thus, we believe that this assumption should be jettisoned from future formulations of the theory. Various modifications would allow the theory to capture superadditive probability judgments. One possibility would be to introduce a support function that represents probability judgments as an average, rather than a sum (cf. Koehler et al., 2003), of the support associated with cognitively active subsets of the target category. Another way to implement the narrow interpretation hypothesis would be to represent the support for a packed category as the support for the category's most typical instance or instances. This will often equal the maximum of the supports for category subsets.

The value of viewing support as an intermediate component of judgment is demonstrated in a study by Rottenstreich, Brenner, and Sood (1999). Participants read the description of Tversky and Kahneman's (1983) character Linda and were told that Linda worked either as a journalist, a realtor, or an insurance salesperson. In a control group, the mean judged probabilities for the three occupations were greater for "journalist" than for "realtor" and for "insurance salesperson." A joint judgment group judged the probability that Linda was "either a journalist or realtor" and that Linda was an "insurance salesperson." The mean judged probability of the pair "journalist or realtor" was marginally above the probability assigned to "journalist" in the control group and far less than the total probability assigned in the control group to "journalist" and "realtor." One interpretation of this result is that people compared support values rather than probability judgments. Rottenstreich et al. show that the judgments imply higher support for the component ("journalist") than for the union ("journalist or realtor"). This accords with the notion of judgment by representativeness; most people agree that the pair "journalist or realtor" resembles Linda less than the hypothesis "journalist" alone. One consequent prediction is that judgments not made via representativeness (e.g., "what percentage of Americans work as either journalists or realtors?") can reveal upwards adjustment of support, a result obtained by Rottenstreich et al. Such predictions are transparent when support is invoked as an intermediate component of judgment but are difficult to see without it.

Conclusion

Implicit superadditivity may be especially important in situations in which several salient possibilities are long shots. For an individual who is considering an investment in an index of Internet stocks, the index may seem relatively attractive as a whole but relatively unattractive when unpacked into a list of individual components, many of which are unlikely to have a rosy future. In such situations, whether categories are unpacked will be consequential for judgments under uncertainty. But just as important is which subcategories are unpacked. The support associated with the specific selection can shape the assessment of the entire category.

References

Barsalou, L. W. (1985). Ideals, central tendency, and frequency of instantiation as determinants of graded structure in categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 11,* 629–654.

Brenner, L., & Koehler, D. (1999). Subjective probability of disjunctive hypotheses: Local-weight models for decomposition of evidential support. *Cognitive Psychology, 38,* 16–47.

Brenner, L., & Rottenstreich, Y. (1999). Focus, repacking, and the judgment of disjunctive hypotheses. *Journal of Behavioral Decision Making, 12,* 141–148.

Fox, C. R. (1999). Strength of evidence, judged probability, and choice under uncertainty. *Cognitive Psychology, 38,* 167–189.

Fox, C. R., & Birke, R. (2002). Lawyers exhibit subadditivity when forecasting trial outcomes. *Law and Human Behavior, 26,* 159–173.

Fox, C. R., & Clemen, R. T. (2003). *Partition dependence in subjective probability assessment.* Unpublished manuscript, Anderson School of Management, University of California, Los Angeles.

Fox, C. R., Rogers, B. A., & Tversky, A. (1996). Options traders exhibit subadditive decision weights. *Journal of Risk and Uncertainty, 13,* 5–17.

Fox, C. R., & See, K. E. (2003). Belief and preference in decision under uncertainty. In D. Hardmand & L. Macchi (Eds.), *Thinking: Psychological perspectives on reasoning, judgment and decision making* (pp. 273–314). New York: Wiley.

Fox, C. R., & Tversky, A. (1998). A belief-based account of decision under uncertainty. *Management Science, 44,* 879–895.

Hadjichristidis, C., Sloman, S. A., & Wisniewski, E. (2001). Judging the probability of representative and unrepresentative unpackings. In Johanna D. Moore (Ed.), *Proceedings of the 23rd Annual Conference of the Cognitive Science Society* (pp. 376–380). Mahwah, NJ: Erlbaum.

Hadjichristidis, C., Stibel, J., Sloman, S. A., Over, D. E., & Stevenson, R. J.

(1999). Opening Pandora's box: Selective unpacking and superadditivity. *Proceedings of the European Society for the Study of Cognitive Systems 16th Annual Workshop.* Siena, Italy: European Conference on Cognitive Science.

Hampton, J. A. (1998). Similarity-based categorization and fuzziness of natural categories. In S. A. Sloman & L. J. Rips (Eds.), *Similarity and Symbols in Human Thinking,* Cambridge, MA: MIT Press, 87–302.

Johnson, E. J., Hershey, J., Meszaros, J., & Kunreuther, H. (1993). Framing, probability distortions, and insurance decisions. *Journal of Risk and Uncertainty, 7,* 35–51.

Koehler, D. J., Brenner, L. A., & Tversky, A. (1997). The enhancement effect in probability judgment. *Journal of Behavioral Decision Making, 10,* 293–313.

Koehler, D. J., White, C. M., & Grondin, R. (2003). An evidential support accumulation model of subjective probability. *Cognitive Psychology, 46,* 152–197.

Macchi, L., Osherson, D., & Krantz, D. H. (1999). A note on superadditive probability judgment. *Psychological Review, 106,* 210–214.

Murphy, G. L. (2003). *The big book of concepts.* Cambridge, MA: MIT Press.

Nickerson, R. (1984). Retrieval inhibition from part-set cuing: A persistent enigma in memory research. *Memory & Cognition, 12,* 531–552.

Redelmeier, D., Koehler, D., Liberman, V., & Tversky, A. (1995). Probability judgment in medicine: Discounting unspecified possibilities. *Medical Decision Making, 15,* 227–230.

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

Rottenstreich, Y., Brenner, L., & Sood, S. (1999). Similarity between hypotheses and evidence. *Cognitive Psychology, 38,* 110–128.

Rottenstreich, Y., & Tversky, A. (1997). Unpacking, repacking, and anchoring: Advances in support theory. *Psychological Review, 104,* 406–415.

Slamecka, N. (1968). An examination of trace storage in free recall. *Journal of Experimental Psychology, 76,* 504–513.

Sloman, S. A. (1991). Part-set cuing inhibition in category-instance and reason generation. *Bulletin of the Psychonomic Society, 29,* 136–138.

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

Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, *90,* 293–315.

Tversky, A., & Koehler, D. J. (1994). Support theory: A nonextensional representation of subjective probability. *Psychological Review, 101,* 547–567.

Received May 23, 2003
Revision received October 22, 2003
Accepted October 27, 2003 ■