

Uncertainty without all the doubt

Forthcoming in *Mind & Language*.
Aaron Norby

Abstract: I investigate whether degreed beliefs are able to play the predictive, explanatory, and modelling roles that they are frequently taken to play. The investigation focuses on evidence—both from sources familiar in epistemology as well as recent work in behavioral economics and cognitive psychology—of variability in agents’ apparent degrees of belief. Although such variability has been noticed before, there has been little philosophical discussion of its breadth or of the psychological mechanisms underlying it. Once these are appreciated, the inadequacy of degrees of belief becomes clear. I offer a theoretical alternative to degrees of belief, what I call the *filter theory*.

Introduction

Call the following the *storage hypothesis*: degreed beliefs are stable attitudes of agents that can be appealed to for predicting and explaining patterns of behavior, judgment, and decision making. The storage hypothesis is a widely held assumption in philosophical work on uncertainty and degreed belief, and it isn’t hard to see why. The need to appeal to *degrees* of belief, over and above ‘outright’, yes-or-no beliefs, is suggested by such facts as that humans have a capacity for being uncertain about the truth of propositions; and this uncertainty can take on a wide range of gradations, making for a correspondingly wide range of behavioral differences. For example, a person can become a little bit more

Acknowledgments: Thanks to Tamar Gendler, Zoltán Szabó, Keith DeRose, Justin Khoo, Josh Knobe, Esther Shubert, Bruno Whittle, members of Yale’s Works in Progress seminar, and two anonymous reviewers for helpful comments on previous drafts.

Address for correspondence: Aaron Norby, Department of Philosophy, Yale University, New Haven, CT 06511, USA.

Email: aaron.norby@yale.edu

or a little bit less uncertain than she was before about whether her car has fuel injection. This difference in uncertainty will, it seems, make a difference to her potential behavior. If she becomes more confident that her car is fuel injected, we can suppose she'll become more likely to buy fuel-injector cleaner when it goes on sale at the store, more willing to tell others that her car is fuel-injected, and more likely to attribute engine trouble to a problem with a fuel injector. And just how much her behavior is changed by her change in uncertainty seems to depend on the degree of the change in uncertainty. It seems a good way to make sense of this is in terms of her having some degree of belief, or credence (I will use 'degree of belief' and 'credence' interchangeably), in the proposition that her car is fuel-injected.¹

While I think it's correct that we need more than outright belief in order to explain human judgment and decision making, I also think that the storage hypothesis is false and that degrees of belief cannot in fact play either the explanatory or predictive roles they are generally thought to play. The reason, I will argue, has to do with the way in which uncertainty in human cognition is subject to a systematic kind of variability or instability that cannot be explained in terms of familiar processes like reasoning or belief-update. This suggests that an agent's degrees of belief (insofar as anyone be said to have them) do not remain stable from situation to situation, and so whatever states or attitudes can be used to predict and explain behavior, they don't include degreed belief in any of its familiar forms.

In the next section, after describing the storage hypothesis in more detail and saying why I think it is so widely held, I will make a few remarks about what it means to reject

¹ It should be clear that what we are after is a stable attitude and not what we might call the 'feeling' of uncertainty, the occurrent sense of a lack of confidence that might accompany the explicit pondering of whether one's car has fuel injection. The two might be tightly connected, but we would have to do some investigating to find out how. See DeRose (2009, pp.190–3) for more on the relevant distinction between two notions of confidence.

it. In the following section (section 2), I'll make the case for rejecting it. And in section 3 I'll offer a theory to replace it, which in a nutshell says that insofar as we have credences, they are to a significant degree constructed or formed anew each time there is a decision or judgment to make. I propose that we appeal to a state other than credence, which for lack of a better term I call *proto-credence*, and which can do the explanatory work credence was meant to do while giving us a more comprehensive and realistic picture of ourselves. My goal is not to dismiss all talk of credences, but rather to dismiss a particular way of thinking about them and their role in decision making—embodied in the storage hypothesis—and to describe what else is needed to understand the way that we make decisions and form judgments.

1 The Storage Hypothesis

The storage hypothesis is much weaker than the name might imply, and correspondingly my claim that it is false is somewhat stronger than it might seem. The term ‘storage hypothesis’ is meant to cover not only the suggestion that degrees of belief (/credences/degrees of uncertainty) are Really Stored in the mind of the agent (e.g., encoded in a sentence-like form somewhere in the brain), but also the weaker suggestion that it is merely *as if* degrees of belief are stored (as, e.g., some who hold that such attitudes are ‘implicitly’ encoded in a connectionist network might say; or as classical behaviorists would claim; or as those who say that credences can merely be used to ‘model’ behavior might claim).²

It is a corollary of the storage hypothesis in both its weak and strong form that an agent’s having a particular degree of belief in a particular proposition makes a systematic impact—in combination with a range of her other attitudes—on that agent’s judgment and decision-making. This corollary is not open to much negotiation, since what it says is only

² It should be clear, then, that the scope of the problem will not be limited to behaviorist accounts of degreed belief.

that an agent's degreed beliefs are not idle with respect to her cognitive and behavioral processes: if it is true, for example, that a subject is to some degree uncertain about whether p is true, then that uncertainty will make a systematic difference to the subject's overall pattern of judgments and decisions. To deny this is to deny that uncertainty and degreed belief should be part of a theoretical account of human cognition.

There are of course a variety of theoretical options for saying exactly *how* degrees of belief impact on agents' judgment and decision making. One particularly clear case is expected utility theory, where degrees of belief are modelled as subjective probabilities and agents are expected to make those decisions that maximize their subjective expected utility (where agents' utility functions are standardly derived from their preferences over pairs of lotteries). But that is just one way of spelling things out. As a descriptive theory, expected utility theory is widely derided (Gigerenzer, 2008; Gilovich *et al.*, 2002; Kahneman & Tversky, 1979; Stanovich, 2011). So, a defender of the storage hypothesis might take degrees of belief to be something other than subjective probabilities; and agents' *decision strategies*—how agents make decisions on the basis of their attitudes—might be taken to be something other than expected utility maximization. For example, it would be reasonable to suppose that agents use some sort of satisficing, rather than maximizing, decision strategies (Simon, 1955), and we might suppose that credences come in only a few degrees rather than being real-valued (for example, 'certain', 'likely', 'about as likely as not', 'unlikely', 'certainly not', could serve as terms for degrees of belief in a theory that recognized five degrees).

And again, one might hold some version of the storage hypothesis while making no claims about the psychological mechanisms underlying the phenomena that are being modelled. For example, one might very well hold that expected utility theory is an excellent way of modelling decision makers and their judgment and choice behavior without also

holding that there is any sense in which decision makers make choices by calculating expected utility. The theory may be pitched at a highly abstract level so that its empirical grounding comes only from the hypothesis that it is *as if* decision makers were making decisions by calculating expected utilities (which are derived from their actual preferences). This requires only predictive success—along with the ability to organize the behavior of a subject into coherent patterns—rather than mechanism-level explanation.

Put simply, the storage hypothesis is that humans have, or it is as if we have, degrees of belief, and our judgment and decision making across a wide range of situations is systematically predicted by them (more precisely, predicted by them in conjunction with some of our other states and processes, such as desires or preferences and intentions). In general, the truth of the storage hypothesis is a default assumption in philosophy.³

Keeping these facts in mind, we can think of the commitment to the storage hypothesis of degrees of belief and uncertainty as a commitment to the view that a person's degrees of belief are like a map, or a collection of maps, representing ways the world might be, along with, for each way it might be, how likely it is to be that way (or, more generally, how much weight each outcome will get in your decision making). You make decisions by consulting this map, using your operative decision strategy to pick a course (cf. Ramsey 1931).

³ The very idea that we can reasonably debate what the constraints are on rational degrees of belief—as is done in formal epistemology—seems to presuppose that we have degrees of belief that are a stable part of our cognitive lives (Christensen, 2004). Eriksson & Hájek (2007) argue that *degree of belief* should be treated as a primitive, but a primitive that earns its keep by being able to explain and predict a wide range of behavior and other psychological processes, such as inference, as well as affective states (p.208). Additionally, amongst those who seem to hold some version of the storage hypothesis but also recognize the kind of uncertainty-ignoring that I'll describe in the next section are Harman (1986, ch.3), Weatherston (2012), and Williamson (2000, pp.98–99), who, though they all have accounts that are in various ways different from one another, discuss degreed belief in terms of its relation to outright belief and to behavior. An exception is Clarke (2013), who argues that subjects' degrees of belief change from context to context depending on what possibilities are relevant. Clarke offers a formal framework for this view, but he does not there attempt to give an account of the nature of the underlying mental states, nor does he argue, as I do, that degreed beliefs might not play any real explanatory role. The importance of going beyond a formal model will come out in my criticism of multiple-sets-of-credence views in section 2.

The critical assumption is that, like holding a map in your pocket, having some degree of belief or other in some proposition is a stable, persistent state. The degree of belief that someone has is not dependent upon the situation that she is in or the decision problem that she is facing (except insofar as those might lead the agent to update her degrees of belief). The assumption that degrees of belief are not situation-bound is what gives accounts that incorporate the storage hypothesis, like expected utility theory, a claim to having predictive purchase and to being able to reveal coherent and interesting patterns in agents' judgments and decision making.

Thus, to give up the storage hypothesis is to allow for the possibility that an agent's degrees of belief cannot be specified independently of a specification of the decision problem she faces. This would be a radical revision to the ordinary way of doing things, since the ordinary way of doing things assumes that an agent's degrees of belief can be used to explain and predict her judgments and decisions because those degrees of belief are in place antecedent to the problems that prompt those judgments and decisions.⁴ My suggestion is that the ordinary way of doing things should be abandoned. Instead, we ought to recognize that degrees of belief are not so much stored and stable—waiting to guide decision-making when called upon like a map in one's pocket—but are rather in some sense shaped and constructed anew when we go to make decisions, not through reasoning but simply through the act of consulting them.⁵

2 The Psychology of Uncertainty

In this section, I'm going to present some evidence that I think strongly supports the claim that it is wrong in general to think that we carry around with us degrees of confidence

⁴ This is of course setting aside the subject's attitudes that are specifically about the concrete decision problems whose manner of resolution we are trying to predict and explain.

⁵ More precisely, through the act of consulting whatever mental states it is we have, if not degrees of belief.

(even intervals or ranges of degrees) in propositions from which it can be predicted (even by an omniscient observer) how likely we will treat those propositions as being when it comes time to make decisions. This evidence is important to my account not solely for its negative force against various incarnations of the storage hypothesis. It will also form the basis for the positive account that I will defend.

The evidence supports both my negative and positive claims by indicating a systematic variability in the way that humans think about possible states of the world, in particular a kind of variability that cannot be readily understood in terms of the familiar concepts of updating or revising one's doxastic state (as occurs via reasoning and with the acquisition of new evidence). The existence of much of this variability is, however, relatively uncontroversial.

In what sense is there an uncontroversial variability in the way we think about the world? First, notice that a big part of what allows us to get around in the world successfully is that we store in memory possible states of the world,⁶ which memories can be recalled at opportune moments and used as considerations in making decisions. Consider, then, a case in which someone, let's call him Lot, has entered a lottery that has a very low probability of winning but a very large payoff if he does win. And imagine that for planning purposes Lot has stored in memory two possibilities: first, the possibility that he wins the lottery (we can think of this as a memory whose content is expressed by 'Lot might win the lottery'⁷), and, second, the possibility that he won't be able to afford a new car this year (a memory whose content may be expressed by 'Lot might not be able to afford a new car this year'). Of course, thoughts of these possibilities cannot guide judgment and decision

⁶ This is a slightly inaccurate way of speaking. The possibilities themselves—whatever they are—are of course not stored in one's memory. Rather, representations—or something that plays roughly the same functional role as mnemonic representation—of them are stored in memory.

⁷ The idea here is only to get the content roughly right; there may very well be better ways to characterize it.

making unless Lot in some sense assigns weights to them corresponding to their relative likelihoods.⁸ But the weights Lot assigns, if he is like most people, will not be the same in all cases. Typically, if both possibilities are brought to his attention, he'll affirm that it's highly likely, though not certain, that he won't be able to afford a new car this year, and that it's highly unlikely, though not impossible, that he'll win the lottery (and thus be able to afford a new car). It's also the case, however, that for most purposes he's going to *ignore* the possibility that he'll win the lottery and treat it as a foregone conclusion that he won't be able to get a new car. If, for instance, he's thinking about whether he'll be driving an expensive sports car to work next year, Lot might seem to totally ignore the possibility that his ticket will win and take it for granted that his ticket will lose, concluding that he won't be driving an expensive sports car because he won't be able to afford a new car at any price. In that context the possibility of winning the lottery may not, unless it's explicitly raised, even cross Lot's mind, and so he'll think, talk, and act just as he would if he were absolutely certain that he was going to lose. *This* is the sense in which there is *prima facie* variability in the way we think about these sorts of possibilities: sometimes Lot thinks about the possibility of not being able to afford a new car as an inevitable eventuality, and sometimes he thinks of it as only being very likely but as also having real, live alternatives.

A similar sort of thing happens with the low but non-zero credence that most of us seem to give to various skeptical scenarios. I think it's very unlikely, though not of probability zero, that I'm a brain in a vat. Most of the time, however, this possibility is the furthest thing from my mind. Correspondingly, I think and talk as though it were certainly false: I think about my body as though it were absolutely certain I had one and about the external

⁸ This is so even if you are using a decision strategy like maximin (choose the option whose worst outcome is the least bad of the worst outcomes of each option), which is typically thought of as a strategy that does not use likelihood information (Sunstein, 2002). For, the possibility must have a non-zero likelihood for it to be considered a possible outcome to be taken into account at all.

world as though I were certain that it is more or less as though it appears to me to be.

Facts of this sort have been noticed before,⁹ but it is surprisingly rare for philosophers to pursue the question of *why* these phenomena occur. After all, if I really do think it's only highly likely, but not certain, that I'm not a handless brain in a vat, why isn't the uncertainty I have about the existence of my hands manifest in all my thoughts about them? Why aren't all such judgments hedged? One might reason that we have a general tendency to ignore low-probability events (Ross & Schroeder, 2012; Weatherston, 2005, 2012). But that seems more to describe the phenomenon than to explain it. More important would be an explanation of what it is about the way we store and recall possibilities in memory that makes these phenomena not only real but so common as to be hardly noticeable.

The reality is that it seems to be due to the fact that the way we form judgments and make decisions is by a process that stochastically retrieves a small number of items from memory, and this small set of items is then used as the basis for subsequent judgments (Giguère & Love, 2013; Higgins, 1996; Stewart *et al.*, 2006).¹⁰ We only ever call up some small subset of the attitudes that, under different circumstances, we might have called up. To put it differently, we don't evaluate options using our attitudes all together as a 'corporate body', but rather only as parts of smaller organs (cf. Quine 1953, p.41). Our degrees of belief are apportioned across those possibilities that we happen to call up on a given occasion. Consequently, whereas on one occasion we might apportion some positive degree of belief to the possibility that p as well as to a mutually exclusive alternative q , on another occasion where q is simply not brought to mind, the possibility that p will be treated as having no alternatives and so as a certainty.

⁹ For descriptions of the lottery case, see Hawthorne (2004); Whittle (2012). In the context of skeptical scenarios, Descartes notices, in the first *Meditation*, our tendency to ignore in most situations what we acknowledge to be possible in others, where he refers to it as 'habitual opinion', which he attributes to the 'law of custom' (CSM 2:15). It is also noticed by Locke (*Essay*, Bk.4, ch.15) and by Hume (*Enquiry*, §V).

¹⁰ An influential model of how this sampling operates is Anderson (2007)'s ACT-R cognitive architecture.

To return to the lottery example, when Lot is explicitly thinking both about the possibility of buying a new car and the possibility of winning the lottery he's entered, he partitions things into the mutually exclusive alternatives of *Lot can't afford a new car* and *Lot wins the lottery and can afford a new car*. He gives a little credence to the possibilities in the latter cell and much more to those in the former. But in most situations, where he's ignoring the possibility of winning the lottery, the possibility of winning the lottery is simply not on the table: he does not call the possibility up from memory and so does not use it as a consideration. Consequently, it is as though all of Lot's credence is apportioned to the possibility that he won't be able to afford a new car and he does not countenance any situations incompatible with it.

This process of stochastic recall and reapportionment of degrees of belief has important implications for how we ought to think about what degrees of belief *are*. I've already said that what I think it shows is that we should not think of degrees of belief as something we have or carry around with us; but notice also that what I've been describing is a process in which (in broad outline) a credence is assigned to a possibility upon its being recalled from memory, where the credence that is assigned is in part dependant upon which other possibilities happen to be brought to mind at the time. The *stable* state of the agent, then, is not the credence but the state realized by the stored memory, which does not itself have a credence attached to it. This underlying state we can think of as a *proto-credence*. I will expand on this idea at greater length in section 3.

One way to try and maintain the basic spirit of the storage hypothesis in the face of the kind of variability illustrated by the lottery and skeptical-belief examples is to say that humans often have *multiple sets* of conflicting credences, each set of which is manifested in distinct situations (Egan, 2008; Lewis, 1996; Stalnaker, 1984; Wedgwood, 2012; Whittle, 2012). But if the basic story that I've been telling is right about how the underlying

psychology of these cases works—the stochastic-recall-and-reapportionment-of-degrees-of-belief story—then invoking multiple sets of credences looks like it’s at best a superficial re-description of the phenomena that doesn’t really tell us at all what credences are like or what contribution they make to our thought and behavior. For, given the cognitive processes as I’ve described them, the likelihood that you assign to, say, your being able to buy a new car, is not due to your calling up some one attitude toward that proposition rather than a different attitude, but is much better described as one in which your attitude is *formed* on the spot. This, I think, is the right way to think of the reapportionment of your credences based on the partition of the space of possibilities provided by the stochastic recall processes involved.

If one were to insist that, even so, thinking of these phenomena in terms of multiple competing sets of credences is a useful abstract *model* of what’s going on, I’d have to say that *of course* we can model such phenomena in terms of multiple sets of attitudes. Trivially, we can do so for any agent no matter how inconsistent, irrational, or unpredictable are its apparent attitudes. I think that high-level models that abstract from the psychological details are extremely useful in both cognitive scientific and philosophical theorizing, but a model that can be applied no matter what the phenomena look like is, it seems to me, useless.

A related strategy to try to preserve the storage hypothesis is to appeal to so-called *outright* belief. A number of philosophers have recently suggested that in cases like those we’ve been discussing, it is outright belief, rather than degreed belief, that is at work where an agent apparently ignores some small-probability event that they think might occur, like Lot ignoring the possibility of winning the lottery; the agent’s degreed belief surfaces, on these views, only when the possibility of the small-probability event’s occurring becomes relevant in some way (Foley, 1993, 2009; Harman, 1986; Holton, 2008; Nagel, 2010; Ross

& Schroeder, 2012; Wedgwood, 2012).

It is important to see, however, that even if we make this terminological choice and refer to such cases as cases of outright belief, we won't have gotten any closer to understanding them or to understanding what uncertainty and degreed belief are. We will still have the question: What *is* uncertainty such that sometimes we ignore it? Saying we have such a tendency and calling it (or its manifestation) 'outright belief' will not help answer the question. In addition, soon I will describe cases that display a variability of degree of belief much like those we've already looked at, but that do not involve the ignoring of small probability events in a way that would be plausibly amenable to analysis in terms of outright belief.

A more difficult question that one might raise in defense of the storage hypothesis is whether in the sorts of examples I've so far used—in lottery cases and in skeptical-belief scenarios—it is reasonable to suppose that the agents are *really* ignoring, in the sense of completely failing to consider, the possibilities that they seem to be ignoring rather than that they, or their decision-making systems, have simply deemed them to be irrelevant to the tasks at hand.¹¹ I have suggested that the former is the most plausible way to understand such cases, but it is perhaps worth pointing out that experimental work is likely needed in order to settle the issue.

Preliminary work from a series of studies designed to test this very question suggests that, indeed, in many cases agents genuinely do ignore low-probability events in a robust sense, even when they have knowledge of those events and even when the events have a high enough likelihood as to make them relevant to the agent's decision making. I describe those results here only in broad outline.

Norby & Phillips (2013) had subjects learn about two novel categories of objects (e.g.,

¹¹ Thanks to an anonymous reviewer for raising this question.

types of made-up cartoon robots), which for simplicity I'll refer to as category *High* and category *Low*. Subjects learn, for example, that 90% of the objects with property *F* are members of category *High* and 10% of objects with property *F* are members of category *Low*. Subjects are then shown a series of ten new objects with property *F*, after which they are asked whether all ten of the objects in the series belong to *High*. Subjects are told that they will receive a small monetary prize for answering correctly. The hypothesis was that subjects would, for each of the ten objects, completely ignore the small (.1) probability that the object is from *Low* rather than *High*. This generated the prediction that subjects would answer that *all ten* objects they saw belonged to *High*, even though the probability of that being the case would be $.9^{10} \approx .35$. If, on the other hand, subjects were not genuinely ignoring the small chance that, for each object, that object belongs to *Low*, they would realize that each of those small chances would aggregate across the series of ten to yield a relatively high probability (.65) that at least one object belongs to *Low*. In fact, a significant majority of subjects answered that all ten objects were members of *High*. Results are thus in line with the hypothesis that subjects often genuinely ignore low-probability events, even when they are plainly relevant to the task at hand, and correspondingly they treat high-probability events as though they were certainties.¹²

I mention this work in order to describe one of the ways in which the questions under discussion can be advanced by further experimental research. Even if the results described are taken at face value, many questions remain—questions, for example, about the mechanisms at work and the scope of our ability to ignore uncertainty. In much of what follows I will describe the overall shape that I think the answers to these questions are likely to take.

¹² A series of control conditions ensured (i) that subjects were not simply making errors in calculating probabilities, and (ii) that subjects did not simply fail to learn that the chance of an object with property *F* being from *High* is only 90%. If, for example, subjects were forced to think of all ten objects at once, rather than serially, they were more likely to answer that at least one of them was from *Low*.

In the next section I'm going to describe an interesting phenomenon called *unpacking*, which is widely attested in research in behavioral economics and that appears to demonstrate the instability of credences through the elicitation of both subadditive and superadditive probability estimates. The purpose of doing so is three-fold. First, it will vividly illustrate that the evidence against the storage hypothesis is not limited to the way we think about lotteries, skeptical scenarios, and other similar situations of interest to epistemologists: the issue is a much broader one. Second, the evidence is in many ways more compelling than the lottery and skepticism cases. Third, it will help me to develop more fully the theory of proto-credences (and credences) that I've already described in rough outline but that requires filling in.

2.1 Unpacking Effects

I've been maintaining that we store possible states of the world in memory, to be used in judgment and decision making, and that when we recall them for use, we do so in a way that leads us to partition the space of possibilities differently in different situations. And, citing recent psychological research, I said that this involves a process of stochastic sampling. This then leads to treating in different situations the same possibility as having different degrees of likelihood. With lotteries, for example, when we fail to bring to mind the possibility of winning, we are led to treat it as a forgone conclusion that we're going to lose. And this is so despite the fact that, to all appearances, we don't do anything to change our opinion on the matter—we don't reason about it and we don't get new evidence. As a result of this, I said, we should be careful not to think of, e.g., Lot's attitude toward the proposition *Lot won't be able to afford a new car* as being something that is ready and waiting to play a particular role in guiding his thought. Rather, the credence that he gives to that proposition—and thus the attitude that he has towards it—depends upon

what *other* possibilities are called to mind at the same time. Importantly, this variability in credence is underwritten by something relatively stable, namely, the stored memories that encode the various possibilities toward which he might take one attitude or another.

The key to this suggestion is the claim that, insofar as we have credences or degreed beliefs, they are surprisingly dynamic, transient, and situation-bound. The goal of this section is to further support this claim with an example that has not gotten much attention in the philosophical literature.

What's known in behavioral economics as *unpacking* is the breaking down of a description of an event into components that are mutually exclusive and jointly exhaustive with respect to the original event. For example, take the event *Colleen going to college*. This can be unpacked into the disjunction of *Colleen goes to U of State* or *Colleen goes to State A&M* or *Colleen goes to college somewhere else*. There is abundant empirical evidence that unpacking events in this way can reveal subjective probabilities that are subadditive.¹³

Suppose, for example, that you were to ask Colleen, 'How likely is it that you're going to go to college?' She'll give some estimate, call it $\mathbf{p}_1$. Suppose, also, that you were to ask Colleen (not in addition to but instead of the first question), 'How likely is it that you're going to go to U of State or State A&M or some other college?'. Call the estimate she would give in response to this question, $\mathbf{p}_2$. And then suppose that you were to ask Colleen to give separate estimates for each of the three possibilities (that is, suppose you were to ask her on separate occasions, 'How likely is it you'll go to U of State?'; 'How likely is it you'll go to State A&M?'; 'How likely is it you'll go somewhere other than U of State or State A&M?'). Call her responses to each of these questions, respectively,

¹³ If $A_1 \dots A_n$ are events that together partition the event A , and $\mathbf{P}(A_k)$ is the subjective probability assigned to event A_k , then a subject's subjective probability assignment is additive iff $\mathbf{P}(A) = \sum_{i=1}^n \mathbf{P}(A_i)$ (more precisely, her assignment is finitely additive iff this equality holds for all events and their finite partitions). The assignment is subadditive iff $\mathbf{P}(A) < \sum_{i=1}^n \mathbf{P}(A_i)$, and it is superadditive iff $\mathbf{P}(A) > \sum_{i=1}^n \mathbf{P}(A_i)$. The case where $\mathbf{P}(A) < \mathbf{P}(A_1 \vee \dots \vee A_n)$ is sometimes said to be 'implicitly subadditive'.

p₃, p₄, p₅. When experimenters ask subjects packed and unpacked versions of questions like these (using between-subjects designs), what they find in many cases are patterns of subadditive judgments (Brenner *et al.* , 2002; Rottenstreich & Tversky, 2002; Tversky & Koehler, 1994). That is, their findings suggest that Colleen’s various estimates would show the following pattern: $p_1 < p_2 < p_3 + p_4 + p_5$. That is, her estimate of the disjunction would be higher than that for the ‘packed’ statement, and her estimates of each of the individual disjuncts would add up to more than her estimate of the disjunction.

Why is this? One explanation that has been offered is that disjunctions are estimated to be more likely than their corresponding ‘packed’, non-disjunctive, descriptions because the explicit mentioning of possibilities makes it more concrete the various ways in which the disjunction could be true. And when the disjunctions are considered separately, it might be that they get, collectively, higher estimates because when considered individually there is no other possibility ‘competing’ with it to be the most likely of a set of mutually exclusive outcomes.¹⁴ As Tversky & Koehler (1994, p.565) put it, continuing our cartographic theme:

Like the measured length of a coastline, which increases as a map becomes more detailed, the perceived likelihood of an event increases as its description becomes more specific. This does not imply that judged probability is of no value, but it indicates that this concept is more fragile than suggested by existing formal theories.

However, things are somewhat more complicated than this, as it has also been found that certain ways of unpacking events leads to *superadditivity*, that is, where the packed

¹⁴ This explanation is framed by the authors of these studies in terms of *support theory* (Brenner *et al.* , 2002; Rottenstreich & Tversky, 2002; Tversky & Koehler, 1994). According to support theory, subjects estimate probabilities by first judging the degree of evidential support for the focal hypothesis as against salient alternatives (where hypotheses are descriptions of events or outcomes). Because the psychological availability of evidence for a particular hypothesis can be influenced by apparently superficial features of the context, like whether the hypothesis is an unpacked or packed description of an event, judged support, and thus probability, of an hypothesis will vary across contexts. In particular, relations of support are theorized to be subadditive, which in turn generates subadditivity in probability judgments.

description is judged to be more likely than the unpacked version. In one set of experiments, it was shown that this can occur when events are unpacked into atypical instances (Sloman *et al.* , 2004). For example, unpacking *buying a kitchen utensil* into *buying a strawberry slicer or buying an oyster knife or buying some other utensil* might lead to probability estimates for the disjunction that are lower than estimates for the packed description, because strawberry slicers and oyster knives are relatively atypical utensils and so might make the prospect of buying a utensil seem more unlikely.

Likewise, more recently it has been suggested that unpacking an event into events that are relatively complicated or hard to process can lead to judgments more in line with superadditivity (Redden & Frederick, 2011). In one experiment, subjects indicated that they found a gamble described as one in which they would get \$50 if a fair six-sided die turned up an even number as being more attractive than a gamble described as one in which they would get \$50 if a fair six-sided die turned up a 1, 4, or 6 (Redden & Frederick, 2011, Study 2). Although the one gamble is not an unpacking of the other, the idea is that events that are composed of a relatively unnatural or complex set of events will be judged less likely than one that is relatively more natural and easier to process (the experimenters also included studies testing unpacking strictly conceived). The difficulty associated with the more complex event might be interpreted by subjects as indicating a lower degree of likelihood.

What’s interesting is that subjects in this experiment not only rated the gamble on the more natural event (the die coming up even) as preferable to the less natural one, they also judged the more natural event to be more likely, indicating that relative likelihood was driving preferences. What’s even *more* interesting, from our perspective, is that after indicating how much they ‘liked’ various gambles and rating their likelihoods on a seven-point scale (1 being ‘definitely lose’ and 7 being ‘definitely win’), subjects were given the

chance to estimate the probabilities of winning each gamble. Most subjects estimated the probabilities of both gambles to be 50%. Indeed, even when *only* those subjects who gave the correct probability estimates were included in the analysis, these subjects still tended to say that the natural event is more likely than, and a preferable gamble to, the less natural event (Redden & Frederick, 2011, p.161).

The point of this is not to be able to say, ‘Look how dumb and irrational we are,’ but is rather to make us notice how difficult it is to answer certain central questions about degrees of belief while we’re still trying to work within the framework of the storage hypothesis. How uncertain is Colleen about whether she’s going to go to college, and how will that uncertainty manifest itself in her decision making, that is, how uncertain will she treat the proposition that she goes to college as being? Unpacking effects present another way, in addition to those we’ve already seen, that these questions don’t seem to have good answers within standard ways of thinking about degrees of belief. In certain ways, she seems to have one level of uncertainty, and in other ways she seems to have other levels of uncertainty.¹⁵

Put another way, what unpacking suggests is that our credences are, at least to some degree, dynamic, transient and situation-bound. What they also suggest is that this dynamism is due to a process of credence *formation* based in part on which possibilities—as well as *how* those possibilities—are called up from memory, and that they are not due to something like the storage of multiple sets of competing credences or other attitudes. The evidence is better explained by the former hypothesis than the latter.

¹⁵ Can it be objected that unpacking phenomena involve irrationality, and that the storage hypothesis—and philosophical theories of degreed belief more generally—are designed to describe us only insofar as we are or approximate rational creatures? One thing to say in response is that we should prefer a theory that describes us more accurately rather than less, including our foibles and imperfections, so long as they are systematic foibles and imperfections. Second, I think it’s hard to say definitively if unpacking effects really do reveal irrationality. Insofar as they may be the result of how we operate within the cognitive limitations that we have—that is, insofar as they might represent an optimal trade-off between accuracy and the allocation of limited cognitive resources—we may want to demur in using the word ‘irrational’. This debate, however, is difficult to resolve. See Stanovich (2011, ch.1) for review.

For example, when Colleen thinks about how likely it is that she'll go to U of State, she may not bring to mind any other colleges that would serve as alternatives over which her credence could be spread. Now, if her credence in the proposition that she'll go to U of State were a stable attitude of hers, waiting to be used in her judgment and decision making—as implied by the storage hypothesis—then we wouldn't expect her *not thinking* about State A&M to make any difference to the level of confidence she ends up manifesting in her judgment about U of State. We'd expect the judgment of the likelihood of her going to U of State to be an expression of her preexisting attitude toward that proposition. If, however, we were to think of degrees of belief as formed in response to the items from memory that are recalled in the particular situation, then we wouldn't be surprised to find that unpacking makes a difference to the likelihood that Colleen seems to assign to various possibilities, and that these differences make a difference to her judgments and decisions.

Likewise, if we started to think of credences as formed and not stored, we wouldn't be surprised to find that someone might assign different likelihoods to *buying a kitchen utensil* and to *buying a strawberry slicer or buying an oyster knife or buying some other utensil*, even though it is obvious to the person that these are the same possible states of the world.

To put all of this in general terms: it seems our judgments might reflect what we think about something relative to some small part of what else we think, but there may be no sense in which our judgments reflect what we think, full stop. We are capricious creatures with malleable attitudes. We need a theory of our doxastic states that reflects this.

I should say that, although I've presented what I take to be some of the evidence that best illustrates what is wrong with the storage hypothesis, I do not take it to settle the issue, nor do I take it to be all the evidence that there is. Alternative approaches have produced evidence supporting conclusions complementary to those that I've reviewed. For example, Stewart *et al.* (in press) present an important line of research, which I will mention

only briefly, suggesting that the weight that agents attach to particular probabilities is influenced by the set of probability values that are readily available to working memory, and this influence induces a high degree of variability in how likely the relevant events are treated as being. For example, the extent to which an 80% chance of winning \$100 appears to be preferable to a 70% chance of winning \$100 depends on which other likelihoods or probability values have recently been thought of, even if they have nothing to do with the choice at hand.

In particular, Stewart *et al.* present evidence suggesting that probabilities are weighted according to their rank within the set of available probabilities. So, for example, if a subject has recently encountered a set of probabilities clustered *between* 70% and 80%, those probabilities (i.e, 70% and 80%) will be treated as being farther apart from each other than if the set of recently encountered probabilities is clustered below 70% or above 80%. The authors hypothesize that the reason for this is that when there are few readily available probability values falling between 70% and 80%, those two values will be similarly ranked. If, however, a subject has in mind many probability values falling between the two, then 70% and 80% will receive ranks that are farther apart from each other, and so the subjective weights given to the two values will be farther apart from each other in that condition. In short, how ‘good’ a chance an 80% probability of winning \$100 appears to be is in part a function of how that probability compares to the set of probabilities one has in mind. And that set can vary across circumstances.¹⁶

I suggested earlier—and in the next section I will elaborate on the idea—that the degree of belief that an agent assigns to an event depends upon what other events are brought to mind at the time. The work by Stewart *et al.* suggests the complementary idea that it can also depend upon which probabilities are readily available. For present purposes, the

¹⁶ Stewart *et al.* make similar claims for the weighting functions for amounts of money as well as for temporal delays.

important idea in both cases is that degrees of belief appear to be both variable and, rather than being stored, they are to some degree constructed from situation to situation.

One issue that nevertheless remains open is the scope of the variability in degrees of belief that I’ve been discussing. One may wonder whether the examples I’ve chosen point not to a general variability but rather to special cases outside of which there is a great deal of stability. It might be that there are conditions under which agents will manifest robust stability in their likelihood judgments—perhaps, for example, whenever likelihoods are thought of in terms of frequencies instead of probabilities, since we seem to have somewhat more facility with frequencies (Cosmides & Tooby, 1996).¹⁷ I think that to answer this question requires further empirical investigation, but two points are worth mentioning. First, even if the overall variability turns out to be somewhat limited, we still should seek out a theory of degreed belief that can account for what variability there is. Second, it is plausible that the variability demonstrated in the findings of Tversky & Koehler (1994), Redden & Frederick (2011), and Stewart *et al.* (in press) are each the result of distinct mechanisms, suggesting that the sort of variability at issue is not limited to a narrow set of circumstances.

We need, then, a theory of our doxastic states that reflects their malleability.

3 Proto-credences

I have tried to point to cases where agents do seem to have some kind of doxastic attitude toward a proposition—where it seems right to say, for example, that they are confident but not certain, and where it seems wrong to say that they have no opinion at all—but where there is no level or degree of uncertainty which they are disposed, in general, to take toward the proposition. We might put this by saying that our doxastic attitudes can fail

¹⁷ Thanks to an anonymous reviewer for bringing this issue to my attention as well as for suggesting the relevance of thinking in frequency format to the question of stability in degreed belief.

to be settled, or, at least, the role that an attitude is going to play in decision making can fail to be settled: how likely we're going to treat a proposition as being is subject to the influence of small perturbations from the way that we conceptualize the space of possibilities.

I've already indicated the basic outline of the theory I favor for explaining these facts and for getting a better understanding of the nature of degreed belief in light of unpacking and related effects. What I've been suggesting is that the correct theory involves explaining the likelihood that an agent assigns, on a given occasion, to a possible state of the world in terms of recall processes by which that possibility is brought to mind and has its likelihood assigned to it. The likelihood assignment is sensitive to various situational factors, including to which other possibilities are called to mind at the time. In this section, I'm going to flesh this theory out and offer a way to define traditional credences in terms of it. First, I'll introduce a distinction between two stages of the decision-making process. I'll then describe what I call *proto-credence*—a state that is a stable state of the agent but that cannot be reduced to credences or degreed beliefs. I characterize proto-credences themselves in terms of the first stage of the two-stage decision-making process, and define credence in terms of proto-credence.

First, then, the distinction between decision *making* and decision *set-up*. Whenever someone needs to make a decision/solve a problem/figure something out, there are always at least two decisions to make/problems to solve/things to figure out. There is the explicit decision to be made—should I go out or stay in?—and then there is the problem of how to approach and make that decision—which options should I consider? What contingencies do I need to take into account? Not only must a decision be made, but, prior to that, it must somehow be determined what courses of action and possible states of the world will be taken into consideration. This latter process we can call decision *set-up*, and we

can distinguish it from decision *making*, which we can understand to be the process of using the considerations selected by the decision set-up process to actually make a choice. In humans, decision set-up processes are non-trivial because our ability to discriminate amongst possible courses of action and store opinions about what the world is like outruns our ability to take all of those options and opinions into account in any given situation, even if we somehow antecedently limit ourselves to those that are topically ‘relevant’ to the decision at hand. It is my opinion that a general failure on the part of philosophers to acknowledge the importance of decision set-up is one reason that the problems raised here about degreed belief have been under-appreciated.

I’ve argued that credences do not have the connection to decision making and judgment that they are often thought to have, and that as a result they are inadequate both for explaining what our doxastic attitudes are and for explaining and predicting the decisions we make. What I propose to put in their place are what I have already called, for lack of a better term, *proto-credences*. I allow that there is a derivative sense in which we can be correctly said have credences, but the work they are generally meant to do is done in my theory by proto-credences. Thus, a version of what’s called the Bayesian Challenge, run by defenders of credence against defenders of so-called outright belief, might on a view like the one I propose be run against credences themselves. The Bayesian Challenge is that, once you have credences, which can purportedly do everything outright beliefs can do but more, we should forget about outright belief as far as serious theorizing goes (Jeffrey, 1970; Kaplan, 1996). This new ‘Bayesian’ Challenge would say that proto-credences can do everything that credences were supposed to do but better, and they’re psychologically realistic to boot.¹⁸

¹⁸ The reason for offering to define credences is that I do not want to close off the possibility that appeal to credences, even if they are epiphenomenal, may be useful for conceptualizing certain epistemological puzzles. At least, once we have a realistic picture of what credences are.

The theory that I want to propose goes like this. We explain what proto-credences are and what it is to be uncertain in terms of *inputs to the decision set-up process*. A proto-credence, on this view, is in the first place an input to a decision set-up process, not to decision making. Thus, this view posits a process that stands in between one's stable states and decision making. This is the critical feature of the approach that allows us to understand how deliberate decision making can be partially disengaged from an agent's doxastic states, while nevertheless offering a substantive account of those attitudes' role in cognition.

Before laying out the picture in detail, it's important that one more feature of decision set-up be established. It is crucial for what comes next that we keep this in mind: decision set-up processes are not all deliberate decision making processes; *they are, in large part, automatic, non-conscious processes*. Here's why this must be so. As we've already seen, all human deliberate decision-making processes require prior non-trivial decision set-up processes to occur. So, if all decision set-up processes were deliberate decision making processes, then all decision set-up processes would require prior non-trivial decision set-up processes. But then those decision set-up processes would themselves require prior non-trivial decision set-up processes, and so on. Thus, the start of any decision would begin an infinite regress.¹⁹ The conclusion is that decision set-up is in no small part automatic,²⁰ and we can tentatively draw the further conclusion (though it is not essential to the view), on the basis of its automaticity, that it is also largely unconscious (Stanovich, 2011).²¹

¹⁹ For other versions of this argument, see Joyce (1999, pp.70–77), Resnik (1987, pp.10–11), and Ross & Schroeder (2012, p.9). The first two citations are made in the third of these.

²⁰ Calling these processes 'automatic' does not explain *how* they operate or how the inputs to the process are selected. What it indicates is only that they are not deliberative procedures.

²¹ This argument also tells us why we can't explain an agent's ignoring possibilities in terms of something like supposition. The reason is that supposition is something that is taken up deliberately, with intention by the agent. Since decision set-up, and thus many instances of ignoring possibilities in decision making, is automatic and so non-deliberate and non-intentional, we can't understand these phenomena in terms of an attitude like supposition.

Decision set-up processes, as I've explained, are the means by which it is determined which considerations an agent will take into account in a given episode of deliberate decision making. We should think of these processes as implementing a set of rules for taking a highly complex range of possibilities that *might* be taken into consideration and on each occasion paring that set down to just those possibilities that will on that particular occasion actually be considered. In other words, the picture is this: there is a set of possibilities that might be taken into consideration; a largely unconscious, automatic decision set-up process takes that set as input and produces as output a simplified version, a sampling, of that larger possibility space. This more limited sample of possibilities is then what is used for the actual making of the decision for which it was produced. The sample is treated, for the purposes of decision making, as the entirety of the possibility space of which it is a sample. Based on the whole of what is in the sample, likelihoods are then assigned (again, largely unconsciously) to the represented possibilities. Thus, for example, if the possibility that *I won't be able to afford a new car* is called up but no possibility in which I can afford one (say, in which I win the lottery) is, then the possibility that *I won't be able to afford a new car* is going to be assigned, in that situation, a likelihood corresponding to certainty. On this proposal, we can analogize decision set-up to a *filtering* process, and the total set of possibilities an agent might take into account to the material the filter sifts through; on any given occasion, the possibilities that make it through the filter are those that are taken into account in explicit decision making, and those that don't are not.

What I propose is that we identify agents' proto-credences with these complex sets—more precisely, mental representations—of possibilities that serve as inputs to the decision set-up process. A set of proto-credences is composed of a representation of all those possibilities that an agent thinks might be actual, and this representation serves as input to the

decision set-up process.²² This is the heart of the *filter theory* of decision making. On the filter theory, probabilities *can* be assigned to proto-credences, but they do not correspond to subjective probabilities of the agent. Instead, what can be assigned are probabilities that the various possibilities in the agent's proto-credence set will be brought into working memory to be used for decision making. In addition, different probabilities would have to be assigned for different situations. This is how behavior can be predicted with the filter theory. For example, the likelihood of thinking of an alternative to *I won't be able to afford a new car* is higher when there is a recently purchased lottery ticket in my hand than when there isn't. This is, of course, a very different sense in which probabilities can be assigned than is meant in theories of subjective probability.

At this point, traditional credences can be defined. Credences, of course, come in degrees. Thus, the possibilities that are represented by an agent's set of credences must be weighted somehow, and that weighting scheme is itself to be spelled out in terms, once again, of a functional relationship with decision set-up. To have a particular degree of credence in a particular proposition is for the state that serves as input to the agent's decision set-up process to influence that process in a characteristic way (and to have a fitting representational content); more specifically, it is for the decision set-up process to potentially result in deliberate decision making in which the given proposition is taken as being likely to the degree corresponding to the degree of credence. To use the filter analogy, to have a proto-credence in a proposition is for that proposition to potentially slip through the decision set-up filter and into explicit decision making. To have a particular degree of (traditional) credence is for that proposition to potentially slip through and then be treated as having the degree of likelihood that corresponds to the degree of credence.

²² We might think of this representation as one big representation of all the possibilities the agent can distinguish, or as a collection of more specific representations. I'm neutral here. Also, although I am uncritically taking on an ontology that includes mental representations, the theory could be cast in terms that avoided representation talk.

Thus, on this view people might have more than one credence in a given proposition, since they might treat that proposition as having different degrees of likelihood on different occasions, depending on how the decision set-up process goes. Note, however, that this is different from having an interval-valued credence, which would reflect not an ambiguity in one's degree of credence but rather a lack of precision. Instead, we might want to identify an agent's degree of *proto-credence* in a proposition with the set²³ of likelihoods which she will potentially treat the proposition as having.

This is fairly abstract, so let's again return to Lot and his lottery, but, for now, let's put aside his tendency to ignore the possibility of winning. Suppose that we want to describe very broadly Lot's proto-credence in the proposition that he won't be able to afford a new car by saying that he thinks it's likely but not certain that he won't be able to afford a new car.²⁴ Since Lot thinks it's likely but not certain, there are possibilities in which he can't afford a new car that he thinks might be actual and possibilities in which he can that he thinks might be actual. So, the mental state corresponding to his set of proto-credences represents possibilities in which he can't afford a new car and also possibilities in which he can. What it is for these to be possibilities that he has a proto-credence in is for the state that represents them to interact with Lot's decision set-up process in such a way that the latter potentially includes them as considerations in the explicit decision making that results from decision set-up. In other words, it's for each of those possibilities to potentially make it past the decision set-up filter and into decision making.

To give the appropriate weights to each of those possibilities, we advert to the likelihood which Lot will potentially treat them as having, and this allows us to explain these weightings in functional terms. That is, to think that p is likely is for the decision set-up

²³ Not the range, since proto-credences may not be ordered by degrees in a meaningful way.

²⁴ I'm choosing the coarse description 'likely but not certain' because the filter theory is neutral with respect to what the correct 'grain' of degree of belief is. Whether there are infinitely many degrees of belief or only a few, the model can accommodate.

process to potentially result in deliberate decision making in which p is treated as merely likely and not- p is explicitly taken into account as possibly true. Thus, for Lot to think that the proposition that he won't be able to afford a new car is likely but not certain to be true, is for his decision set-up process to potentially result in deliberate decision making in which Lot treats it as likely but not certain that he won't be able to afford a new car and as unlikely but not ruled out that he will. But, as I mentioned earlier, because a proto-credence can give rise to treating a proposition with different degrees of likelihood in different situations, we will want to identify his degree of proto-credence not with any of these particular likelihoods but with the set of them.

More generally—and here is the important point—what constitutes Lot's being *uncertain* about whether he will be able to afford a new car is that: he is in a representational state that serves as input to his decision set-up process and in some circumstances²⁵ this decision set-up process results in deliberate decision making in which both the possibility that he won't be able to afford a new car and the possibility that he will are taken into account (because he takes into account the possibility of winning the lottery).

If one wants to retain traditional credences in addition to proto-credences but also wants to be able to assign unique credences, an agent's 'true' credence might be identified in terms of some privileged situation in which the relevant attitude plays a behavior guiding role. Perhaps this would be a situation in which the agent is maximally reflective and careful; or perhaps the 'true' credence could be identified with the lowest likelihood with which the agent is disposed to treat the proposition as having. My inclination, however, after all of the evidence and phenomena we've seen, is to hold that there is no particular credence that should be identified as the one the agent has.

²⁵ How is 'some' restricted? It is restricted only to circumstances in which the decision set-up process proceeds normally, where it is only the decision set-up process (and not, say, reasoning) that influences how likely the relevant possibility is treated as being, and where the input has an appropriate intentional content.

Any degree of proto-credence or credence that we want to include can be incorporated into the model in the description of what the agent's deliberate decision making is potentially like. One must only be able to specify, for each degree one wants to invoke, what it is to take a possibility into account in decision making in a way that directly corresponds to that degree. How exactly to do this—decision theoretically, perhaps, but certainly there are many other ways—is left open.

On the filter theory, we can, figuratively, think of every possibility that an agent has some proto-credence in as a little 'shout' to the decision set-up process that says, 'take me into account!' The decision set-up process lets through only those whose shouts are loud enough in each case in which that process is activated. The filter is thus more like an electronic signal filter than like, say, an air filter. For someone to think a proposition (merely) likely is for there to be an internal shout that is loud enough to make it past the filter and result in explicit decision making that treats that proposition as (merely) likely. The loudness of a shout, on this way of thinking, corresponds *not* to the degree of credence but to the likelihood of the content of the shout making it through the decision set-up process. The empirical questions this raises concern how those shouts are implemented and what determines the threshold loudness for making it through the filter. The work canvassed in §2—the research on unpacking effects as well as the remarks on stochastic sampling from memory—was meant in part to speak, in a very partial way, to these empirical questions. There is, however, a lot that is yet unknown.

Importantly, on this theory an agent's being uncertain about a proposition is perfectly compatible with many instances of decision making in which the decision set-up process and its input state result in her ignoring the possibility that that proposition is true. On the filter theory we say that what it is for an agent to have a degree of belief in a proposition is for her to have a decision set-up process and an input to that process which, in some

circumstances, cause her to take the possibility of that proposition’s being true into account in her decision making. To go back to Lot: his decision set-up process is such that some of the filterings of his proto-credences it produces let through the possibility that he will win the lottery and be able to afford a new car, and some do not. And on the filter theory this allows us to say that he has some degree of credence that he will win the lottery and be able to afford a new car even if he often completely ignores it.

As I indicated earlier, there is in fact a good deal of empirical evidence suggesting that what I’m here calling ‘decision set-up processes’ operate by a kind of filtering operation of the sort I’ve described, taking the most plausible, easily accessible possibilities (those with the loudest ‘shouts’) that are relevant to the task at hand, and filtering the alternatives out of consideration.²⁶ If that’s right, then how likely Lot is to ignore the possibility that he will win the lottery will not depend directly on things like how much is at stake, but rather more generally on factors that influence how likely he is to think of alternative possibilities, such as his baseline tendency to reflect on his immediate judgments (Frederick, 2005) and how recently he last thought about the lottery (Anderson, 2007), which themselves might be influenced by the perceived severity of what’s at stake. These are two of the many factors, part of a still-emerging picture from empirical psychology, that influence whether the shout of a given possibility will make it past the decision set-up filter. And these are facts that the filter theory, but not the storage hypothesis, accounts for.²⁷

²⁶ See Higgins (1996, for extensive, if early, review of evidence concerning factors that influence the accessibility of ideas to thought), Förster & Liberman (2007), as well as Giguère & Love (2013), Gilovich *et al.* (2002), Hilbig (2012), Morewedge & Kahneman (2010) Stanovich (2011), Stewart *et al.* (2006), and Kahneman & Frederick (2005, for review)) amongst many others. However, it is also important to note that many observed ‘distortions’ of probability information will be due not to selective retrieval from memory but to immediate features of presented outcomes and risks, for example affective richness (Harris *et al.* , 2009; Loewenstein *et al.* , 2001; Rottenstreich & Hsee, 2001; Sunstein, 2002) and psychological salience (Bordalo *et al.* , 2012).

²⁷ The approach to credences I’ve laid out has been influenced by the theory of *preference construction*: the idea that we do not carry preferences or values around with us, ready to be using in making decisions, but rather form them in relatively arbitrary ways when the situation demands it. See Lichtenstein & Slovic (2006), for extensive review.

What I'm attempting to give is a theory on which our fundamental doxastic states are attitudes whose impact on decision making is shaped by the psychological context in which they're elicited. The idea that we can in general say what a person thinks about this or that proposition, in a way that has any real significance for what that person will do and how they will make decisions, is mistaken. The attitudes that we carry around with us, insofar as they can be called attitudes, do not translate directly into behavior or even straightforwardly into facts about how likely we think this or that possibility is. As I said earlier: we're capricious creatures; and the filter theory allows our understanding of uncertainty and degrees of belief to reflect that.

4 Conclusion

In §1, I suggested that amongst those who countenance degreed beliefs, it is widely assumed that they (the degreed beliefs) are—or can be reasonably modelled as—stored and stable states of the agents who have them. I called this the storage hypothesis. The hypothesis forms the basis for the view that the behavior, judgment, and decision making of an agent can be understood in terms of her degrees of belief. There is little obvious reason for appealing to degrees of belief if they cannot play this role. In §2, I argued that the storage hypothesis is called into question by the existence of a systematic instability in our apparent degrees of belief, which is due to the way that processes governing recall from memory work. Based on familiar examples like lotteries and skeptical beliefs, as well as current psychological research on unpacking effects and selective recall processes, I suggested that the best explanation of our judgment and decision making says that we partition the space of possibilities differently in different situations, and that correspondingly we reapportion our uncertainties in each of those situations. And in §3, I laid out the filter theory, where

I explained how to define traditional credences in terms of proto-credences, which are themselves understood functionally in terms of decision set-up processes. The filter theory, I said, is much better at accounting for the phenomena canvassed in §2 than is the storage hypothesis, and so we should consider rejecting the storage hypothesis in favor of the filter theory.

A. Norby
Department of Philosophy
Yale University

References

- Anderson, John R. 2007. *How can the human mind occur in the physical universe?* Oxford University Press.
- Bordalo, Pedro, Gennaioli, Nicola, & Shleifer, Andrei. 2012. Salience theory of choice under risk. *The Quarterly Journal of Economics*, 1243–85.
- Brenner, Lyle A., Koehler, Derek J., & Rottenstreich, Yuval. 2002. Remarks on Support Theory: Recent advances and future directions. In: Gilovich, Thomas, Griffin, Dale, & Kahneman, Daniel (eds), *Heuristics and biases: The psychology of intuitive judgment*. Cambridge University Press.
- Christensen, David. 2004. *Putting logic in its place*. Oxford University Press.
- Clarke, Roger. 2013. Belief is credence one (in context). *Philosophers' Imprint*, **13**(11), 1–18.
- Cosmides, Leda, & Tooby, John. 1996. Are humans good intuitive statisticians after all? Rethinking some conclusions from the literature on judgment under uncertainty. *Cognition*, **58**(1), 1–73.
- DeRose, Keith. 2009. *The case for contextualism, vol.1*. Oxford University Press.
- Descartes, René. 1984. *The philosophical writings of Descartes*. Cambridge University Press. 3 vols. Translated by John Cottingham, Robert Stoothoff and Dugald Murdoch.
- Egan, Andy. 2008. Seeing and believing: Perception, belief formation, and the divided mind. *Philosophical Studies*, **140**(1), 47–63.
- Eriksson, Lina, & Hájek, Alan. 2007. What are degrees of belief? *Studia Logica*, **86**, 183–213.

- Foley, Richard. 1993. *Working without a net*. New York: Oxford University Press.
- Foley, Richard. 2009. Belief, degrees of belief, and the Lockean thesis. In: Huber, F., & Schmidt-Petri, C. (eds), *Degrees of belief*. Springer.
- Förster, Jens, & Liberman, Nira. 2007. Knowledge activation. *Pages 201–31 of: Kruglanski, Arie W., & Higgins, E. Tory (eds), Social psychology: Handbook of basic principles, 2nd ed.* New York: Guilford Press.
- Frederick, Shane. 2005. Cognitive reflection and decision making. *Journal of Economic Perspectives*, **19**(4), 25–42.
- Gigerenzer, Gerd. 2008. *Rationality for mortals*. Oxford University Press.
- Giguère, Gyslain, & Love, Bradley C. 2013. Limits in decision making arise from limits in memory retrieval. *PNAS*, **110**(19), 7613–18.
- Gilovich, Thomas, Griffin, Dale, & Kahneman, Daniel (eds). 2002. *Heuristics and biases: The psychology of intuitive judgment*. New York: Cambridge University Press.
- Harman, Gilbert. 1986. *Change in view*. MIT Press.
- Harris, Adam J.L., Corner, Adam, & Hahn, Ulrike. 2009. Estimating the probability of negative events. *Cognition*, **110**, 51–64.
- Hawthorne, John. 2004. *Knowledge and lotteries*. New York: Oxford University Press.
- Higgins, E. Tory. 1996. Knowledge activation: Accessibility, applicability, and salience. In: Higgins, E. Tory, & Kruglanski, Arie W. (eds), *Social psychology: Handbook of basic principles*. Guilford Press.
- Hilbig, Benjamin E. 2012. Good things don't come easy (to mind). *Experimental Psychology*, **59**(1), 38–46.
- Holton, Richard. 2008. Partial belief, partial intention. *Mind*, **117**(465), 27–58.
- Hume, David. 1999. *An Enquiry concerning human understanding*. Oxford University Press.
- Jeffrey, Richard C. 1970. Dracula meets Wolfman: Acceptance vs. partial belief. In: Swain, Marshall (ed), *Induction, acceptance, and rational belief*. D. Reidel.
- Joyce, James. 1999. *The foundations of causal decision theory*. Cambridge University Press.

- Kahneman, Daniel, & Frederick, Shane. 2005. A model of heuristic judgment. *In*: Holyoak, Keith J., & Morrison, Robert G. (eds), *The Cambridge handbook of thinking and reasoning*. New York: Cambridge University Press.
- Kahneman, Daniel, & Tversky, Amos. 1979. Prospect Theory: An analysis of decision under risk. *Econometrica*, **47**(2), 263–92.
- Kaplan, Mark. 1996. *Decision theory as philosophy*. Cambridge University Press.
- Lewis, David. 1996. Elusive knowledge. *Australasian Journal of Philosophy*, **74**(4), 549–567.
- Lichtenstein, Sarah, & Slovic, Paul (eds). 2006. *The construction of preference*. New York: Cambridge University Press.
- Locke, John. 1975. *An essay concerning human understanding*. Oxford University Press Clarendon.
- Loewenstein, George F., Weber, Elke U., Hsee, Christopher K., & Welch, Ned. 2001. Risk as feelings. *Psychological Bulletin*, **127**(2), 267–86.
- Morewedge, Carey K., & Kahneman, Daniel. 2010. Associative processes in intuitive judgment. *Trends in Cognitive Sciences*, **14**(10), 435–40.
- Nagel, Jennifer. 2010. Epistemic anxiety and adaptive invariantism. *Philosophical Perspectives*, **24**(1), 407–35.
- Norby, Aaron, & Phillips, Jonathan. 2013. Unpublished data.
- Quine, Willard Van Orman. 1953. Two dogmas of empiricism. *In*: *From a logical point of view*. New York: Harper & Row.
- Ramsey, F. P. 1931. *The foundations of mathematics and other logical essays*. Routledge and Kegan Paul.
- Redden, Joseph P., & Frederick, Shane. 2011. Unpacking unpacking: Greater detail can reduce perceived likelihood. *Journal of Experimental Psychology: General*, **140**(2), 159–67.
- Resnik, Michael D. 1987. *Choices: An introduction to decision theory*. University of Minnesota Press.
- Ross, Jacob, & Schroeder, Mark. 2012. Belief, credence, and pragmatic encroachment. *Philosophy and Phenomenological Research*.
- Rottenstreich, Yuval, & Hsee, Christopher K. 2001. Money, kisses, and electric shocks: On the affective psychology of risk. *Psychological Science*, **12**(3), 185–90.

- Rottenstreich, Yuval, & Tversky, Amos. 2002. Unpacking, repacking, and anchoring: Advances in support theory. *In: Gilovich, Thomas, Griffin, Dale, & Kahneman, Daniel (eds), *Heuristics and biases: The psychology of intuitive judgment*. Cambridge University Press. Originally published 1997.*
- Simon, Herbert A. 1955. A behavioral model of rational choice. *The Quarterly Journal of Economics*, **69**(1), 99–118.
- Sloman, Steven, Rottenstreich, Yuval, Wisniewski, Edward, & Hadjichristidis, Constantinos. 2004. Typical versus atypical unpacking and superadditive probability judgment. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **30**(3), 573–82.
- Stalnaker, Robert. 1984. *Inquiry*. Bradford MIT Press.
- Stanovich, Keith. 2011. *Rationality and the reflective mind*. New York: Oxford University Press.
- Stewart, Neil, Chater, Nick, & Brown, Gordon D.A. 2006. Decision by sampling. *Cognitive Psychology*, **53**, 1–26.
- Stewart, Neil, Reimers, Stian, & Harris, Adam J.L. in press. On the origin of utility, weighting, and discounting functions: How they get their shapes and how to change their shapes. *Management Science*.
- Sunstein, Cass R. 2002. Probability neglect: Emotions, worst cases, and law. *The Yale Law Journal*, **112**(1), 61–107.
- Tversky, Amos, & Koehler, Derek J. 1994. Support Theory: A nonextensional representation of subjective probability. *Psychological Review*, **101**, 547–567.
- Weatherson, Brian. 2005. Can we do without pragmatic encroachment? *Philosophical Perspectives*, **19**, 417–43.
- Weatherson, Brian. 2012. Knowledge, bets, and interests. *Pages 75–103 of: Brown, Jessica, & Gerken, Mikkel (eds), *Knowledge ascriptions*. Oxford University Press.*
- Wedgwood, Ralph. 2012. Outright belief. *Dialectica*, **66**(3), 309–329.
- Whittle, Bruno. 2012. Belief, information, and reasoning. *Philosophical Perspectives*, **26**, 431–46.
- Williamson, Timothy. 2000. *Knowledge and its limits*. New York: Oxford University Press.