

Judging the Probability of Representative and Unrepresentative Unpackings

Constantinos Hadjichristidis (constantinos.hadjichristidis@durham.ac.uk)

Department of Psychology, University of Durham
Durham, DH1 3LE, UK

Steven A. Sloman (steven_sloman@brown.edu)

Department of Cognitive & Linguistic Sciences, Brown University
Box 1778, Providence, RI 02912, USA

Edward J. Wisniewski (edw@uncg.edu)

Department of Psychology, University of North Carolina at Greensboro
Box 26164, Greensboro, NC 27402-6164, USA

Abstract

The hypothesis that category descriptions are interpreted narrowly, in terms of representative instances, is examined by comparing probability judgments for packed descriptions of events to judgments for coextensional unpacked descriptions. The representativeness of the unpacked instance was varied along with the type of unpacking (direct vs. priming). In contrast to the prediction of Support Theory (Tversky & Koehler, 1994), we found no evidence that unpacking has a nonnegative effect on probability judgments (subadditivity). Instead, we found a negative effect with low representative direct unpackings (superadditivity). Our data suggest that probability judgments are proportional to the typicality of instances in the description.

Introduction

People frequently assess the probability of uncertain events such as the chance of rain or the success rate of a medical treatment. Such probability assessments are important because they determine not only whether to plan a barbecue, but also whether or not to have surgery. Judgments concerning rain or effectiveness of a treatment are categorical in the sense that the event being judged could be instantiated in many ways. We consider whether the typicality of the instances used in a categorical description affects the judged probability of corresponding events.

Normative models of probability judgment assume *description invariance*: the probability of an event does not depend on how the event is described. This assumption is descriptively invalid (Tversky & Kahneman, 1986). People give lower probability ratings for the packed hypothesis "Death from homicide, rather than accidental death" than for the coextensional unpacked hypothesis "Death from homicide by an acquaintance or by a stranger, rather than accidental death" (Rottenstreich & Tversky, 1997).

Support for Support Theory?

To accommodate this fact, Tversky and Koehler (1994) proposed a descriptive theory of probability judgment, Support Theory, that suggests that subjective probabilities are assigned not to events, but to descriptions of events or *hypotheses*. Probability judgments are hypothesized to be mediated by evaluations of evidence for and against a hypothesis. Specifically, the judged probability of a hypothesis H rather than an alternative hypothesis A is given by:

$$(1) \text{ Judged } P(H, A) = s(H) / [s(H) + s(A)]$$

where $s(X)$ is a global measure of support for hypothesis X. It is a sum of the (weighted) support of all representative instances of X that are available to the judge at the time of evaluation.

A key assumption of Support Theory is that exhaustively unpacking a hypothesis H into mutually exclusive and exhaustive sub-hypotheses ($H_1 \vee \dots \vee H_n$) can increase the support for H:

$$(2) \quad s(H) \leq s(H_1 \vee \dots \vee H_n)$$

This assumption is motivated on the grounds that unpacking may bring additional instances to mind, or increase the salience of the unpacked instances. Either or both of these effects would increase the perceived support for a hypothesis.

Taken together, Support Theory's assumptions predict implicit subadditivity: The judged probability of an implicit (or packed) hypothesis H is no greater than the judged probability of a coextensional unpacked hypothesis. Implicit subadditivity has been observed several times (see Rottenstreich & Tversky, 1997).

However, not all the data are so supportive. Hadjichristidis et al. (1999) showed that selectively unpacking hypotheses into components that enjoy low levels of support results in the opposite phenomenon, implicit superadditivity. 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." In a series of follow-up studies we have consistently found implicit superadditivity with novel categories unpacked with atypical instances. We have consistently failed to find implicit subadditivity, even when events were unpacked using representative instances, instances that enjoy high levels of support. In sum, contrary to support theory's predictions, these data suggest that unpacking does not always increase subjective probability judgments.

The Supported Theory

A parsimonious interpretation of our data is based on Support Theory's own assumption that people interpret category-based hypotheses narrowly, in terms of representative instances. Unpacking unrepresentative instances induces superadditivity by making instances of very low support part of what is judged. Unpacking representative instances leaves probability judgments unaffected because packed categories are interpreted in terms of representative instances.

Unlike Support Theory, we suggest that unpacking can decrease support. According to the present proposal, people assess the likelihood of a category-based hypothesis by thinking about instances in which the event is expected to occur (i.e., by bringing to mind representative instances, instances enjoying high levels of support). This dovetails with Kahneman and Miller's (1986) proposal that norms—contrast events for judgments of surprise, blame, etc.—are constructed according to the availability and representativeness of exemplars. Our proposal is that the determinants of exemplar retrieval control not only how contrast events are conceived, but how focal ones are too. Moreover, the mere availability in memory of an instance is not sufficient to change judgments of likelihood, the instance must be one of the objects of judgment.

Study

The current study tests our hypothesis by crossing representativeness (high- vs. low-representative instances) with type of unpacking (direct vs. priming) in a between-participants design. A separate group of participants was asked to provide estimates for corresponding packed hypotheses. The dependent measure was subjective probability judgment. Table 1

gives one illustration from each of the five experimental conditions.

Table 1: An example stimulus from each of the five conditions. The sentence in bold-faced letters is a description that preceded evaluations in all conditions.

| | |
|-------------------------|---|
| | Sarah is a very energetic and happy eight year old who loves playing with her stuffed animals. |
| <u>Packed</u> | How likely is it that Sarah hates some types of pets (as opposed to loving all pets)? _____ |
| <u>Direct High Rep</u> | How likely is it that Sarah hates tarantulas or some other types of pets (as opposed to loving all pets)? _____ |
| <u>Direct Low Rep</u> | How likely is it that Sarah hates horses or some other types of pets (as opposed to loving all pets)? _____ |
| <u>Priming High Rep</u> | Same as packed but prior to making the judgment primed with a list of words including "tarantulas" |
| <u>Priming Low Rep</u> | Same as packed but prior to making the judgment primed with a list of words including "horses" |

Direct unpacking refers to a conventional unpacking manipulation. Participants in the direct unpacking conditions were asked to judge categories from which one of their instances had been unpacked. Based on previous findings, we expected to find a negative effect of unpacking unrepresentative instances (i.e., implicit superadditivity), and no effect of unpacking representative instances.

Participants in the priming unpacked conditions were asked to judge packed hypotheses after being primed with either representative or unrepresentative instances. Priming consisted of asking participants to study the instances for 1 min. for a later memory test. We reasoned that priming would make the critical instances highly available in memory at the time of judgment without specifically making them the objects of judgment. If merely making an instance available in memory increases the likelihood that it will be considered during judgment of a category that is superordinate to it, then superadditivity should be observed in the Priming Low Representativeness condition (i.e., probability judgments should be higher in the Packed condition than the Priming Low Representativeness condition) and additivity should be observed in the Priming High Representativeness Condition. However, if the availability of an instance is not sufficient, if a narrow interpretation of categories is so ingrained that making atypical instances available in memory does not influence how people conceive of the category being judged, then the priming unpacking conditions should produce additive judgments. That is, we should see no effect of the priming manipulation.

The availability hypothesis predicts a main effect of representativeness and no effect of type of unpacking.

The narrow interpretation hypothesis predicts a Representativeness by Unpacking interaction due to a negative effect of low representativeness in the direct unpacking condition and no other differences.

Method

162 first-year students participated in the experiment, 76 sampled at the University of Durham (UK), and 86 at the University of North Carolina at Greensboro (US). Participants were presented with booklets containing eight examples from one of the 5 experimental conditions, followed by 16 items asking for judgments of representativeness. The judgments of representativeness were obtained as a validation check on the assignment of examples in the high- and low-representativeness conditions.

Results

Representativeness judgments

Table 2 presents mean representativeness estimates for both populations for High Rep and Low Rep conditions. As expected, ratings for "High Representativeness" items were much higher than those for "Low Representativeness" items.

Table 2: Mean Population by Representativeness subjective representativeness estimates.

| | High Rep | Low Rep |
|-------------------|----------|---------|
| <u>Greensboro</u> | 65.7 | 34.3 |
| <u>Durham</u> | 61.4 | 30.0 |

To make sure our UK and US population samples were comparable, we conducted a 2 (Population) by 2 (Representativeness) repeated-measures ANOVA across items. The main effect of representativeness was highly significant ($F(1,14)=23.33, p<.001$). There was also a main effect of population, US probability judgments were about 4 percentage points higher than UK judgments ($F(1,14)=4.55, p<.06$). Most importantly, no interaction was observed ($F<1$). The results justify the assignment of items to High- or Low-representativeness conditions.

An examination of the judgments for each item showed that the direction of representativeness judgments for one were opposite to our expectations. This item was excluded from subsequent analyses.

Probability Judgments

Population To test whether population influenced probability judgments, we performed a 2 Population by 5 Experimental condition ANOVA across participants. Only Experimental condition reached significance

($F(4,151)=5.18, p<.001$). The data for the two populations were combined for subsequent analyses.

Unpacking by Representativeness Table 3 presents mean subjective probability judgments for each Unpacking (direct vs. priming) and Representativeness (high vs. low) condition. The mean of the direct low-representativeness cell is the lowest; means of the other cells are about equal.

Table 3: Mean probability judgments by Type of unpacking and Representativeness.

| | High Rep | Low Rep |
|----------------|----------|---------|
| <u>Direct</u> | 57.3 | 44.6 |
| <u>Priming</u> | 55.1 | 55.0 |
| <u>Packed</u> | 56.2 | |

The data were analyzed by two 2 Unpacking by 2 Representativeness analyses of variance, one by participants (F_1) and one by items (F_2). Unpacking had a significant main effect by participants but not by items ($F_1(1,124)=3.96, p<.05; F_2(1,6)=1.87, p>.22$). Representativeness had a significant main effect by participants ($F_1(1,124)=9.71, p<.005$) but only a marginal effect by items ($F_2(1,6)=3.03, p<.14$). The interaction was significant by participants ($F_1(1,124)=9.37, p<.005$) but only marginally by items ($F_2(1,6)=3.08, p<.14$).

Two-tailed t-tests compared each Unpacking by Representativeness condition to the rest. The only tests reaching significance were those comparing the direct low-representativeness condition to each of the others.

Superadditivity Mean probability ratings for each of the four Unpacking by Representativeness conditions were compared to the mean rating for the packed condition ($M=56.2$) to detect deviations from additivity. The only condition that deviated substantially from additivity was the direct low-representativeness condition that demonstrated superadditivity: $t(61) = 3.27, p <.005$ (participants); $t(6) = 2.51, p <.05$ (items). The item analysis for the direct high-representativeness condition suggested a small amount of subadditivity: $t(6) = 2.43, p <.06$; but $t < 1$ (participants).

Discussion

The present study investigated the hypothesis that people interpret category descriptions narrowly, in terms of representative instances, when making subjective probability judgments by crossing representativeness with type of unpacking. We found a representativeness by unpacking interaction due to a negative effect of low representativeness in the direct unpacked condition. Only the direct low-representativeness ratings substantially deviated from

additivity: they were superadditive. Predictions were confirmed with both British and US samples.

The present data replicated Hadjichristidis et al.'s (1999) finding that directly unpacking unrepresentative instances induces implicit superadditivity. One account of these findings is that unpacked instances are treated as a pragmatic cue for determining what the experimenter means by the category label. When asked about "tarantulas or some other type of pet," people might infer that the experimenter has a different category in mind than when asked only about "pets." This account is indeed consistent with our data, but only if construed in a way equivalent to our hypothesis. The data suggest that people interpret categories narrowly and the explicit inclusion of atypical instances broadens the normal interpretation. But our methodology rules out the interpretation that we are merely asking people to judge a different category in the Low Representativeness condition. In every case, categories were described in the current study by clearly stating the alternative hypothesis (e.g., in Table 1, the judged event is always stated along with "as opposed to loving all pets"). Therefore, although we believe our effect depends on how categories are interpreted, it does not represent a mere task demand induced by pragmatic biases. Rather, it represents a central and generalizable aspect of probability judgment of categorical events.

Support theory Our finding that direct unpacking of low representative instances induces superadditivity disconfirms support theory's prediction that unpacking cannot have a negative effect on probability judgments. Our conclusion is independently supported by Macchi, Osherson, and Krantz (1999) who found that unpacking low-support instances resulted in explicit superadditivity for binary partitions.

Implicit subadditivity is not a robust phenomenon. Rottenstreich and Tversky (1997) themselves predicted it three times, but only observed it twice. Implicit subadditivity obtained in the Trial problem, which pitted the hypothesis "the trial will not result in a guilty verdict" against the disjunction "the trial will result in a not guilty verdict or a hung jury". It also obtained in the Homicide problem, which pitted "death from homicide" against "death from homicide by an acquaintance or a stranger". In both cases, they explained subadditivity in terms of enhanced availability. In the first problem, participants might not have considered the hung jury possibility in the packed condition. In the second, the unpacked hypotheses may have brought a host of possible causes of death to mind (e.g., crimes of passion) that would not have been available in the packed condition. In sum, support theorists have themselves identified a key factor that limits the generality of implicit subadditivity.

We believe that minor modifications would allow Support Theory to capture superadditive probability judgments. Here are some possible changes:

1. Allow for negative support.
2. Stick to nonnegative support, but modify the global support function (make it average rather than summed support).
3. Allow that unpacked instances replace instances that would otherwise have been available at the time of judgment.

Dynamic global support functions? A further possibility is that implicit subadditivity and superadditivity reflect different functional relations between the support attached to packed categories and the support attached to their unpacked instances (global support functions). Subadditivity may involve summing of support across instances, whereas superadditivity may involve averaging of support. In Rottenstreich and Tversky's (1997) examples, support is based on subjective impressions of frequencies or reasons, whereas in our "natural fuzzy category" examples (e.g. pets, restaurants), support is based on similarity. The support for the unpacked hypothesis of the Trial problem, for instance, seems to involve estimating the relative frequency of a "not guilty verdict" and a "hung jury" and adding them up. In contrast, the support for the unpacked hypotheses in our examples seems to involve estimating the similarity of the category instances to the description and averaging them out. Corroborating evidence that the similarity-based global support function may average support comes from Rottenstreich, Brenner, and Sood (1999) who showed that similarity-based likelihood judgment gives rise to nonmonotonicities: the support of a disjunction is less than that of one of its components. They presented participants with a description of Linda: an outspoken, socially conscious, and single woman. "Linda is a journalist" enjoyed higher support than the disjunctive hypothesis "Linda is a journalist or a realtor". Nonmonotonicities cannot be explained by a global support function that adds support, but could by one that averages support.

In sum, the global support function may change dynamically depending on the particular base of support –e.g. objective frequencies, similarity, and reasons. Similarity-based likelihood judgment may involve averaging; frequency-based likelihood judgment may involve summing. Supporting this hypothesis, Rottenstreich et al. (1999) showed that case judgments (e.g., the Linda example) that presumably involve similarity-based reasoning give rise to nonmonotonicities, whereas class judgments (e.g., the probability that a randomly selected American is a

journalist) that presumably promote frequency-based reasoning, do not. These suggestions all stay close to the spirit of Support Theory because we find its assumption that probability judgments are mediated by judgments of evidence to be appealing and worth maintaining in the next generation of descriptive theory (see also Fischhoff, Slovic, & Lichtenstein, 1979).

Decision-making Much everyday decision-making depends on subjective assessments of probability. For instance, the premium one is willing to pay for health insurance depends on a subjective assessment of the likelihood of getting hospitalized for the cases that the health insurance covers (see Johnson et al., 1993). The events for which an insurance provides coverage can be described in many ways. The results we report here can presumably be extended to the domain of decision-making.

Acknowledgments

We want to thank David Over and Rosemary Stevenson for helpful discussions on Support Theory and Marianne Harrison for running pilot studies on this project. Constantinos Hadjichristidis was supported by the ESRC grant No. R000239074 on *Belief Revision and Uncertain Reasoning*.

References

- Fischhoff, B., Slovic, P., & Lichtenstein, S. (1978). Fault trees: Sensitivity of estimated failure probabilities to problem representation. *Journal of Experimental Psychology: Human Perception and Performance*, 4, 330-334.
- Hadjichristidis, C., Stibel, J. M., Sloman, S. A., Over, D. E., & Stevenson, R. J. (1999). Opening Pandora's box: selective unpacking and superadditivity. *Proceedings of the European Conference on Cognitive Science* (pp. 185-190). Siena, Italy.
- Johnson, E. J., Hershey, J., Meszaros, J., & Kunreuther, H. (1993). Framing, probability distortions, and insurance decisions. *Journal of Risk and Uncertainty*, 7, 35-51.
- Kahneman, D. & Miller, D. T. (1986). Norm Theory: comparing reality to its alternatives. *Psychological Review*, 93, 136-153.
- Macchi, L., Osherson, D., & Krantz, D. H. (1999). A note on superadditive probability judgment. *Psychological Review*, 106, 210-214.
- 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.
- Tversky, A. & Kahneman, D. (1986). Rational choice and the framing of decisions. *Journal of Business*, 59, 251-278.
- Tversky, A. & Koehler, D. J. (1994). Support theory: A nonextensional representation of subjective probability. *Psychological Review*, 101, 547-567.