Predictions from Uncertain Beliefs

Samuel G. B. Johnson[1], Thomas Merchant[2], & Frank C. Keil[1]
(samuel.johnson@yale.edu, thomas_merchant@brown.edu, frank.keil@yale.edu)
[1]Dept. of Psychology, Yale University, 2 Hillhouse Ave., New Haven, CT 06520 USA
[2]Dept. of Cognitive, Linguistic, and Psychological Sciences, Brown University, 190 Thayer St., Providence, RI 02912 USA

Abstract

According to probabilistic theories of higher cognition, beliefs come in degrees. Here, we test this idea by studying how people make predictions from uncertain beliefs. According to the *degrees-of-belief* theory, people should take account of both high- and low-probability beliefs when making predictions that depend on which of those beliefs are true. In contrast, according to the *all-or-none* theory, people only take account of the most likely belief, ignoring other potential beliefs. Experiments 1 and 2 tested these theories in explanatory reasoning, and found that people ignore all but the best explanation when making subsequent inferences. Experiment 3A extended these results to beliefs fixed only by prior probabilities, while Experiment 3B found that people can perform the probability calculations when the needed probabilities are explicitly given. Thus, people's intuitive belief system appears to represent beliefs in a 'digital' (true or false) manner, rather than taking uncertainty into account.

Keywords: Explanation; abduction; causal reasoning; belief; prediction; diagnosis; probability; uncertainty.

Introduction

Our beliefs often entail other beliefs. Knowing an object's category helps us to make predictions about that object (Anderson, 1991; Murphy, 2002). If a furry object is a rabbit, it might hop; if it's a skunk, it might smell. Likewise, causal beliefs facilitate predictions (Waldmann & Holyoak, 1992). If the house is smoky because Mark burned the cookies, then we have an unpleasant dessert to look forward to; if it's smoky because Mark dropped a cigarette in the bed, then we may have bigger problems.

However, beliefs are often accompanied by uncertainty. If we see a furry object from a distance, we may be only 70% confident that it is a rabbit rather than a skunk; if we are awoken from a nap by smoke, we may think there is a 20% chance that the house is burning down. In such cases of uncertain beliefs, accurate inference about those beliefs' consequences requires these possibilities to be weighted and combined. This can be done using the tools of probability theory. Here, we test whether people use probabilities to represent beliefs as coming in degrees, or whether people might instead use shortcuts, representing beliefs as though they are either true or false.

Imagine there is a 70% chance that the furry object is a rabbit (possibility *A*), and a 30% chance that it is a skunk (*B*). What is the probability that it will hop (*Z*)? Suppose 80% of rabbits hop, while only 2% of skunks hop. That is:

$$P(A) = .70, P(B) = .30, P(Z|A) = .80, P(Z|B) = .02.$$

Then the probability of hopping (*Z*) can be calculated as:

$$P(Z) = P(Z|A)P(A) + P(Z|B)P(B) = (.8)(.7) + (.02)(.3) = .6$$

Anderson (1991) argued that people follow this principle in category-based prediction. That is, when estimating the likelihood that an object has a feature, people consider the various possible categorizations of that object, and then weight the conditional probability of the feature given those categories by the probability of each category.

But people usually do not consider all possible categorizations of an object, but focus on the single most likely category (Murphy & Ross, 1994). In our example, people would ignore the possibility that the object is a skunk, and 'round up' the rabbit probability to 100%:

$$P(Z) = P(Z|A)P(A) + P(Z|B)P(B) = (.8)(1) + (.02)(0) = .8$$

That is, people only consider the conditional probability of a new feature given the most likely category, as though they believe that the object *must* belong to that category.

This result has been found consistently across many studies. For example, Murphy and Ross (1994) presented participants with exemplars belonging to categories of drawings by different children, which varied in color and shape. Participants were then told about a new exemplar (e.g., a triangle), and asked to categorize it. Because the training exemplars included 5 triangles, of which 3 were drawn by the same child (Bob), virtually all participants responded that the new triangle was likely drawn by Bob (with about 60% confidence). Participants then predicted the color of the new exemplar. Participants based these predictions only on the distribution of colors within the most likely category (Bob), as though the 60% chance of the exemplar belonging to that category had been 'rounded up' to 100%. That is, people relied only on the single best categorization, ignoring the 40% chance that the exemplar belonged to a different category.

These findings may be unique to categorization. Categories are discrete representations (Dietrich & Markman, 2002)—an object is a rabbit or a skunk, not both. This basic underlying logic of categorization may account for people's reluctance to entertain multiple possible categorizations, in which case we would not expect similar results in other cognitive domains.

However, beliefs might be represented in an all-or-none ('digital') manner not only in categorization, but across cognition. Such a result would be surprising from the standpoint of probabilistic theories of cognition (e.g., Oaksford & Chater, 2009). On a common philosophical interpretation of probability, the purpose of probabilities is to reflect 'degrees of belief' (Jeffrey, 1965)—indeed, some philosophical theories hold that only logical tautologies should be assigned a probability of 1, and only logical contradictions a probability of 0 (Kemeny, 1955). If people do not represent beliefs in degrees (as 'graded'),

Figure 1: Causal structure used in all experiments. White indicates a variable that was observed, and grey indicates a variable that is unknown.

this poses serious difficulties for claims that people perform Bayesian updating using normative principles.

In the current experiments, we tested whether people make predictions from uncertain beliefs in an all-or-none or a graded manner. We followed the logic of the Murphy and Ross (1994) experiments, but rather than categorizing objects, participants either inferred causal explanations from observations (Experiments 1 and 2) or inferred the most likely possibility given base rates (Experiment 3).

Participants learned about causal systems with the structure depicted in Figure 1. That is, two explanations (A and B) could account for some data (X), and these explanations had different implications for some novel prediction (Z). We structured the problems so that explanation A would be seen as more probable than explanation B. We tested whether people rely only on A or instead consider both A and B, by varying the conditional probability of Z given each explanation [i.e., $P(Z|A)$ and $P(Z|B)$]. If people integrate across all possible explanations, both manipulations should have an effect on judgments of $P(Z)$. In contrast, if people rely only on the most likely explanation, then manipulating $P(Z|B)$ should have no effect at all on $P(Z)$.

Experiment 1

In Experiment 1, we relied on people's known preference for simple explanations (Lombrozo, 2007). Other things being equal, people prefer to explain data with one cause rather than two causes. Thus, we expected that when a simple and complex explanation can both account for a set of observations, people will make subsequent inferences as though only the simple explanation were possible. For example, participants learned about a simple ecosystem in a lake (letters in brackets not in original):

Juga snails [A] cause lakes to lose sculpin fish [X] and lose crayfish [Y].
Scuta snails [B] cause lakes to lose sculpin fish [X].
Aspera snails [C] cause lakes to lose crayfish [Y].

Thus, if a lake had lost both sculpin fish and crayfish

(effects X and Y), the juga snails explanation (A) would be more compelling than the conjunctive scuta plus aspera snails explanation (B and C combined), even though either explanation could account for the data. Thus, we would expect participants to infer the simple explanation (A), given that they are told that X and Y are observed.

To see whether people would make subsequent inferences that ignored the possibility that the complex explanation was true, participants learned about another effect, bacteria proliferation (Z), which occurs with different probabilities, depending on the cause. In the *low/low* condition, this effect had a low probability regardless of the cause (underlining not in original):

When a lake has juga snails [A], it occasionally has bacteria proliferation [Z].
When a lake has both scuta snails [B] and aspera snails [C], it occasionally has bacteria proliferation [Z].

The *high/low* condition was like the *low/low* condition, except that $P(Z|A)$ remained high while $P(Z|B,C)$ was low:

When a lake has juga snails [A], it usually has bacteria proliferation [Z].
When a lake has both scuta snails [B] and aspera snails [C], it occasionally has bacteria proliferation [Z].

Since participants would infer that A is the best explanation, we would expect a difference between the *low/low* and *high/low* conditions in judgments of $P(Z)$, reflecting the higher value of $P(Z|A)$. Finally, the *low/high* condition was the reverse of the *high/low* condition, with a low value of $P(Z|A)$ and a high value of $P(Z|B,C)$:

When a lake has juga snails [A], it occasionally has bacteria proliferation [Z].
When a lake has both scuta snails [B] and aspera snails [C], it usually has bacteria proliferation [Z].

If participants ignore the possibility that the complex explanation is true (effectively placing 100% of their confidence in the simple explanation), then we would expect no difference between the *low/low* and the *low/high* conditions in ratings of the likelihood of Z, since the conditional probability given the conjunctive explanation would be irrelevant. Conversely, if they weight all possible explanations in a normative manner (Anderson, 1991), then they should differentiate between the *low/low* and *low/high* conditions.

Method

We recruited 120 participants from Amazon Mechanical Turk for Experiment 1; 8 were excluded from analysis because they incorrectly answered more than one-third of a set of true/false check questions.

Each participant completed three items—one each in the *low/low*, *high/low*, and *low/high* conditions. For the snail item, participants first read about the effects of A, B, and C on X and Y, using the above wording. They then read about the effects of these causes on Z, with either the above *low/low*, *high/low*, or *low/high* wording. Next, participants indicated their favored explanation:

Crescent Lake has a loss of sculpin fish [*X*] and crayfish [*Y*]. Which do you think is the most satisfying explanation for this?

Participants answered this question as a forced-choice between "Crescent Lake has juga snails" [*A*] and "Crescent Lake has scuta snails and aspera snails" [*B* and *C*]. Finally, participants were asked to rate the probability of *Z* ("What do you think is the probability that Crescent Lake has bacteria proliferation") on a scale from 0 to 100.

Three vignettes were used (snails, bacteria, and fungus), and condition (*low/low*, *high/low*, or *low/high*) was balanced with vignette using a Latin square. Items were completed in a random order, and all questions for each item appeared on the same page.

Results and Discussion

Most participants (78 out of 112) preferred the simpler explanation for all three items. Because our hypotheses are predicated on the assumption that participants inferred the simple explanation, we focus on these participants' responses in analyzing the results of all experiments. However, the results of all experiments are similar if all participants are included who passed the check questions.

Figure 2 shows the mean estimates of P(*Z*), across the three conditions. When both the simple explanation and the complex explanation corresponded to a low probability of *Z* (it "occasionally" leads to *Z*) in the *low/low* condition, mean judgments were 50.74 (*SD* = 23.63). But when the simple explanation instead corresponded to a high probability of *Z* (it "usually" leads to *Z*) in the *high/low* condition, mean judgments were much higher [*M* = 71.69, *SD* = 18.27; $t(77) = 7.27$, $p < .001$, $d = 0.82$, $BF_{10} > 1000$][1]. Thus, manipulating the P(*Z*|*A*) had a dramatic effect on judgments of P(*Z*). This result is consistent with either graded or digital beliefs, since *A* was the single best explanation for the data.

Much more surprisingly, however, manipulating P(*Z*|*B*,*C*) had *no effect* on the perceived probability of *Z*: There was no difference between the *low/low* condition and the *low/high* condition [*M* = 48.53, *SD* = 21.06; $t(77) = -0.80$, $p = .43$, $d = -0.09$, $BF_{01} = 8.18$]. That is, those participants who (reasonably) believed that the simple explanation was more likely than the complex explanation reasoned as though the simple explanation were *certain* and the complex explanation were *impossible*: Participants ignored the possibility that the complex

---

[1] Because null effects were predicted for some comparisons, all *t*-tests in this paper are accompanied by Bayes Factor (*BF*) analyses (Rouder, Speckman, Sun, Morey, & Iverson, 2009), with a scale factor of 1. *BF*s can quantify evidence either against or in favor of a null hypothesis. When the evidence favors the alternative hypothesis, we denote this 'BF_{10}', and when the evidence favors the null hypothesis, we denote this 'BF_{01}'. For example, "$BF_{10} = 7.0$" means that the data would be 7 times likelier under the alternative than under the null, while "$BF_{01} = 4.0$" means that the data would be 4 times likelier under the null than under the alternative.

Figure 2: Results of Experiments 1 and 2.

explanation was true when estimating P(*Z*).

These results suggest that, just as in category-based prediction (Murphy & Ross, 1994), people base predictions from uncertain explanations off of only their preferred explanation, ignoring the possibility that other explanations could be correct. This is a flagrant violation of probability theory, as all possible explanations must be weighted in making subsequent inferences (Anderson, 1991). Indeed, such behavior seems to defeat the very *point* of probabilistic inference, which is to allow for degrees of belief rather than all-or-none acceptance of propositions (Jeffrey, 1965).

However, two aspects of this experiment might be cause for concern. First, we obtained participants' explanatory ratings as a forced-choice, perhaps creating some experimenter demand to focus on the explanation the participant selected. Although Murphy and Ross (1994) found similar results regardless of whether participants were asked to categorize the exemplar, this is nonetheless a reasonable concern about this experiment.

Second, participants may have thought that the simple explanation was so much more probable than the complex explanation that they were *right* to ignore the complex explanation in estimating the probability of *Z*. That is, suppose participants thought there were a 99% chance of *A*, and a 1% chance of *B* and *C* (this is not so unreasonable, since Lombrozo, 2007 found a very strong simplicity bias, exceeding what is normatively appropriate). In that case, the contribution of P(*Z*|*A*) should be 99 times greater than that of P(*Z*|*B*,*C*), and our experimental set-up may not be sufficiently sensitive to detect such a small effect of P(*Z*|*B*,*C*).

Experiment 2

In Experiment 2, we avoided these concerns by asking participants to estimate the probability of *A* [P(*A*|data)] and of *B* and *C* [P(*B*,*C*|data)] rather than making a forced choice between the two explanations. First, this avoided experimenter demand to focus only on one explanation, and, if anything, would seem to encourage participants to weight both explanations. Second, this measurement allowed us to determine how much larger the effect of P(*Z*|*A*) should be, relative to the size of P(*Z*|*B*,*C*), and to compare performance to this normative benchmark.

Method

We recruited 120 participants from Amazon Mechanical Turk for Experiment 2; 6 were excluded because they incorrectly answered more than one-third of the check questions, and 12 because their total probability ratings for at least one item were not between 80% and 120%.

The procedure for Experiment 2 was the same as Experiment 1, with two changes. First, rather than asking which explanation participants favored as a forced-choice, they were asked to rate the probability of each explanation given the evidence [i.e., P(A) and P(B,C)], and were instructed to ensure the probabilities added up to 100%. Second, the question about the probabilities of the explanations was asked on one page, then the question about the probability of Z was asked on a separate page. This change was made to avoid demand for consistency across the two sets of questions. The probability information was repeated at the top of both pages.

Results and Discussion

Most participants (72 out of 102) rated P(A) at least as high as P(B,C) for all items. Among those participants, the mean estimate of P(A) was 65.88 (SD = 16.33) and the mean estimate of P(B,C) was 34.06 (SD = 16.30). Thus, despite their belief that the simpler explanation was more probable, participants allocated substantial probability to the complex explanation.

Nonetheless, the results of Experiment 2 were similar to those of Experiment 1 (Figure 1). Participants gave higher estimates of P(Z) in the *high/low* than in the *low/low* condition [M = 70.60, SD = 18.62 vs. M = 60.24, SD = 22.66; t(71) = 3.41, p = .001, d = 0.40, BF$_{10}$ = 18.85], though this effect is about half as large, compared to Experiment 1. This difference appears to be due to task demands, although it is not clear whether it is a demand in Experiment 1 to focus more on P(Z|A), or a demand in Experiment 2 to focus less on P(Z|A), relative to a condition in which participants did not make any explicit judgments about the explanations. In any case, however, this result is robust across both tasks, and the true effect size likely lies somewhere in the middle.

Most critically, there is once again no difference between the *low/low* condition and the *low/high* condition [M = 59.31, SD = 23.23; t(71) = -0.33, p = .74, d = -0.04, BF$_{01}$ = 10.22]. Thus, once again, while participants were happy to use P(Z|A) in estimating the probability of Z, they completely ignored P(Z|B,C). This occurred even though participants indicated that there was about a one-third chance that explanation A was true.

Could these results be driven by the assumption that the causes are not mutually exclusive? That is, perhaps participants are assuming that A could have occurred along with combinations of B and C, in which case the evidence is a much better signal for A than for B and C. However, this explanation is untenable in light of participants' explicit ratings of the explanations, which indicated considerable credence in the B,C explanation.

To further rule out this possibility, we can also compare each participant's estimate of P(Z) to a normative standard, calculated from that participant's own probability ratings of P(A), P(B,C), and P(Z). We included all 102 participants who passed the check questions in this analysis. Normatively, P(Z) can be calculated as:

$$P(Z) = P(Z|A)P(A) + P(Z|B,C)P(B,C)$$

Since we used verbal labels ("occasionally" and "usually") rather than precise probabilities, we must use an indirect method to calculate predicted values of P(Z). From the *high/low* and *low/low* conditions, we estimated each participant's implicit probability difference between "usually" and "occasionally" (M = 9.33, SD = 38.04). This allowed us to calculate how large the difference between the *low/high* and *low/low* conditions should be. Participants' difference scores between the *low/low* and *low/high* conditions were substantially smaller than these normative values, derived from their other ratings [M = 5.80, SD = 23.85 for the difference between actual and normative judgments; t(101) = 2.46, p = .016, d = 0.24, BF$_{10}$ = 1.43]. This analysis of individual participants thus corroborates the overall pattern of means, indicating that participants underweighted (in fact, did not weight at all) the complex explanation in estimating P(Z).

Experiments 3A and 3B

In our final experiment, we aimed to test whether people would underweight the probability of any unlikely belief in making subsequent inferences, or whether this effect was confined to explanatory inferences (such as causal and category-based reasoning). Thus, instead of manipulating the probability of two competing beliefs (A and B) by varying their plausibility as explanations (e.g., by making A a simple explanation and B a complex explanation), we instead manipulated the base rates of A and B, by asserting that A had a 65% chance of being true while B had a 35% chance (the same base rates for A and B that people gave for the simple and complex explanations in Experiment 2). If participants' beliefs are 'digital' only when they must infer a category or cause, then they would rely on both P(Z|A) and P(Z|B) when making subsequent inferences about Z. But if *any* two incompatible beliefs are resolved in a digital fashion (so that either A or B is believed all-or-none), then participants would continue to ignore P(Z|B) in estimating P(Z). We tested this in Experiment 3A.

A second goal was to ensure that participants were not making normative errors simply because they are incapable of performing the mathematics. Thus, Experiment 3B gave participants all four numbers needed to calculate P(Z) [i.e., P(Z|A), P(Z|B,C), P(A), and P(B,C)]. If participants make more normative inferences here, it would suggest that participants know that the likelihood of Z given low-probability beliefs is relevant, but do not use it spontaneously when forced to perform a task without the benefit of complete probability information.

Method

We recruited 120 participants from Amazon Mechanical Turk for Experiment 3A, and a different group of 119 participants for Experiment 3B; 5 were excluded because they incorrectly answered more than one-third of the check questions (2 and 3 from Experiments 3A and 3B, respectively), and 9 because their total probability ratings for at least one item were not between 80% and 120% (6 and 3 from Experiments 3A and 3B, respectively).

Rather than manipulating participants' inferences using simplicity, participants in Experiment 3A were simply told the prior probability of each cause. They first read about the probability of Z given either cause A or cause B. For example, in the *low/low* condition, participants read:

> When a lake has Juga snails [A], it <u>occasionally</u> has bacteria proliferation.
> When a lake has Scuta snails [B], it <u>occasionally</u> has bacteria proliferation.

The *high/low* and *low/high* conditions differed as in Experiments 1 and 2 (changing "occasionally" to "usually" either for A or B, respectively). Next, participants were told the prior probabilities of the causes:

> Crescent Lake has a 65% chance of having Juga snails and a 35% chance of having Scuta snails.

These probabilities were adjusted across the three vignettes to match the probabilities of the simple and complex explanations obtained empirically in Experiment 2. Then participants were asked to rate the probability that the lake had each kind of snail, just as in Experiment 2. Finally, participants rated the probability of Z, using the same scale as Experiments 1 and 2. Counterbalancing and randomization were the same as in Experiments 1 and 2.

Experiment 3B was identical to Experiment 3A, except that the conditional probabilities were also numerically specified. Specifically, the word "occasionally" was always followed by the parenthetical "(about 20% of the time)" and the word "usually" was always followed by the parenthetical "(about 80% of the time)."

Results and Discussion

Figure 3 plots the results for both Experiments 3A and 3B. As in the other experiments, most participants rated the probability of A higher than the probability of B for all three items (91 out of 112 for Experiment 3A and 72 out of 113 for Experiment 3B). Ratings of P(A) and P(B) were tightly clustered around the values given in the problem (for P(A), $M = 65.40$, $SD = 2.08$ and $M = 65.38$, $SD = 2.54$ for Experiments 3A and 3B; for P(B), $M = 34.72$, $SD = 2.46$ and $M = 34.61$, $SD = 2.63$). Thus, judgments of P(A) and P(B) were very similar here to judgments of P(A) and P(B,C) in Experiment 2.

For Experiment 3A, inferences about P(Z) were similar to Experiment 2. Participants gave somewhat higher estimates of P(Z) in the *high/low* than in the *low/low* condition [$M = 71.49$, $SD = 19.01$ vs. $M = 66.04$, $SD = 25.68$; $t(90) = 2.03$, $p = .045$, $d = 0.21$, $BF_{01} = 1.65$], although this effect was surprisingly small. Most
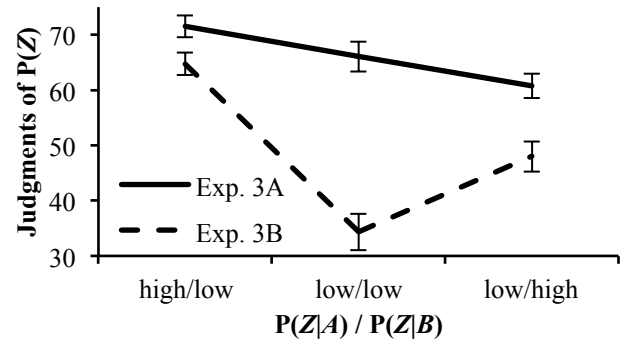

Figure 3: Results of Experiments 3A and 3B.

importantly, however, estimates of P(Z) in the *low/high* condition were no higher than in the *low/low* condition and were, if anything, somewhat lower [$M = 60.71$, $SD = 21.02$; $t(90) = -1.78$, $p = .078$, $d = -0.19$, $BF_{01} = 2.58$]. That is, once again, people did not take into account the possibility of the low-probability alternative (B) when estimating P(Z). This result suggests that people adopt beliefs in an all-or-none manner not only when the belief is the result of a categorization or a causal inference, but even if the belief is determined by prior probability alone.

These results stand in contrast to those of Experiment 3B. This experiment differed from Experiment 3A only in giving precise values of P($Z|A$) and P($Z|B$), so that participants could in principle calculate P(Z) exactly. Here, participants differentiated not only between the *high/low* and the *low/low* conditions in their ratings of P(Z) [$M = 64.75$, $SD = 17.63$ vs. $M = 34.33$, $SD = 27.90$; $t(71) = 11.81$, $p < .001$, $d = 1.39$, $BF_{10} > 1000$], but also gave higher estimates of P(Z) in the *low/high* condition [$M = 47.99$, $SD = 23.36$; $t(71) = 4.88$, $p < .001$, $d = 0.58$, $BF_{10} > 1000$]. Thus, people are *aware* that the probabilities of lower-probability beliefs are relevant. They simply do not spontaneously use those probabilities if they are not given explicitly in the problem.

This difference between Experiments 3A and 3B was also evident when we compared participants' responses to normative benchmarks. We used the same strategy as in Experiment 2 to calculate, based on each participant's other probability ratings, how large that participant's difference in P(Z) ratings should be between the *low/high* and *low/low* conditions. Whereas participants in Experiment 3A underutilized P($Z|B$) by a substantial margin [$M = 8.95$, $SD = 21.00$; $t(111) = 4.51$, $p < .001$, $d = 0.43$, $BF_{10} = 786.09$], participants in Experiment 3B were better calibrated and underutilized P($Z|B$) to a smaller degree [$M = 5.09$, $SD = 19.87$; $t(112) = 2.72$, $p = .008$, $d = 0.26$, $BF_{10} = 2.61$].

General Discussion

Do beliefs come in degrees? The current studies suggest that they may not—that when making predictions from uncertain beliefs, those beliefs are treated as either true or false, without reflecting the uncertainty that people profess when asked explicitly. In Experiments 1 and 2,

people acknowledged that a simple explanation had a 65% chance of accounting for some observations, while a complex explanation had a 35% chance. In making subsequent predictions dependent on the correct explanation, however, people ignored the lower-probability complex explanation, treating the simple explanation instead as though it were certainly true. In Experiment 3A, participants even ignored low-probability beliefs when the prior probabilities were given explicitly.

However, when participants in Experiment 3B were given all relevant probability information, they were able to take low-probability possibilities into account. Although further work will be necessary to pinpoint the reason for this effect of task context, one possibility is that when all relevant probability information is given, participants are able to treat the inference as a math problem rather than relying on their intuitive belief systems. Even if participants are unable to produce the precise Bayesian solution, they may recognize that all four pieces of information are relevant and scale their responses in qualitatively appropriate ways. Future research might also examine other conditions that may lead participants to combine multiple potential beliefs, such as priming participants with problems with two equally likely possibilities, where neither can be ignored.

If people represent beliefs implicitly as all-or-none, then why do they nonetheless profess uncertainty when asked explicitly? That is, why do participants not claim that there is a 100% chance that the simple explanation is true, when asked explicitly? One possibility is that beliefs such as 'there is a 65% chance of possibility *X*' can be represented explicitly but that when we must rely on such beliefs for subsequent inferences, they must be converted to the 'digital' format. For example, when people are planning what to wear during the day, they are clearly able to represent explicitly the possibility that there is a 65% chance of rain. But when they must use that belief implicitly in subsequent reasoning (e.g., to determine whether the road will be slippery), people appear unable to use probabilities in a graded manner.

This possibility is consistent with the *singularity principle* (Evans, 2007), according to which people focus on one possibility at a time in hypothetical thinking. For example, when told about a cause that can lead to an effect, people ignore other possible causes that could be in operation, focusing on the focal cause when estimating the probability of the effect (Fernbach, Darlow, & Sloman, 2011). The current results show that people even neglect alternative causes in predictive reasoning when one cause is merely more *likely*, rather than certain.

These results are challenging for probabilistic theories of cognition (Anderson, 1991; Oaksford & Chater, 2009), in that the very purpose of probability is to reflect degrees of uncertainty (Jeffrey, 1965). Graded beliefs are critical for Bayesian updating, or modifying one's beliefs in light of new evidence. The current results point to differences between implicit and explicit representations of beliefs,

since people can simultaneously profess 65% confidence in an explanation, but treat it in subsequent inference as though they are 100% confident. Thus, people may look more or less like Bayesians depending on the nature of the task and the associated cognitive architecture.

Abductive (data-to-explanation) and predictive (explanation-to-predicted-data) reasoning are critical to diverse cognitive processes, including not just causal reasoning and categorization, but also decision-making, perception, and social cognition. Therefore, an important goal for future research will be to test whether a digital belief architecture is confined to high-level cognitive tasks (such as causal reasoning and categorization), or whether it might instead be a common architectural constraint across many cognitive domains.

Acknowledgments

References

Anderson, J.R. (1991). The adaptive nature of human categorization. *Psychological Review*, *98*, 409–429.

Dietrich, E., & Markman, A.B. (2003). Discrete thoughts: Why cognition must use discrete representations. *Mind & Language*, *18*, 95–119.

Evans, J.St.B.T. (2007). *Hypothetical thinking: Dual processes in reasoning and judgement*. New York, NY: Psychology Press.

Fernbach, P.M., Darlow, A., & Sloman, S.A. (2011). Asymmetries in predictive and diagnostic reasoning. *Journal of Experimental Psychology: General*, *140*, 168–185.

Jeffrey, R.C. (1965). *The logic of decision*. New York, NY: McGraw-Hill.

Kemeny, J.G. (1955). Fair bets and inductive probabilities. *Journal of Symbolic Logic*, *20*, 263–273.

Lombrozo, T. (2007). Simplicity and probability in causal explanation. *Cognitive Psychology*, *55*, 232–257.

Murphy, G.L. (2002). *The big book of concepts*. Cambridge, MA: MIT Press.

Murphy, G.L., & Ross, B.H. (1994). Predictions from uncertain categorizations. *Cognitive Psychology*, *27*, 148–193.

Oaksford, M., & Chater, N. (2009). Précis of *Bayesian Rationality: The Probabilistic Approach to Human Reasoning*. *Behavioral and Brain Sciences*, *32*, 69–120.

Rouder, J.N., Speckman, P.L., Sun, D., Morey, R.D., & Iverson, G. (2009). Bayesian *t* tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, *16*, 225–237.

Waldmann, M.R., & Holyoak, K.J. (1992). Predictive and diagnostic learning within causal models: Asymmetries in cue competition. *Journal of Experimental Psychology: General*, *121*, 222–236.