UNCERTAINTY AND UNDERMINING

James Pryor
NYU
19 February 2007

1. Let's understand **dogmatism about perception** to be the view that we sometimes have immediate but defeasible justification to believe things are as they perceptually appear to us. That is, we have some justification that doesn't "come from" our justification to believe other premises. When one has any justification of that sort, I'll say one has **dogmatic support** for the relevant beliefs.

Dogmatism is a claim about a possible *epistemic position*, not about the *metaphysics of what puts us* in that position. So, for example, it leaves it open whether the intrinsic nature of a perceiving subject's state is the same as that of a hallucinating subject's state.

Further, dogmatism claims only that *there is* an epistemic position of the sort described, that we *sometimes occupy*. It leaves it open *exactly which* subjects are in this epistemic position. Is the dogmatic support available only to those whose perceptual appearances are reliable? Is it available only when one's perceptual appearances constitute genuine perceptions? Or is it, rather, available only when one's appearances *don't* constitute genuine perceptions—whereas subjects who *are* perceiving are in *better* epistemic positions? I have views about these questions, but I've defined dogmatism so as to leave them all open.

Similarly, I've left it open what conceptual resources are needed, for one to have dogmatic support to believe anything.

2. In my earlier defenses of dogmatism, I operated in terms of having or lacking justification for all-out belief. I did characterize dogmatic support as *prima facie*, able to be opposed or undermined by other evidence. The idea of opposing already suggests that justification comes in degrees. But I never said how dogmatism should be understood when we're working explicitly with degrees of justification, or justification for degrees of belief. Ultimately, I do think we should be prepared to work with those notions.

Are degrees of justification and justification for degrees of belief the same? It's not obvious that they should be. And even if they are the same, it's not obvious that what they come to should be the standard probabilistic story about degrees of belief that Bayesians employ. However, that is a reasonable place for us to start. If we need to alter the standard picture, we can discover that as we proceed.

Now, Bayesians construe updating as a matter of becoming certain that some hypothesis is true. And dogmatists will insist that perceptual updating *needn't* involve becoming certain of anything. One might not have any *opinion about* what experiences one is having. Instead of updating on *hypotheses about* one's experience, the dogmatists will say we should update on the mere *having of* the experiences.

I'm sympathetic to this complaint, but perhaps we can work around it. Let's restrict our attention to subjects who both have, and are conscious of having, experiences of the sort that are alleged to provide dogmatic support. These reflective subjects are presumably no *worse off* than subjects who have no opinions about their experiences. So it will be worthwhile to consider how the dogmatist will model their epistemic state.

For convenience, I'll refer to experiences that provide dogmatic support as just "experiences." I hope, though, to leave the questions we started with open. For all the definitions say, there may be phenomenological duplicates of subjects with dogmatic support, who don't themselves have dogmatic support. For expository convenience, I exclude their states from what I'll refer to with "experience."

Even when subjects *are* aware of what experiences they're having, the dogmatist will still be uncomfortable with a Bayesian conception of perceptual updating. The Bayesian story invites understanding one's confidence about the world as only mediately justified, that is, supported in part by *premises about* what experiences one is having. Claims about how mediate one's justification is aren't part of the official Bayesian story, but as we'll see, this is what the official story most naturally models. We'll have to do some work to develop a model where one's confidence about the world can be immediately justified, though still defeasible.

3. What hypotheses *do* get dogmatic support? In my earlier work, I said that experiences as of P gave one immediate but defeasible justification to believe P. In the formal discussion we'll have below, I'll have experiences in the first place providing dogmatic support to the hypothesis that you're, by having those experiences, thereby perceiving something to be the case. The effect on hypotheses about how things simply *are* in the world, regardless of whether you're perceiving them, will be epistemically downstream. It's natural to think that being an "epistemically downstream effect" amounts to the same thing as being a "mediately justified consequence." If so, then the view I develop here will be that hypotheses about *how things simply are in the world* are only mediately justified. What gets immediate support are hypotheses about *how you are, by having certain experiences, now perceiving things to be*.

This may still fit the letter of my earlier work. One way for it to do so would be if the contents of experiences were self-referential claims about what those experiences constitute perceptions of. That is, if the P such that the experience is as of P is: *by having this experience, I'm now perceiving Q to be the case*. That P is the kind of proposition our formal treatment will allow to be dogmatically supported. Alternatively, one might try to pull apart the notions of being an "epistemically downstream effect" and being a "mediately justified consequence." I will not attempt here to settle this.

I trust that, even if the view I develop below violates the letter of my earlier proposals, it's still true to their spirit. Hypotheses about what one's perceiving to be so are also claims about the world.[1]

---

[1] Remember that we're talking explicitly about subjects who are *aware of* the experiential source of the dogmatic support they're acquiring. I'm modeling *them* as being such that their epistemically most upstream uptake is that they're now perceiving Q to be the case. I don't know how to model subjects who get justification from their experiences *without* being aware that they're doing so. But I think there are many such subjects.

I'll say that a hypothesis U **undermines** the dogmatic support an experience gives you to believe you're perceiving Q iff U entails that *you have* that experience but *aren't* perceiving Q. (In Pryor 2004, I called these "non-perceiving hypotheses.") Hypotheses whose truth would merely make it to some degree *likely* that you're having the experience, but not thereby perceiving, don't count as underminers in this technical sense—though evidence for them may manage to raise the likelihood of U, and thereby have epistemically downstream undermining effects.

4. Let PERC be the hypothesis that you are, *by* having some experiences η that dogmatically support PERC, now perceiving that Q. Let E be the hypothesis that you are having—or will have, I will ignore issues of tense—those experiences η. Though the experiences dogmatically support the claim that you're perceiving, they're also compatible with your not perceiving. Let U be a hypothesis that says you have η but aren't thereby perceiving. So U is an underminer for the support that η gives you to believe PERC. Finally, suppose you're not yet certain you'll have η.

Now standard probability theory tells us that:

(Theorem 1) $p(U|E) > p(U)$.[2]

Since PERC entails not-U, another theorem of standard probability theory is that $p(\text{PERC}|E) \leq p(\text{not-U}|E)$. Since Theorem 1 is equivalent to $p(\text{not-U}|E) < p(\text{not-U})$, it follows that:

(Theorem 2) $p(\text{PERC}|E) < p(\text{not-U})$.

These two theorems relate static quantities. The Bayesian understands them to have dynamic upshots: first, updating on evidence E should make one more confident that U, not less; and second, one's posterior confidence in PERC, after updating on evidence E, should never exceed one's prior confidence in not-U.

This seems to pose two challenges to the dogmatist. We need to proceed cautiously, though, since it's not yet clear what the probabilistic commitments of dogmatism are.

The first challenge arises like this. The dogmatist says that having η tends to support PERC in some way that it doesn't also tend to support U. It has some kind of *bias* towards PERC. Since PERC is incompatible with U, that suggests that updating on the hypothesis that you've had η should tend to justify you in believing *not*-U. (And indeed, I've argued that, in the right conditions, an experience as of hands *will* give you justification to believe you're not a handless brain in a vat; see Pryor 2004.) However, Theorem 1 says that $p(U|E)$ is *higher* that $p(U)$. If we accept the Bayesian construal of "justifying" as "probability raising," it follows that E justifies *U* rather than not-U. One of these views must be wrong. Roger White puts the Bayesian viewpoint like this:

Neither do I know how to model subjects who are aware of their experiences but who have false beliefs about their epistemological effect. But as I discuss in Pryor 2004 and Pryor Wright festschrift, I think there are subjects of this sort, too.

[2] Proof: $p(U|E) = p(E|U)p(U)/p(E)$. Since U entails E, $p(E|U) = 1$. Since you're not yet certain of E, $p(E)$ will be < 1. It follows that $p(U|E)>p(U)$.

> Dogmatism has the consequence that when it appears to me that there is a hand before me, I can gain justification, perhaps for the first time, for believing that it is not a fake-hand, that I am not a brain in a vat, and so on. Now if I gain justification for a hypothesis, then my confidence in its truth should increase. But arguably when it appears to me that something is a hand, my confidence that it is not a fake hand should *decrease*. For since this is just what a fake-hand would look like, the degree to which I suspect it *is* a fake should *increase*.[3]

The second challenge goes like this. The dogmatist says that η can justify you in believing PERC even if you're not antecedently justified in believing not-U. But Theorem 2 puts an upper cap on the amount of justification you can get for PERC. It says that the hypotheses that you've had η can only raise your justification to believe PERC to a given level if you're *already* justified above that level in believing not-U. Here again is White:

> …p(Q|E) < p(not-U). So its appearing to me that this is a hand can render me justifiably confident that it is a hand, only if I am already confident that it is not a fake-hand.[4]

These two challenges make the following assumptions:

(a) Degrees of belief should conform to standard probability theory, and so should validate Theorem 1 and Theorem 2.
(b) The justificatory effects of having η, at least for a subject who's aware of and certain that she has η, should be the same as Bayesian conditionalizing on the hypothesis that she has η.
(c) Updating on E cannot *justify* a hypothesis if its effect is to *lower* that hypothesis's probability.
(d) If E can justify PERC to a given level only when you're *already* justified above that level in believing not-U, then your justification to believe PERC *does* require antecedent justification to believe not-U.

I have doubts about all four of these assumptions.

The doubts about (a) will emerge below (in §§5–7, and §§15–18).

I already voiced some uncertainty about (b) in §2. In the formal model I develop, it will be useful to press dogmatic support into the mold of updating on a hypothesis; but as we'll see, the updating will be best understood to be non-Bayesian.

I won't pursue my doubts about (c) here, except to point out that (c) is incompatible with the intuitive thought that any justification to believe P&Q should count as *some* justification to believe Q. For we can raise the probability of P&Q while

---

[3] White 2006, 531.

[4] White 2006, 534. I've substituted my hypothesis-names for White's. His passage concerns the probability of Q, rather than the probability of PERC (=I am now perceiving that Q). But the claims he makes apply to both hypotheses.

See also Schiffer 2004, Silins ??, and Wright forthcoming.

lowering the probability of Q. There's a clash of intuitions here, and I think they need to negotiate more with each other.[5]

As to (d), the dogmatist's central thesis concerns *justificatory dependence*. By contrast, Theorem 2 is about an *upper bound* on how much your experience can justify you in believing PERC. It's not clear how those issues are connected. The fact that we use "already" in the antecedent of (d) and "antecedent" in the consequence seems to me to invite more confusion than illumination. It'd be wrong to *identify* the epistemic phenomenon the dogmatist is arguing for with a violation of Theorem 2. Dogmatic support is supposed to be present to some degree even in the perceptual updating of subjects who *do* have excellent evidence that not-U, and so don't violate Theorem 2.

White understands the dogmatist as at least committed to the *possibility* of violating Theorem 2; and this will turn out to be correct. But that commitment isn't easily read off of our present formulations of dogmatism. It will take a lot of work to unearth it. And I hope that work will simultaneously have the effect of making (a) and (b) look less compulsory.

5. It's difficult to know what the probabilistic commitments of dogmatism are because dogmatism is stated and motivated using vocabulary that has no easy translation into the standard probabilistic model. The two most important pieces of vocabulary are the notions of agnosticism and of undermining defeat.

By **agnosticism** I mean the kind of doxastic attitude that's an appropriate response to a lack of evidence. This is different than **epistemic indifference**, which can be appropriate even when one has a great deal of evidence on either side. Epistemologists should be careful not to conflate:

(i) subjects who have no doxastic attitude at all towards an undermining hypothesis U, e.g. because they can't entertain U;

(ii) subjects who do entertain U, but are wholly uninformed about it, and so agnostic; and

(iii) subjects who think U is as likely as not, in response to a great deal of evidence on either side that balances out that way, or in response to knowledge that U has a 50% objective chance.

I've discussed subjects in situation (i) elsewhere.[6] In the present discussion I will set them aside. We'll assume that our subjects always have *some* doxastic attitude towards the undermining hypotheses we consider. That leaves subjects who are in situation (ii), or situation (iii), or at points intermediate between them.

As I've developed dogmatism, it makes very different proposals about situation (ii) and situation (iii). It says that the immediate justification your experiences give you *isn't* defeated or undermined by the mere epistemic possibility that U—by U's failing to be conclusively ruled out by your evidence. But *evidence* that you're in U does undermine. This suggests that any confidence you have in U that's *not* based in evidence should have no undermining effect. So the confidence you have in U *prior* to acquiring any evidence must interact differently with your perceptual justification than the

---

[5] Part of the problem is that the Bayesian is thinking about all things considered justification, whereas the intuitive thought need not be.

[6] Pryor 2004, Pryor Wright festschrift.

confidence you base on evidence. We need to keep track of how much of your confidence is a response to evidence and how much is a leap in the evidential dark. The standard probabilistic framework gives us no guarantee of being able to untangle those.

In his critique of dogmatism, White describes a game where three cards are labeled “Leave his hands alone,” “Cut off his hands,” and “Cut off his hands and replace them with plastic replicas.” Your captors anesthetize your arms and put you to sleep, then draw a card randomly and carry out its instructions. You awaken and seem to see hands. White claims that at this point your probability that you do still have hands should be ½. I’m disposed to agree with him, because this sounds like a case where you know that the objective chance of your having hands, given that you seem to see hands, is ½. This knowledge comes from your evidence about the situation. Dogmatism only gives a contrary verdict when subjects *don’t* know what the objective chances are, and their views about how likely their experiences are to be perceptions are based on incomplete evidence. White says:

> The only apparent difference between the card game and an ordinary case of judging whether someone has a hand is the following. How we ought to distribute our credence among the possibilities seems more straightforward in the card game… In a regular case it is not so clear what the relative plausibility of these hypotheses is. But this does not appear to make any difference to how our convictions should be altered in the light of experience.[7]

However, this last claim is exactly what dogmatism as I’ve developed it denies. In fact, it was my *guiding thought* in developing dogmatism that confidence based on evidence and confidence (or any doxastic attitude) not so based should be epistemically different.

So we need to settle on some formal representation of this difference, even to be able to *state* the distinctive claims of dogmatism.

6. Keynes drew something like the distinction I’m drawing. He used “risk” to designate evidentially-based confidence that falls short of certain knowledge. He used “uncertainty” to designate something like *lacking an evidential basis* for confidence. He explains this latter notion like this:

> By ‘uncertain’ knowledge, let me explain, I do not mean merely to distinguish what is known for certain from what is only probable. The game of roulette is not subject, in this sense, to uncertainty; nor is the prospect of a Victory bond being drawn. Or, again, the expectation of life is only slightly uncertain. Even the weather is only moderately uncertain. The sense in which I am using the term is that in which the prospect of a European war is uncertain, or the price of copper and the rate of interest twenty years hence, or the obsolescence of a new invention, or the position of private wealth owners in the social system in 1970. About these matters there is no scientific basis on which to form any calculable probability whatever. We simply do not know.[8]

---

7 White 2006, 536.

8 Keynes, “The General Theory of Employment,” Quarterly Journal of Economics 51 (1937) 209-23. Reprinted in *The Collected Writings of John Maynard Keynes*, ed. DE

One common refinement to Bayesianism is to model doxastic states with *sets* of probability functions, rather than single probability functions. (Commonly, they're required to be convex sets.) Then one's credence in a proposition can be more or less *spread out*, as well as being more or less *high*. That may help us model the kinds of evidentially ungrounded agnosticism that Keynes describes. However, we should contrast *that* use of sets of probability functions from their use to model *imprecision* in one's confidence. I don't think they can play both of these explanatory roles at the same time.

7. The second important notion for the dogmatist is the notion of undermining defeat. Dogmatists don't just think that immediate perceptual justification can be defeated; they think it can be **undermined**. Intuitively, that's the kind of defeat that says, not that your environment is some way *other* than it appears, but rather that you're not in a position to *perceptually tell* what your environment is like.

This notion has been formulated under different names by a variety of theorists in a variety of contexts.[9] It's difficult to explain what it amounts to in Bayesian terms. I don't mean that the Bayesian can't represent the defeating that goes on in paradigm cases of undermining; but rather that it's difficult for him to identify *what makes them* undermining. The Bayesian will regard *any* kind of defeating evidence as increasing the confidence you should assign to your environment's being other than it appears. Moreover, *any* evidence that makes it more likely that things are other than they appear should increase the likelihood that your earlier evidence is unreliable. So it seems like

---

Moggridge, (London: Macmillan), vol 14, pp. 109-23, at pp. 114-15? Thanks to Brian Weatherson for bringing this article to my attention.

9 The earliest sustained discussion I know is John Pollock, *Knowledge and Justification* (Princeton Univ. Press, Princeton, 1974), Ch. 2 and 5. But I think it was already then common to distinguish between skeptical scenarios that are incompatible with the *truth* of what you purport to know, and scenarios that are merely incompatible with your knowing it. The notion of undermining defeat may already be implicit in the attempt to assign that contrast any significance.

Here's a sampling of other sources: John Pollock, "Reliability and justified belief," *CJP* 14 (1984); John Pollock and Joseph Cruz, *Contemporary Theories of Knowledge*, 2nd ed. (Rowman & Littlefield, Totowa, NJ, 1999), pp. 195-97; William Alston, "An internalist externalism," *Synthese* 74 (1988), reprinted in *Epistemic Justification* (Cornell Univ. Press, Ithaca, 1989), 237-45, at p. 238; Crispin Wright, "Strict finitism," *Synthese* 51 (1982), reprinted in *Realism, Meaning and Truth*, 2nd ed. (Blackwell, Oxford, 1993), 107-75, at pp. 117-20; Crispin Wright, "Rule-following, meaning and constructivism," in Charles Travis, ed. *Meaning and Interpretation* (Blackwell, Oxford, 1986), 271-97, at pp. 290-1; Crispin Wright, "Skepticism and dreaming: imploding the demon," *Mind* 100 (1991), 87-115, at pp. 94-95; Stewart Cohen, "Knowledge, context, and social standards," *Synthese* 73 (1987); Alvin Goldman, *Epistemology and Cognition* (Harvard Univ. Press, Cambridge, Mass., 1986), Ch. 4; and Robert Audi, *The Structure of Justification* (Cambridge Univ. Press, Cambridge, 1993), pp. 142-44.

he'd regard defeating evidence as on a par: it *all* tends to confirm opposing alternatives, and it *all* tends to confirm that your initial justification was unreliable. There may be more subtle ways for a Bayesian to capture the intuitive notion of undermining. (One of my students, Matt Kotzen, has a proposal.) But it will take some care; especially if we want to model the facts that some evidence *undermines more* than others, and that much evidence is *a mix* of undermining and opposing components.

I do think the Bayesian will be able to give *some* account of the intuitive notion of undermining. He may not be able to model *everything* we want to say informally with this notion, but I expect he'll be able to go some distance. What primarily worries me is not *how much* the Bayesian will be able to say, but rather *what* he will say. I expect him to tell us that undermining always works like this:

> When you get evidence E, you also get defeasible support to believe Q, because of your prior confidence in some hypothesis about how E and Q are reliably linked. Undermining the support E gives you to believe Q is just getting evidence that should lower your confidence that E and Q are so linked.

Now, no doubt we often are in epistemic positions of that sort. But what this sounds like is a case of undermining justification for Q that was *mediated* by a belief about how E and Q are linked. It's central to the dogmatist's thinking that undermining needn't always be mediated in that way. The dogmatist says you can undermine justification for Q even when that justification is *immediate*. Perhaps that will have the *consequence*, if you're aware that E describes the source of your justification, that you should be less confident that E and Q are linked. But it seems wrong to identify your weaker confidence in a linking belief as *the mechanism* through which your justification for Q is eroded. It should be possible to *directly undermine* the immediate justification you have for Q, without opposing any premise that supports Q for you.

The Bayesian doesn't say this is impossible. Officially, he doesn't say *anything* about relations of mediacy or justificatory dependence. But neither does his formalism offer anything that looks like an intuitive model of direct undermining. This is another reason why it's hard to see how exactly to translate the dogmatist's distinctive claims into the standard Bayesian picture.

8. Let's make a fresh start. I propose to model epistemic states as a combination of two parts: a mass distribution, which changes as the subject acquires new evidence, and a range of inheritance plans, which does not change.

What is a mass distribution? Let W be a finite space of hypotheses, and call W's most determinate epistemically possible hypotheses its "atoms."[10] For illustration, let's

---

[10] Formally, W is a finite σ-algebra on some partition Ω of epistemic possibilities (that is, a finite set of Ω's subsets including Ω and closed under complementation and union). W's atoms are its minimal non-empty elements. I will designate non-atomic elements of W like {a,b} as $a \vee b$.

I do not know whether there are obstacles to extending the formal model I set out below to infinite hypothesis spaces.

say W has four atoms: a, b, c, and d. Since these are W's most determinate hypotheses, they are jointly incompatible. W will contain 2^4 hypotheses altogether:

a∨b∨c∨d					
a∨b∨c	a∨b∨d	a∨c∨d	b∨c∨d		
a∨b	a∨c	b∨c	a∨d	b∨d	c∨d
a	b	c	d		
∅					

That's as fine-grained as W represents the world. Hypothesis a∨b∨c∨d is represented by W to be epistemically certain; let's designate this hypothesis ⊤. ∅ is epistemically impossible; let's designate it ⊥.

I'll say that i is a subhypothesis of h when i ⊆ h. For example, a and b and a∨b are each subhypotheses of a∨b.

A mass distribution on hypothesis space W is an assignment of reals to each hypothesis in W, such that: (i) ⊥ is assigned 0, (ii) every other hypothesis is assigned a real ≥ 0, and (iii) all the assignments sum to 1. For example, here is one mass distribution:

		⊤ $^{1}/_{3}$			
a∨b∨c 0	**a∨b∨d** 0	**a∨c∨d** 0	**b∨c∨d** 0		
a∨b $^{1}/_{3}$	**a∨c** 0	**b∨c** 0	**a∨d** 0	**b∨d** 0	**c∨d** 0
a $^{1}/_{6}$	**b** $^{1}/_{6}$	**c** 0	**d** 0		
		⊥ 0			

The amount a distribution assigns to a hypothesis is that hypothesis's **specific mass**, or **mass** for short. In the example above, a∨b has a specific mass of $^{1}/_{3}$. I'll write that as:

$m(a \vee b) = {}^{1}/_{3}$.

A hypothesis's **cumulative mass** is the sum of the specific masses of all its subhypotheses. In our example, a∨b has a cumulative mass of $^{1}/_{3} + {}^{1}/_{6} + {}^{1}/_{6} = {}^{2}/_{3}$. We won't need to talk about these much; specific masses are more fundamental to what we'll be doing.

Intuitively, what h's specific mass represents is the amount of *evidentially-based* confidence one's epistemic state recommends specifically in h, and not in any of h's more determinate subhypotheses. So a **pre-evidential epistemic state**—one prior to the effects of any evidence—will have all its mass assigned to hypothesis ⊤. Someone in that state would have no *evidential* basis for further dividing her confidence. In the example above, on the other hand, you've acquired some evidence with a net effect of assigning $^1/_3$ of your confidence to a∨b's being true; and you lack any evidential basis for further dividing that confidence between a and b.

To say you have no *evidential* basis for further dividing your confidence isn't to say that you have no basis at all. This is where inheritance plans come in. An **inheritance plan** is a complete specification of how to divide your confidence in the absence of evidence. It will tell you how the mass that's assigned higher up in a distribution should be "inherited by" more determinate subhypotheses. We'll explain how questions of inheritance arise shortly. For now, trust that we sometimes do need to know how to divide up confidence that a mass distribution assigns to non-atoms. One inheritance plan might tell you to split ⊤'s mass evenly between a, b, and c∨d; to give 90% of a∨b's mass to a, and the rest to b; and so on. An inheritance plan settles all such questions.

In the usual case, a subject won't have *one fully determinate* plan for how to divide her confidence in the absence of evidence. This is why I represent epistemic states as a combination of a mass distribution and *a range of* inheritance plans.

9. Inheritance plans participate in three important operations on an epistemic state.

One of these operations is *conditionalizing on a supposed hypothesis*. I understand this to be deriving how a subject who's in the epistemic state should be doxastically inclined when reasoning under the given supposition. For example, a subject may think c most likely to be true, but be prepared to reason hypothetically on the supposition that a∨b∨d. When adopting a supposition, subjects may find that their evidence still only gives a basis for assigning *some* of their confidence; so we should distinguish between mass assigned to atoms and mass not so assigned in the conditional case, too. I'll write m(a∨b | a∨b∨d)=$^4/_5$ to mean that the subject's specific mass for a∨b, conditional on the supposition that a∨b∨d, is $^4/_5$. I'll use m(•) and m(• | ⊤) interchangeably to designate unconditional masses.

A second operation on an epistemic state is *updating on evidence*. This has internal connections to conditionalizing, but the connections are complex and it's best to regard these as different operations. Conditionalizing is a matter of *reading off* what your *current* epistemic state says about things, against some suppositional background. Updating is a matter of *changing* your epistemic state in response to evidence.

Think of the hypotheses in our diagram as having buckets attached. Let the specific mass you have assigned to each hypothesis be a quantity of sand in that hypothesis's bucket. When you update on evidence, you'll be gaining confidence that some hypothesis is true. We'll have to plug into our formalism a value indicating *which* hypothesis you're most directly gaining confidence in. Other hypotheses may get more credible, too, in response. But the effect on them will be "downstream." We need there to be one hypothesis that your evidence most directly supports. We'll also need to plug a second value into our formalism, indicating *how much* more confident you should become in that hypothesis.

Let's walk through an example. Suppose you get evidence of strength *J* for hypothesis a∨b. *J* will translate into a quantity θ of new sand you get to pour into the buckets of a∨b's subhypotheses: a, b, and a∨b. You divide this new sand among those subhypotheses in such a way that the conditional masses m(a | a∨b), m(b | a∨b), and m(a∨b | a∨b) don't change. Then you renormalize so that the total amount of sand in your distribution again adds up to 1. In other words:

$m'(\bot)$ stays 0			
$m'(a)$	$= (m(a)$	$+ \theta\, m(a \mid a\vee b)$	$) / (1+\theta)$
$m'(b)$	$= (m(b)$	$+ \theta\, m(b \mid a\vee b)$	$) / (1+\theta)$
$m'(a\vee b)$	$= (m(a\vee b)$	$+ \theta\, m(a\vee b \mid a\vee b)$	$) / (1+\theta)$
$m'(c)$	$= (m(c)$		$) / (1+\theta)$
$m'(a\vee c)$	$= (m(a\vee c)$		$) / (1+\theta)$
$m'(b\vee c)$	$= (m(b\vee c)$		$) / (1+\theta)$
etc.			

How to calculate the conditional masses m(• | a∨b), and what relation there is between *J* and θ, are matters we have yet to settle. But our intuitive picture should already motivate the idea that as the evidence you're updating on gets better and better, θ will get larger and larger. The limit as $\theta \rightarrow \infty$ will represent getting evidence that makes a∨b epistemically certain. In that case:

$$m'(a) \;=\; \lim_{\theta \rightarrow \infty} \left[\left(m(a) + \theta\, m(a \mid a\vee b) \right) / (1+\theta) \right] = m(a \mid a\vee b)$$

This is an intuitively natural result. It means that your conditional mass in a, given the *supposition* that a∨b, is the confidence you should end up having in a were you to acquire *certain evidence* that a∨b is true.

I've described the operations of conditionalizing and updating. A third operation on epistemic states is *flattening*. A mass distribution plus a single inheritance plan will determine (though usually not be determined by) a probability function on the hypothesis space. Flattening is the operation of reading off what that probability function is. Since subjects usually have *a range* of inheritance plans, their epistemic states will determine a corresponding range of probability functions. Those ranges of probability functions will turn out to be too coarse-grained, by themselves, to do all the work we want. Epistemic states will sometimes need the extra structure I've given them. But for other purposes we'll only need the probability functions that are the "flattenings" of a state.

I've described three operations: conditionalizing, updating, and flattening. I'll set out the details of how these work shortly. First I want to describe an intuitive constraint that I take to govern them. This is the constraint that these operations be "path-independent."

One way for that constraint to operate is this: if you were to conditionalize an epistemic state on the hypothesis a∨b∨c, and then conditionalized the result on the hypothesis a∨b, you should get the same result as if you had conditionalized the state directly on a∨b in the first place.

A second way for the constraint to operate is that evidence should commute: if you get evidence of strength *J1* directly supporting hypothesis h1, and then evidence of strength *J2* supporting hypothesis h2, you should end up with the same result as if you had acquired the evidence in the other order.

A third way for the constraint to operate is that updating and conditionalizing should be commutable: if you conditionalize a state on the hypothesis a∨c, and then update the result on evidence of strength *J* for a, you should get the same result as if you had first updated the unconditionalized state on evidence of strength *J* for a (or for any hypothesis whose intersection with a∨c is a), and then conditionalized on a∨c.

These constraints are perhaps negotiable; if it turned out that we couldn't meet them all, then perhaps we could reconcile ourselves to that outcome. But I take it to be a strength in a formal model that it does meet them. And the model I'm setting out will manage to do so.

10. Our next step is to think harder about how to calculate conditional masses. I'll describe three different approaches we might take to that. They correspond to three different structures that inheritance plans might have. This range of *theoretical options* about inheritance plans should not be confused with a subject's *own range* of inheritance plans. I think that *at most one* of these theories about the structure of inheritance plans can be correct; and *all* of a subject's inheritance plans must have the structure that the correct theory states.

A further complication is that we'll need to talk separately about "classical" cases, where there's no dogmatic support in one's evidence, and cases where there is a dogmatic support. In the following sections, I'll describe how things go when everything is classical. Later we'll consider how things change in the context of dogmatic support.

Let S be a subject whose opinion conforms to her epistemic state. Suppose we want to calculate S's *conditional* opinion, relative to the hypothesis that a∨b∨c. We'll need to settle to what extent the mass S has assigned to light grey cells can be "inherited by" darker grey cells:

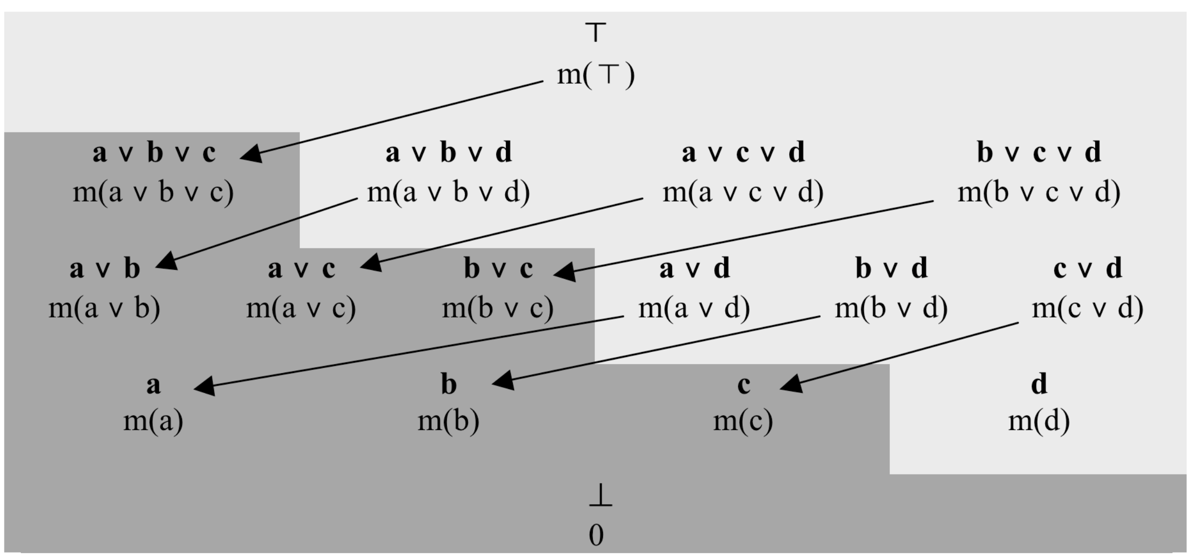

One natural proposal is that this confidence should be inherited along the paths where I've drawn arrows. That is, when conditionalizing on the hypothesis a∨b∨c, the hypothesis a gets the benefit of *all* the mass that had formerly been assigned to a∨d; the hypothesis a∨b∨c gets *all* the mass that had formerly been assigned to ⊤; and so on. At the end we renormalize. I'll call this a Dempster-Shafer inheritance structure.

If that's the right structure for inheritance to have, then S's conditional masses m(• | avbvc) should look like this:

$$
\begin{aligned}
&m(\bot \mid a{\vee}b{\vee}c) &&= 0 \\
&m(a \mid a{\vee}b{\vee}c) &&= (\, m(a) + m(a{\vee}d) \,) / \sigma \\
&m(b \mid a{\vee}b{\vee}c) &&= (\, m(b) + m(b{\vee}d) \,) / \sigma \\
&m(a{\vee}b \mid a{\vee}b{\vee}c) &&= (\, m(a{\vee}b) + m(a{\vee}b{\vee}d) \,) / \sigma \\
&m(c \mid a{\vee}b{\vee}c) &&= (\, m(c) + m(c{\vee}d) \,) / \sigma \\
&m(a{\vee}c \mid a{\vee}b{\vee}c) &&= (\, m(a{\vee}c) + m(a{\vee}c{\vee}d) \,) / \sigma \\
&m(b{\vee}c \mid a{\vee}b{\vee}c) &&= (\, m(b{\vee}c) + m(b{\vee}c{\vee}d) \,) / \sigma \\
&m(a{\vee}b{\vee}c \mid a{\vee}b{\vee}c) &&= (\, m(a{\vee}b{\vee}c) + m(\top) \,) / \sigma
\end{aligned}
$$

where σ is a renormalizing factor (= 1-m(d)).

This Dempster-Shafer inheritance structure has no free parameters; so if it's the right way for inheritance to go, then subjects will only ever need, or have, a single inheritance plan, rather than a range of different inheritance plans.

A second proposal retains the idea that confidence is inherited along the paths where I've drawn arrows. But on this proposal, subhypotheses only inherit *a fraction* of the confidence assigned to their superhypotheses. Let α be a probability function on our space of hypotheses. Then S's conditional masses might look like this:

$$
\begin{aligned}
&m(\bot \mid a{\vee}b{\vee}c) &&= 0 \\
&m(a \mid a{\vee}b{\vee}c) &&= (\, m(a) + {}^{\alpha(a)}/_{\alpha(a\vee d)}\; m(a{\vee}d) \,) / \sigma' \\
&m(b \mid a{\vee}b{\vee}c) &&= (\, m(b) + {}^{\alpha(b)}/_{\alpha(b\vee d)}\; m(b{\vee}d) \,) / \sigma' \\
&m(a{\vee}b \mid a{\vee}b{\vee}c) &&= (\, m(a{\vee}b) + {}^{\alpha(a\vee b)}/_{\alpha(a\vee b\vee d)}\; m(a{\vee}b{\vee}d) \,) / \sigma' \\
&m(c \mid a{\vee}b{\vee}c) &&= (\, m(c) + {}^{\alpha(c)}/_{\alpha(c\vee d)}\; m(c{\vee}d) \,) / \sigma' \\
&m(a{\vee}c \mid a{\vee}b{\vee}c) &&= (\, m(a{\vee}c) + {}^{\alpha(a\vee c)}/_{\alpha(a\vee c\vee d)}\; m(a{\vee}c{\vee}d) \,) / \sigma' \\
&m(b{\vee}c \mid a{\vee}b{\vee}c) &&= (\, m(b{\vee}c) + {}^{\alpha(b\vee c)}/_{\alpha(b\vee c\vee d)}\; m(b{\vee}c{\vee}d) \,) / \sigma' \\
&m(a{\vee}b{\vee}c \mid a{\vee}b{\vee}c) &&= (\, m(a{\vee}b{\vee}c) + {}^{\alpha(a\vee b\vee c)}/_{\alpha(\top)}\; m(\top) \,) / \sigma'
\end{aligned}
$$

where σ′ is another renormalizing factor. Here the mass assigned to avbvd gets divided into two parts: one part, sized ${}^{\alpha(a\vee b)}/_{\alpha(a\vee b\vee d)}$, going to the intersection of avbvd and the hypothesis you're conditionalizing on, and the rest going to the complement:

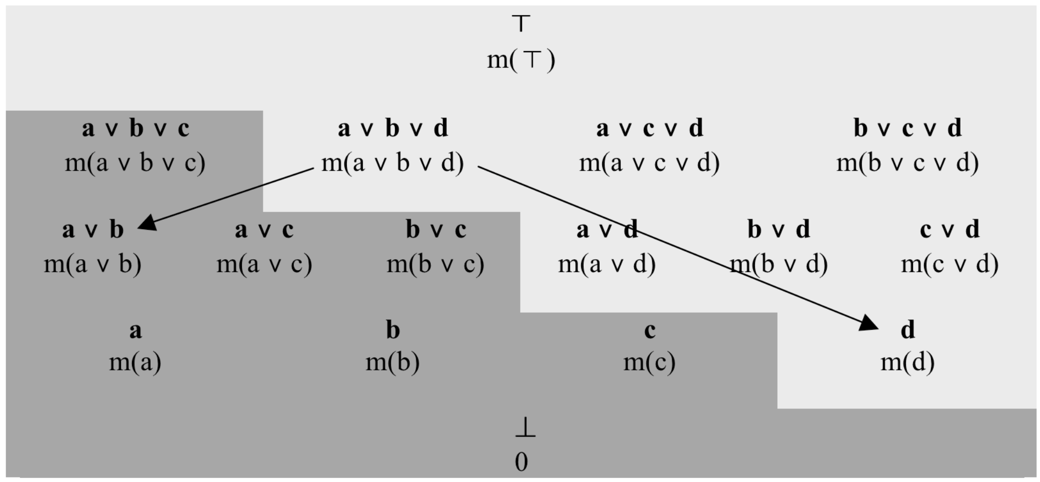

When we're calculating the subject's conditional masses m(• | a∨b∨c), we only include that part of a∨b∨d's mass that the dark grey a∨b∨c region has inherited.

I'll call this an α-based inheritance structure. You get a different inheritance plan for each probability function α. So a subject's range of inheritance plans could be thought of as a set of α functions.

A third proposal works somewhat differently. Here the mass that S assigned to a∨b∨d gets divided into *six* parts, one for each of that hypothesis's more determinate possible subhypotheses. So that mass gets inherited along each of the solid paths in this diagram:

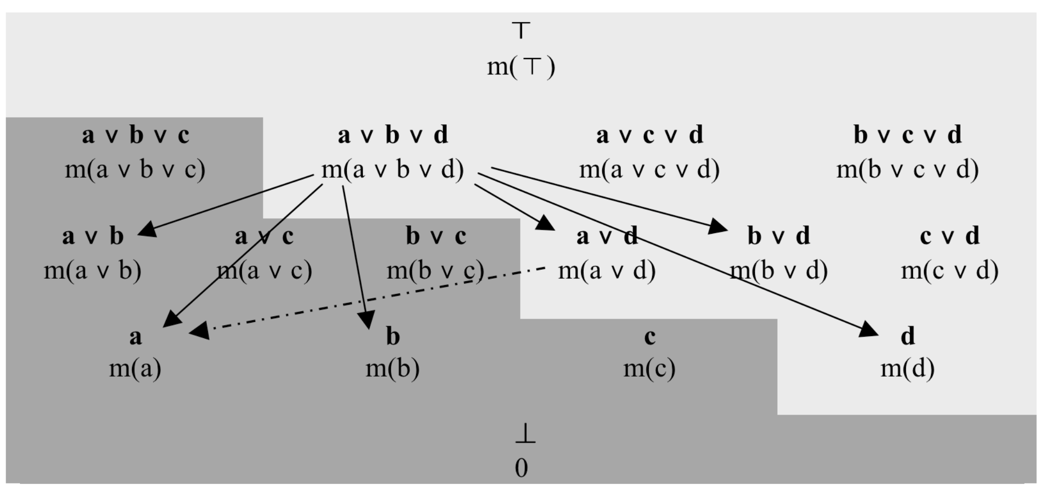

I'll write γ(a∨b, a∨b∨d) to indicate the portion of a∨b∨d's specific mass that should get inherited by a∨b. The only constraints this third proposal imposes on the values:

γ(a, a∨b∨d)
γ(b, a∨b∨d)
γ(a∨b, a∨b∨d)

$\gamma(d, \quad a \vee b \vee d)$
$\gamma(a \vee d, \; a \vee b \vee d)$
$\gamma(b \vee d, \; a \vee b \vee d)$

is that they're each ≥ 0, and add up to 1. This proposal should be understood recursively: when hypothesis a inherits from a∨d (along the dotted path in the diagram), it should *also* inherit any mass that a∨d itself inherited from its superhypotheses. Spelled out, this proposal says S's conditional masses should look like this:

```
m(a | a∨b∨c) = (
      m(a)
    + γ(a, ⊤)* m(⊤)
    + γ(a, a∨b∨d) * ( m(a∨b∨d ) + γ(a∨b∨d, ⊤) * m(⊤) )
    + γ(a, a∨c∨d) * ( m(a∨c∨d ) + γ(a∨c∨d, ⊤) * m(⊤) )
    + γ(a, a∨d) * (
            m(a∨d)
          + γ(a∨d, ⊤) * m(⊤)
          + γ(a∨d, a∨b∨d) * ( m(a∨b∨d ) + γ(a∨b∨d, ⊤) * m(⊤) )
          + γ(a∨d, a∨c∨d) * ( m(a∨c∨d ) + γ(a∨c∨d, ⊤) * m(⊤) )
          )
) / σ″

m(a∨b | a∨b∨c) = (
      m(a∨b)
    + γ(a∨b, ⊤)* m(⊤)
    + γ(a∨b, a∨b∨d) * ( m(a∨b∨d ) + γ(a∨b∨d, ⊤) * m(⊤) )
) / σ″

etc.
```

where σ″ is another normalizing factor. On this proposal, you get a different inheritance plan with each γ function; so we can regard a subject's range of inheritance plans as a range of γ functions.

As it turns out, the formalism we get from the third proposal suffices for modeling the other two proposals, too. Set γ as follows:

γ(a∨b∨d, ⊤)	=	1
γ(a∨c∨d, ⊤)	=	1
γ(a∨b, ⊤)	=	–1
γ(a∨d, ⊤)	=	–1
γ(a, ⊤)	=	1
γ(a∨b, a∨b∨d)	=	1
γ(a∨d, a∨b∨d)	=	1
γ(a, a∨b∨d)	=	–1
γ(a∨d, a∨c∨d)	=	1
γ(a, a∨c∨d)	=	–1

$\gamma(a, a \vee d) = 1$

and the conditional masses just specified reduce to:

$$m(a \mid a \vee b \vee c) = (\ m(a) + m(a \vee d)\ ) / \sigma''$$
$$m(a \vee b \mid a \vee b \vee c) = (\ m(a \vee b) + m(a \vee b \vee d)\ ) / \sigma''$$

which is what the Dempster-Shafer proposal said. Similarly, if we set γ like this:

$$\begin{aligned}
\gamma(a \vee b \vee d,\ \top) &= \alpha(a \vee b \vee d) / \alpha(\top) \\
\gamma(a \vee c \vee d,\ \top) &= \alpha(a \vee c \vee d) / \alpha(\top) \\
\gamma(a \vee b,\ \top) &= -\,\alpha(a \vee b) / \alpha(\top) \\
\gamma(a \vee d,\ \top) &= -\,\alpha(a \vee d) / \alpha(\top) \\
\gamma(a,\ \top) &= \alpha(a) / \alpha(\top)
\end{aligned}$$

$$\begin{aligned}
\gamma(a \vee b, a \vee b \vee d) &= \alpha(a \vee b) / \alpha(a \vee b \vee d) \\
\gamma(a \vee d, a \vee b \vee d) &= \alpha(a \vee d) / \alpha(a \vee b \vee d) \\
\gamma(a, a \vee b \vee d) &= -\,\alpha(a) / \alpha(a \vee b \vee d)
\end{aligned}$$

$$\begin{aligned}
\gamma(a \vee d, a \vee c \vee d) &= \alpha(a \vee d) / \alpha(a \vee b \vee d) \\
\gamma(a, a \vee c \vee d) &= -\,\alpha(a) / \alpha(a \vee c \vee d)
\end{aligned}$$

$$\gamma(a, a \vee d) = \alpha(a) / \alpha(a \vee d)$$

we get the conditional masses stated by the α-based proposal. Notice that to do this, we need to let γ sometimes take a negative value. The third proposal doesn't itself supply any intuitive motivation for that. So I regard the Dempster-Shafer proposal and the α-based proposals as *competitors* to our third proposal, rather than special cases of it. Don't think that their intuitive motivation comes from any way of telling the third story but with "negative inheritance." Rather, it comes from the stories I set out earlier. It's just that this "trick" of sometimes letting γ take negative values allows us to use *the equations* appropriate to the third proposal to cover all three proposals.

In what follows, then, I will just use the γ-based equations for conditional masses. Moreover, I'll for a time write as though there were only a *single* inheritance plan γ. In the full story, though, everything we say needs to be applied to *each* of the γ functions that are in a subject's range of inheritance plans.

11. What are the general equations for all of this? We can build conditional masses out of a recursive 3-place function φ (i, h, H), where i,h each $\subseteq$ H and $i \cap h \neq \varnothing$. This is defined as follows:

$$\varphi(i, h, H) = m(i \mid H) + \sum_{\substack{i \subset j \subseteq H \\ j \not\subseteq h}} \gamma(i,j)\, \varphi(j, h, H)$$

One's conditional mass m(i | h) will be equivalent to φ(i, h, H), only normalized so that the m(• | h) add up to 1. The relevant normalization factor will be: $\sum_{j \subseteq h} \varphi(j, h, H)$. H in these equations can be any hypothesis such that φ(i, h, H) is defined, including ⊤. The results will be the same.

Here's a few examples of how this plays out:

$$
\begin{aligned}
\varphi(\top, a \vee b \vee c, \top) &= m(\top) \\
\varphi(a \vee b \vee d, a \vee b \vee c, \top) &= m(a \vee b \vee d) + \gamma(a \vee b \vee d, \top)\, \varphi(\top, a \vee b \vee c, \top) \\
\varphi(a \vee c \vee d, a \vee b \vee c, \top) &= m(a \vee c \vee d) + \gamma(a \vee c \vee d, \top)\, \varphi(\top, a \vee b \vee c, \top) \\
\varphi(a \vee d, a \vee b \vee c, \top) &= m(a \vee d) \\
&\quad + \gamma(a \vee d, a \vee b \vee d)\, \varphi(a \vee b \vee d, a \vee b \vee c, \top) \\
&\quad + \gamma(a \vee d, a \vee c \vee d)\, \varphi(a \vee c \vee d, a \vee b \vee c, \top) \\
&\quad + \gamma(a \vee d, \top)\, \varphi(\top, a \vee b \vee c, \top) \\
\varphi(a, a \vee b \vee c, \top) &= m(a) \\
&\quad + \gamma(a, a \vee d)\, \varphi(a \vee d, a \vee b \vee c, \top) \\
&\quad + \gamma(a, a \vee b \vee d)\, \varphi(a \vee b \vee d, a \vee b \vee c, \top) \\
&\quad + \gamma(a, a \vee c \vee d)\, \varphi(a \vee c \vee d, a \vee b \vee c, \top) \\
&\quad + \gamma(a, \top)\, \varphi(\top, a \vee b \vee c, \top)
\end{aligned}
$$

etc.

That gives us a general algorithm for conditional masses.

Our algorithm for *flattening* a mass distribution into a probability function takes a similar form. The basic idea is that, when we're calculating the probability of a hypothesis, we let it and its subhypotheses inherit *every* portion they're entitled to from *all* their superhypotheses. So, for instance, the probability of a would be:

$$
\begin{aligned}
&\ \ m(a) \\
&+ \gamma(a, \top) * m(\top) \\
&+ \gamma(a, a \vee b \vee c) * (\, m(a \vee b \vee c) + \gamma(a \vee b \vee c, \top) * m(\top)\,) \\
&+ \gamma(a, a \vee b \vee d) * (\, m(a \vee b \vee d) + \gamma(a \vee b \vee d, \top) * m(\top)\,) \\
&+ \gamma(a, a \vee c \vee d) * (\, m(a \vee c \vee d) + \gamma(a \vee c \vee d, \top) * m(\top)\,) \\
&+ \gamma(a, a \vee b) * (\\
&\qquad\quad m(a \vee b) \\
&\qquad + \gamma(a \vee b, \top) * m(\top) \\
&\qquad + \gamma(a \vee b, a \vee b \vee c) * (\, m(a \vee b \vee c) + \gamma(a \vee b \vee c, \top) * m(\top)\,) \\
&\qquad + \gamma(a \vee b, a \vee b \vee d) * (\, m(a \vee b \vee d) + \gamma(a \vee b \vee d, \top) * m(\top)\,) \\
&\quad) \\
&+ \gamma(a, a \vee c) * (\\
&\qquad\quad m(a \vee c) \\
&\qquad + \gamma(a \vee c, \top) * m(\top) \\
&\qquad + \gamma(a \vee c, a \vee b \vee c) * (\, m(a \vee b \vee c) + \gamma(a \vee b \vee c, \top) * m(\top)\,) \\
&\qquad + \gamma(a \vee c, a \vee c \vee d) * (\, m(a \vee c \vee d) + \gamma(a \vee c \vee d, \top) * m(\top)\,) \\
&\quad) \\
&+ \gamma(a, a \vee d) * (\\
&\qquad\quad m(a \vee d) \\
&\qquad + \gamma(a \vee d, \top) * m(\top) \\
&\qquad + \gamma(a \vee d, a \vee b \vee d) * (\, m(a \vee b \vee d) + \gamma(a \vee b \vee d, \top) * m(\top)\,) \\
&\qquad + \gamma(a \vee d, a \vee c \vee d) * (\, m(a \vee c \vee d) + \gamma(a \vee c \vee d, \top) * m(\top)\,)
\end{aligned}
$$

)

In general, we can define $p(h \mid H) = \sum_{j \subseteq h} \varphi(j, h, H)$. This you'll notice is exactly the normalization factor for conditional masses. So the normalized conditional mass m(i | h) will = φ (i, h, H) / p(h | H), for any suitable H.

I've characterized p(•|•) as a probability function, but in fact whether it's a probability function depends on the structure of inheritance plans. On the α-based and γ-based ways of thinking about inheritance, p(•|•) does turn out to be a probability function. But on the first, Dempster-Shafer way of thinking about inheritance, p(•|•) instead turns out to be what Shafer calls a *plausibility function*. This is something like an upper probability. On the Dempster-Shafer model, the plausibility of a + the plausibility of not-a will usually exceed 1. It's surprising that a single formal model of confidence should be suited to represent both Dempster-Shafer and traditional Bayesian thinking. To preserve generality, I propose to call our p(•|•) functions "plausibilities," and to let the further formal properties of a plausibility function be settled by the operative inheritance structure. If masses ought to be inherited in the Dempster-Shafer way, then our plausibility functions will be Dempster-Shafer plausibilities. If masses ought to be inherited in the α-based or γ-based ways, instead, then they will be probabilities.

I'll remind you that our present claims about conditional mass are all about "classical" cases, where no dogmatic support is present. Given the way we're understanding conditional mass, it follows that for any i ⊆ j ⊆ h:

$$p(i \mid j)\, p(j \mid h) = p(i \mid h)$$

This is true for Dempster-Shafer plausibilities as well as for traditional probabilities. It will fail to be true, though, when we move to non-classical cases, with dogmatic support. (In fact, I think that different distributions of dogmatic support map onto the plausibility functions p(•|•) where some such inequality is violated; see fn. 18, below.)

12. At the end of §9, I said that one of the "path-independence" constraints I wanted to respect was that: if you conditionalize an epistemic state on the hypothesis a∨c, and then update the result on evidence of strength *J* for a, you should get the same result as if you had first updated the unconditionalized state on evidence of strength *J* for a (or for any hypothesis h such that h ∩ a∨c = a), and then conditionalized on a∨c.

Let's see what that constraint teaches us. Suppose you first get evidence for the hypothesis a∨b, which gives you quantity θ of new mass to add to your distribution and renormalize. We have yet to determine what the relation is between your evidence and θ; so for now θ is an unknown. As we said in §9, your unconditional masses should update like this:

$$
\begin{aligned}
m'(a) &= (\, m(a) + \theta\, m(a \mid a \vee b) \,) / (1+\theta)\\
m'(b) &= (\, m(b) + \theta\, m(a \mid a \vee b) \,) / (1+\theta)\\
m'(a \vee b) &= (\, m(a \vee b) + \theta\, m(a \vee b \mid a \vee b) \,) / (1+\theta)\\
m'(c) &= (\, m(c) \,) / (1+\theta)\\
m'(a \vee c) &= (\, m(a \vee c) \,) / (1+\theta)
\end{aligned}
$$

etc.

If we then conditionalize your updated mass distribution m′ on the hypothesis avc, using the algorithm from §§10–11, and renormalize, we derive:

$$m'(a \mid avc) = (\, m(a \mid avc)\, p(avc) + \theta\, (\, m(a \mid avb) + \gamma(a, avb)\, m(avb \mid avb)\,)\,) \,/\, (\, p(avc) + \theta\, p(a \mid avb)\,)$$

$$= (\, m(a \mid avc)\, p(avc) + \theta\, p(a \mid avb)\,)\,) \,/\, (\, p(avc) + \theta\, p(a \mid avb)\,)$$

$$m'(c \mid avc) = m(c \mid avc)\, p(avc) \,/\, (\, p(avc) + \theta\, p(a \mid avb)\,)$$

$$m'(avc \mid avc) = m(avc \mid avc)\, p(avc) \,/\, (\, p(avc) + \theta\, p(a \mid avb)\,)$$

Suppose, on the other hand, we had started with the conditional masses m(• | avc), and updated *them* on the same evidence. In this case, the evidence should be understood to be supporting the intersection x of avb and the hypothesis we're conditionalizing on (avc). So x=a. And whereas before our evidence gave us quantity θ of new mass to add to our distribution, we shouldn't assume that it supplies the same quantity to add to this conditional mass distribution. Let's designate the quantity it does contribute with a second unknown, θ_x. Applying the same updating operation, we get:

$$m'(a \mid avc) = (\, m(a \mid avc) + \theta_x\, m(a \mid x)\,) \,/\, (1+\theta_x)$$

$$m'(c \mid avc) = (\, m(c \mid avc)\,) \,/\, (1+\theta_x)$$

$$m'(avc \mid avc) = (\, m(avc \mid avc)\,) \,/\, (1+\theta_x)$$

Our "path independence" constraint tells us that these results should match the results derived above, gotten by first updating then conditionalizing. Setting them equal, we obtain:

$$p(avc) \,/\, (\, p(avc) + \theta\, p(a \mid avb)\,) = 1 \,/\, (1+\theta_x)$$

Solving for θ_x:

$$\theta_x = \theta\, p(a \mid avb) \,/\, p(avc)$$

Now, choose some value ⅃ such that ⅃ p(avb | ⊤) = θ. Then θ_x will = ⅃ p(avb | ⊤) p(a | avb) / p(avc), which in classical cases will = ⅃ p(a | avc). That's very interesting. It tells us that we can use ⅃ as our invariant *J*, measuring the evidence's intrinsic strength; and that:

> The quantity *θ* of new mass to use, when applying evidence of strength *J* to hypothesis h, conditional on supposition H (which may be ⊤), is equal to *J* p(h ∩ H | H).

This yields the most elegant formal system. However, it has the consequence that evidential strengths don't add linearly: evidence of strength *J1* for h plus evidence of strength *J2* for h are equivalent to evidence of strength *J1+J2+J1 J2* for h. One can move to a different scale to avoid this. (Designate your evidence's linear strength *L*, and set $J = e^L - 1$. Then evidence of (linear) strength *L1* for h plus evidence of strength *L2* for h will be equivalent to evidence of strength *L1+L2* for h.) But that needlessly complicates the formalism.

When you get evidence of strength *J* for avb, it follows that your new plausibilities will be this:

$$p'(a) = (\, p(a) + J\, p(avb)\, (\, m(a \mid avb) + \gamma(a, avb)\, m(avb \mid avb)\,)\,) \,/\, (1+ J\, p(avb)\,)$$

$$= (\ p(a) + J\, p(a\vee b)\ p(a \mid a\vee b)\ ) / (1 + J\, p(a\vee b)\ )$$
$$= \ p(a)\ (1+J) / (1 + J\, p(a\vee b)\ )$$
$$p'(b) = p(b)\ (1+J) / (1 + J\, p(a\vee b)\ )$$
$$p'(c) = p(c) / (1 + J\, p(a\vee b)\ )$$
$$p'(d) = p(d) / (1 + J\, p(a\vee b)\ )$$

In other words, the effect on plausibilities is to multiply the plausibilities of the supported hypothesis a∨b by 1+*J*, and then to renormalize. One can easily see that this operation commutes: if you get evidence of strength *J1* for hypothesis h1, and evidence of strength *J2* for hypothesis h2, you can just multiply the h1-plausibilities by (1+*J1*) and multiply the h2-plausibilities by (1+*J2*), and renormalize. It doesn't matter what order you do it in.

Updating in the way I've described also commutes in its effect on masses, though I'll leave that as an exercise for the reader.

13. The operations I've described may seem computationally complex. However, if we make certain choices about scaling, then it all just turns out to be simple matrix algebra.

We've seen that updating on evidence of strength *J* for hypothesis a∨b should have the following effect on one's unconditional masses:

$$m'(a) = (\ m(a) + J\, p(a\vee b|\top)\ m(a \mid a\vee b)\ ) / (1+J\, p(a\vee b|\top))$$
$$m'(b) = (\ m(b) + J\, p(a\vee b|\top)\ m(a \mid a\vee b)\ ) / (1+J\, p(a\vee b|\top))$$
$$m'(a\vee b) = (\ m(a\vee b) + J\, p(a\vee b|\top)\ m(a\vee b \mid a\vee b)\ ) / (1+J\, p(a\vee b|\top))$$
$$m'(c) = (\ m(c)\ ) / (1+J\, p(a\vee b|\top))$$
$$m'(a\vee c) = (\ m(a\vee c)\ ) / (1+J\, p(a\vee b|\top))$$
etc.

The effect on one's conditional masses m(•|a∨c) will be:

$$m'(a \mid a\vee c) = (\ m(a \mid a\vee c) + J\, p(a \mid a\vee c)\ \ m(a \mid a)\ ) / (1+J\, p(a \mid a\vee c))$$
$$m'(c \mid a\vee c) = (\ m(c \mid a\vee c)\ ) / (1+J\, p(a \mid a\vee c))$$
$$m'(a\vee c \mid a\vee c) = (\ m(a\vee c \mid a\vee c)\ ) / (1+J\, p(a \mid a\vee c))$$

In general, for any $i \subseteq$ each of a∨b and h:

$$m'(i \mid h) = (\ m(i \mid h) + J\, p(h \cap a\vee b \mid h)\ m(i \mid h \cap a\vee b)\ ) \ / (1+J\, p(h \cap a\vee b \mid h))$$

And for any $i^* \subseteq h$ but $\not\subseteq$ a∨b:

$$m'(i^* \mid h) = (\ m(i^* \mid h)\ ) / (1+J\, p(h \cap a\vee b \mid h))$$

If we multiply these through by $(1+J\, p(h \cap a\vee b \mid h))$, we get:

$$m'(i \mid h)\ \ (1+J\, p(h \cap a\vee b \mid h)) = m(i \mid h) + J\, p(h \cap a\vee b \mid h)\ m(i \mid h \cap a\vee b)$$
$$m'(i^* \mid h)\ (1+J\, p(h \cap a\vee b \mid h)) = m(i^* \mid h)$$

Replacing $p(h \cap a\vee b \mid h)$ with $p(h \cap a\vee b \mid \top) / p(h \mid \top)$, and multiplying through by $p(h \mid \top)$:

$$m'(i \mid h)\ \ (\, p(h \mid \top) + J\, p(h \cap a\vee b \mid \top)\,) = m(i \mid h)\, p(h \mid \top) + J\, p(h \cap a\vee b \mid \top)\, m(i \mid h \cap a\vee b)$$

$$m'(i^* \mid h)\, (\, p(h \mid \top) + J\, p(h \cap a\vee b \mid \top)\,) = m(i^* \mid h)\, p(h \mid \top)$$

Then, since m(i | h) will in general = φ (i, h, ⊤) / p(h | ⊤):

$$\varphi'(i, h, \top)\ \ (\, p(h \mid \top) + J\, p(h \cap a\vee b \mid \top)\,) / p'(h \mid \top) = \varphi(i, h, \top) + J\, \varphi(i, h \cap a\vee b, \top)$$

$$\varphi'(i^*, h, \top)\, (\, p(h \mid \top) + J\, p(h \cap a\vee b \mid \top)\,) / p'(h \mid \top) = \varphi(i^*, h, \top)$$

p′(h | ⊤) will = (p(h | ⊤) + J p(h ∩ a∨b | ⊤)) / (1+J p(a∨b | ⊤)), so:

$$\varphi'(i, h, \top)\ \ (1 + J\, p(a\vee b \mid \top)\,) = \varphi(i, h, \top) + J\, \varphi(i, h \cap a\vee b, \top)$$

$$\varphi'(i^*, h, \top)\, (1 + J\, p(a\vee b \mid \top)\,) = \varphi(i^*, h, \top)$$

Multiplying through by an arbitrary constant c, we get the following:

$$c'\, \varphi'(i, h, \top)\ \ = c\, \varphi(i, h, \top) + J\, c\, \varphi(i, h \cap a\vee b, \top)$$

$$c'\, \varphi'(i^*, h, \top) = c\, \varphi(i^*, h, \top)$$

where $c' = (1 + J\, p(a\vee b \mid \top)\,)\, c$. These results are easily represented with matrix algebra.

Where h ⊆ H , suppose **Φ**hH is the vector $c \begin{bmatrix} \varphi(a,h,H) \\ \varphi(b,h,H) \\ \varphi(a\vee b,h,H) \\ \ldots \\ \varphi(h,h,H) \end{bmatrix}$. Even without knowing the constant c, we can derive the masses m(• | h) by normalizing that vector. Let **Φ***H be concatenation of **Φ**hH for every h ⊆ H: $c \begin{bmatrix} \varphi(a,a,H) \\ \varphi(b,b,H) \\ \varphi(a,a\vee b,H) \\ \varphi(b,a\vee b,H) \\ \varphi(a\vee b,a\vee b,H) \\ \ldots \\ \varphi(H,H,H) \end{bmatrix}$. The elements of **Φ***H will have some order, e.g., such that their first two arguments are <a,a>, <b,b>, <a,a∨b>, <b,a∨b>, <a∨b,a∨b>,…,<H,H>. Call that the **serialization** of **Φ***H. Let λ be the length of that sequence, and define **J**hH(y) to be the sum of the λ × λ identity matrix and the following λ × λ matrix, where <rowi,rowj> is the pair of arguments at the row's position in the **Φ***H-serialization, and <coli,colj> is the pair at the column's position:

$$\begin{bmatrix} \ldots & \ldots & \ldots \\ \ldots\text{if rowj} \cap \text{h} = \text{colj and rowi=coli then y; else 0}\ldots \\ \ldots & \ldots & \ldots \end{bmatrix}$$

For instance, using the serialization described above, **J**a(a∨b)(y) =

$1+y$	0	0	0	0
0	1	0	0	0
y	0	1	0	0
0	0	0	1	0
0	0	0	0	1

.

If we multiply **J**hH(y) × **Φ***H, we get the vector c' $\begin{bmatrix} \varphi'(a,a,H) \\ \varphi'(b,b,H) \\ \varphi'(a,a\vee b,H) \\ \varphi'(b,a\vee b,H) \\ \varphi'(a\vee b,a\vee b,H) \\ \ldots \\ \varphi'(H,H,H) \end{bmatrix}$, where

$c' = (1+y\ p(h \mid H))\ c$, and φ′ is the result of updating on evidence of strength y for hypothesis h. That vector can be unpacked and normalized to yield m′(• | h) for any h ⊆ H. So essentially, updating is just a matter of multiplying matrices.

Updating commutes because the construction of these **J** matrices is such that, for any hypotheses h1, h2 ⊆ H and strengths y1, y2: **J**h1H(y1) × **J**h2H(y2) always = **J**h2H(y2) × **J**h1H(y1).

In classical cases, like we've been considering, conditionalizing and flattening can also be made into matrix operations.

As before, let h ⊆ H. Define |h| to be the number of non-empty subhypotheses h has; and similarly for |H|. Understand rowi to be the hypothesis which is at the row's position in the serialization of **Φ**hH; and similarly for coli. Then define **Q**hH to be the following |h| × |H| matrix:

$$\left[\text{if rowi} = \text{coli then 1; else if rowi} \subset \text{coli} \not\subseteq h \text{ then} \sum_{\substack{e \in \text{powerchain of} \\ \{\text{rowi}\} \cup \{k \mid \text{rowi} \subset k \subseteq \text{coli and } k \not\subseteq h\}}} \prod_{\substack{x \text{ and } y \text{ adjacent} \\ \text{elements in } e}} \gamma(x,y)\right]$$

where the powerchain of a set S of hypotheses is every element of S's powerset that contains min(S) and max(S) and whose members are pairwise ⊃ or ⊂ each other; and x and y are adjacent elements in e iff: x ∈e, y∈e, x ⊂ y, and there's no z∈e such that x ⊂ z ⊂ y.

So, for instance, **Q**a(a∨b) will =

1	0	$\gamma(a, a\vee b)$

and **Q**(a∨b)(a∨b∨c) will =

1	0	0	0	$\gamma(a,avc)$	0	$\gamma(a,avbvc) + \gamma(a,avc)\ \gamma(avc,avbvc)$
0	1	0	0	0	$\gamma(b,bvc)$	$\gamma(b,avbvc) + \gamma(b,bvc)\ \gamma(bvc,avbvc)$
0	0	1	0	0	0	$\gamma(avb,avbvc)$

These **Q** matrices can be used to conditionalize: just let **Φ**jH = **Q**jh × **Φ**hH.

The construction of **Q** matrices is such that **Q**jh × **Q**hH will always = **Q**jH. This corresponds to the "path-independence" of conditionalizing that I stated as a desideratum in §9.

Where $j \subseteq h$, let ∑Qjh be the 1 × |h| matrix that's the result of adding up each of the columns of **Q**jh. Define ||j|| to be the number of atomic subhypotheses j has, and define **Ψ**jh to be the ||j|| × |h| array $\begin{bmatrix} \cdots \\ \sum Qih \\ \cdots \end{bmatrix}$, where i ranges over the atomic subhypotheses of j. Then the plausibilities p(•|j) of a mass distribution corresponding to **Φ**hH will = the normalization of **Ψ**jh × **Φ**hH.

14. There's a thread of literature arguing that the right way to understand Jeffrey Conditionalization for degrees of belief is as an operation that changes the "Bayes Ratio" of one's belief in a specified way. That is how we're supposed to settle when two subjects should count as having updated on "equivalent evidence." For instance, if Adam has probability function pa and acquires some evidence that directly supports h, then we calculate the Bayes Ratio:

$$\beta = \frac{pa'(h) \,/\, pa'(\text{not-h})}{pa(h) \,/\, pa(\text{not-h})}.$$

We'd need to induce an update with the same Bayes Ratio in Betty's probability function pb, for her to count as acquiring evidence of the same strength for h. That is, for Betty too, it would need to be the case that:

$$\beta = \frac{pb'(h) \,/\, pb'(\text{not-h})}{pb(h) \,/\, pb(\text{not-h})}.$$ [11]

How do things play out with the plausibility updates I've described? Well, for an update on evidence of strength J for avb, the Bayes Ratio is:

$$\frac{p'(avb) \,/\, p'(cvd)}{p(avb) \,/\, p(cvd)} = \frac{p(avb)\,(1+J) \,/\, p(cvd)}{p(avb) \,/\, p(cvd)} = (1+J).$$

This will be constant no matter what p(•) was like to start with (we just require that the relevant ratios be defined). So the updating procedure I've described turns out to be Jeffrey Conditionalization!

More carefully: I've described an updating procedure on a structure richer than a probability function. That rich structure can be "flattened" into a "plausibility function." If we make certain theoretical choices about how inheritance works, then plausibility functions are just probability functions, and the effects that updating induces on those probability functions are just the ones that Jeffrey Conditionalization does.

---

[11] Cite Field, Wagner, Jim Hawthorne. They argue that by construing "equivalent evidence" in this way—and only by doing so—can we make Jeffrey Conditionalization commutative.

If, on the other hand, we understand inheritance in the Dempster-Shafer way, then the updating procedure I've described constitutes an application of Dempster's "Rule of Combination."

We can understand what I've done so far as providing some extra structure to help us keep track of how much of your confidence is fixed by evidence (that is, is assigned by your mass distribution to more atomic hypotheses), and how much isn't. I've told some intuitive stories about inheritance and updating and so on, that motivate making that extra structure behave in ways that correspond to more familiar models of belief. So far we haven't used the extra structure to do anything novel.

15. It's time now to consider how things should go when we add dogmatic support to our model. Let's construe the atoms of W like this:

		$\top$			
a∨b∨c		**a∨b∨d**	**a∨c∨d**		**b∨c∨d**
a∨b Q is true and now perceivable by me	**a∨c**	**b∨c = E** I am now having experiences η as of perceiving Q	**a∨d**	**b∨d**	**c∨d** Q is not now perceivable by me (Q may be true or false)
a …but I don't now even seem to perceive Q		**b = PERC** I am now perceiving that Q	**c = U** undermining scenario (Q may be true or false)		**d** …and I don't now even seem to perceive Q (it may be true or false)
		$\bot$			

We could subdivide U and d even further, but this is enough for our purposes. Our main concern will be with the subhypotheses of E: PERC, U, and their disjunction, E itself.

The dogmatist says that η has some dogmatic bias towards PERC over U. This bias gets undermined by any *evidentially based* confidence you have in U—in our model, the relevant factor will be m(U | E)—but not by any confidence you have in U that's not based on evidence—in our model, that will be any mass that U merely "inherits" from m(E | E).

Here's one way we might model that. Suppose you get dogmatic evidence of strength *K*. Treat that as evidence for E, only instead of the classical updating operation:

$$\begin{aligned}
m'(\text{PERC}) &= (\, m(\text{PERC}) + K\, p(\text{E} \mid \top)\; m(\text{PERC} \mid \text{E}) \,) / \sigma''' \\
m'(\text{U}) &= (\, m(\text{U}) + K\, p(\text{E} \mid \top)\; m(\text{U} \mid \text{E}) \,) / \sigma''' \\
m'(\text{E}) &= (\, m(\text{E}) + K\, p(\text{E} \mid \top)\; m(\text{E} \mid \text{E}) \,) / \sigma''' \\
m'(\ldots) &= (\, m(\ldots) \,) / \sigma'''
\end{aligned}$$

do this instead:

$$m'(\text{PERC}) = (\, m(\text{PERC}) + K\, p(\text{E} \mid \top)\, (\, m(\text{PERC} \mid \text{E}) + m(\text{E} \mid \text{E}) \,) \,) / \sigma'''$$

$$
\begin{aligned}
m'(U) &= (m(U) + K\, p(E \mid \top)\ \ m(U \mid E)) / \sigma''' \\
m'(E) &= (m(E)) / \sigma''' \\
m'(\ldots) &= (m(\ldots)) / \sigma'''
\end{aligned}
$$

What have we done? In this update, we've let PERC get the benefit of *all* the new mass that U wasn't already entitled to in the classical update. That's a way for the evidence to be "biased" in favor of PERC—to give PERC the benefit of any doubt that isn't already *evidentially* assigned to U. As m(U | E) increases, there will be less and less extra mass of this sort for PERC to benefit from; so the dogmatic effect will tend to diminish. Moreover, notice that m′(U) on the new update comes out the same as it would on the classical update. The dogmatic bias in no way alters *the evidence you have* for U. It only steals from the evidentially uncommitted confidence m(E | E), some part of which U classically claims to *inherit*.

That's the basic idea I will develop. However, it won't quite work in its present form.

One difficulty stems from the fact that we now seem to have two different updating operations. What would distinguish the case where we should update *classically* on evidence of strength *K* for E, and the case where we should update *dogmatically*? Do we need to build a third parameter into our representation of evidence, indicating whether it has any dogmatic bias? Wouldn't it be more natural to make the bias be intrinsic to the experience η, which E is just the hypothesis that you're having? Then it ought to show up in *any* update on E.[12] Yet nothing in our formalism indicates that updates on E should always be dogmatic.

A second difficulty stems from the fact that the dogmatic update I described will raise the cumulative mass of E—that is, m(PERC)+m(U)+m(E)—as much as classical evidence of strength *K* for E would. But what if you haven't *yet* acquired evidence that you'll have the experiences in question? Weatherson argues that in some cases, you might not enjoy the dogmatic benefits of some evidence until you actually acquire the evidence.[13] But even he is reluctant to say that will *always* be the case—that the dogmatic bias in some evidence can *never* show up prospectively, before you've acquired the evidence.[14] I'd like to find a way to represent a dogmatic bias in one's epistemic state *prior to acquiring any evidence* that you in fact do or will enjoy experiences η.

---

[12] Granted, subjects aren't always *aware* of what the proper effects of their evidence should be. That's something they can be reasonably uncertain about. But it's also the kind of reflective failure that models like mine, and the Bayesian's, abstract away from (see fn 1). We only try to represent what epistemic positions are *available* to subjects, should they reflect and epistemologize properly. Hence, if it's *true* that your experiences have a dogmatic bias, then we should take E to be the hypothesis that you have experiences *with* that bias.

[13] Weatherson, "The Bayesian and the Dogmatist."

[14] This came out in discussion of his paper when he visited the NYU Mind & Language seminar.

Now, you may wonder: if I'm going to say there's a bias in your epistemic state *pre-evidentially*, then what distinguishes my dogmatism from epistemologies of perception that I've taken pains elsewhere to oppose, which say we have *a priori entitlement* to believe we're not brains in vats?[15]

I agree with those other epistemologists that we need an anti-skeptical bias somewhere; the question is where. I'm not satisfied with just giving subjects unconditional a priori justification to believe they're not in skeptical scenarios. The constraints that places on their priors seem too arbitrary and ad hoc. It also seems too disconnected from what I take to be the proper *source* of anti-skeptical bias: the distinctive epistemic properties that dogmatists attribute to our experiences. I'd prefer any anti-skeptical bias to be represented as somehow *coming from* our experiences, even if the bias is in place antecedently to our *having* the experiences.

That raises a second point. I think the most interesting question here is not: Is our anti-skeptical bias already in our epistemic state prior to our having experiences? That's one thing we may mean by asking whether the bias is a priori. The more interesting question is: Is the anti-skeptical bias present *in virtue of* the epistemic properties of experience—albeit perhaps experiences we haven't yet enjoyed—or does it rather derive from extra-experiential considerations? This is another thing we may mean by asking whether it's a priori. The view I prefer says that our anti-skeptical bias is a priori in the first sense but not the second: It's there in our state *before* the experiences are, but it's there *because* experiences have the epistemic properties they do.

That's what I'd like to say. Let's see whether we can revise our model of dogmatic updating in a way that accomplishes it.

16. The best solution I've found to the difficulties set out in §15 goes like this. Begin by applying the dogmatic update procedure on one's masses *conditional on E* (and E's subhypotheses) with evidence of strength *K* for E:

$$
\begin{aligned}
m'(\mathrm{PERC} \mid \mathrm{E}) &= \big(m(\mathrm{PERC} \mid \mathrm{E}) + K\, p(\mathrm{E} \mid \mathrm{E})\, (m(\mathrm{PERC} \mid \mathrm{E}) + m(\mathrm{E} \mid \mathrm{E})) \big) / \sigma''' \\
m'(\mathrm{U} \mid \mathrm{E}) &= \big(m(\mathrm{U} \mid \mathrm{E}) + K\, p(\mathrm{E} \mid \mathrm{E})\, m(\mathrm{U} \mid \mathrm{E}) \big) / \sigma''' \\
m'(\mathrm{E} \mid \mathrm{E}) &= \big(m(\mathrm{E} \mid \mathrm{E}) \big) / \sigma'''
\end{aligned}
$$

Leave one's unconditional masses, and one's masses on hypotheses that $\not\subseteq$ E, alone. *K* here represents the strength of the experience's dogmatic bias. Perhaps this will vary with some properties of the experience: e.g. more forceful experiences get higher *K*s. Or perhaps we should always let *K* be maximal (take the limits as $K \to \infty$). I'll take its proper value to be supplied by considerations outside of the formalism.

At this point, you just deploy the matrix algebra laid out in §13! That is, if $\mathbf{\Phi}^{*\top}$ represents your epistemic state prior to the application of any bias, and $\mathbf{B}^{*\top}$ represents it after the biasing operation described in the previous paragraph, then upon gaining new evidence of strength y for hypothesis h, you let your epistemic state be determined by $\mathbf{J}h^{\top}(y) \times \mathbf{B}^{*\top}$, instead of $\mathbf{J}h^{\top}(y) \times \mathbf{\Phi}^{*\top}$. That's all there is to it.[16]

---

[15] This is endorsed by Wright [cite]; Cohen 1999?, 2000?; White 2006; others.

[16] The details are available in a *Mathematica* workbook that I'll place on my website. As we'll see below, one can no longer rely on the **Q** matrices to derive

As you then *do* get evidence for E, some portion of it will work classically and some portion will work dogmatically. How much does which will depend on the size of *K*. In the limit, as you come closer to having certain evidence that E, you'll approach the initial dogmatic update laid out in the previous section.

This model has many intuitively appealing features.

Firstly, the dogmatic bias gets to be present prospectively, in your conditional epistemic state, before you've yet had any dogmatic experiences.

Secondly, no decisions need to be taken about whether to update dogmatically or classically. The *K*-bias in your conditional masses automatically settles what happens when you do get evidence that E.

Thirdly, the model gives us intuitive results about other conditional masses. It turns out we can't use the same conditionalizing algorithms we developed in §§10–11.[17] However, what we can do is this. We let your epistemic state be determined not merely by your unconditional masses, or their encoding into the vector $\mathbf{\Phi}\top\top = \begin{bmatrix} \varphi(a,\top,\top) \\ \varphi(b,\top,\top) \\ \varphi(a \vee b,\top,\top) \\ \dots \\ \varphi(\top,\top,\top) \end{bmatrix}$. Rather we take the concatenation of *all* your conditional and unconditional masses, encoded into the (longer) vector $\mathbf{\Phi}^*\top = \begin{bmatrix} \varphi(a,a,\top) \\ \varphi(b,b,\top) \\ \varphi(a,a \vee b,\top) \\ \varphi(b,a \vee b,\top) \\ \varphi(a \vee b,a \vee b,\top) \\ \dots \\ \varphi(\top,\top,\top) \end{bmatrix}$. Prior to the application of any bias, the conditional parts of this could be derived in the usual way from the unconditional part. But when we apply a dogmatic bias, we end up with a new vector $\mathbf{B}^*\top$. The conditional masses encoded in this vector will diverge from the relations the conditional masses in $\mathbf{\Phi}^*\top$ stand in to the unconditional masses in $\mathbf{\Phi}\top\top$. The degree to which they diverge encodes the amount of dogmatic bias in one's epistemic state.[18] For instance, the vector that encodes the *K*-bias described in our

---

conditional masses; we'll need to keep track of them as part of an expanded epistemic state. However, the **J** and $\mathbf{\Psi}$ matrices continue to work as laid out in §13.

17 Those algorithms guarantee the result stated in §11, that for any $i \subseteq j \subseteq h$: $p(i \mid j)\, p(j \mid h) = p(i \mid h)$. But as we'll see, the model of dogmatic bias I've presented requires that equation to sometimes fail.

18 Indeed, I think we can formally identify *any* divergence whatsoever between one's conditional masses and the results determined by the algorithms of §§10–11 with some amount of dogmatic bias. However, this exploits a feature of my formal model that I haven't spelled out here, where some hypotheses exhibit *competing* biases: as though E

example will = $\begin{bmatrix} \varphi(a,a,\top) \\ \ldots \\ \varphi(PERC,E,\top)+\frac{k}{1+k}\,\varphi(E,E,\top) \\ \varphi(U,E,\top) \\ \frac{1}{1+k}\,\varphi(E,E,\top) \\ \ldots \\ \varphi(\top,\top,\top) \end{bmatrix}$. We then identify one's epistemic state with the combination of *that vector* $\mathbf{B}^*\top$ and a range of inheritance functions. One's conditional masses will always be encoded in one's **B**-vector at that time.

We preserve the intuitive idea floated in §9: your conditional mass in i, given the *supposition* that h, still turns out to be the same confidence you should end up having in i were you to acquire *certain evidence* that h is true.

We get other intuitive results too. Suppose you acquire evidence of strength *J* for E. If you've got the dogmatic biases I've described, then your masses conditional on any hypothesis that intersects *both* PERC and U will update somewhat non-classically. But *your masses conditional on any hypothesis that intersects only one of them* will behave just as they would if you had *classically* updated on evidence of strength *J* for E. For example, conditional on the assumption that you're not perceiving, it's as though the dogmatic bias in your experience weren't even there. This is a consequence of the rules for biasing and updating that we've already laid out.

There is one result that's less intuitive. Because your *masses* conditional on E are non-classical, so too will your *plausibilities* conditional on E (or on any hypothesis that intersects both PERC and U, including $\top$) be non-classical. In particular, in the case we described, and in the absence of any evidence for hypotheses intersecting E:

$$p(PERC \mid E) = p(PERC) / p(E) + K\,\gamma(U, E)\, m(E \mid E)$$

$$p(U \mid E) = p(U) / p(E) - K\,\gamma(U, E)\, m(E \mid E)$$

When you do have evidence for hypotheses intersecting E, the relationship will be more complicated. But it will only be true that $p(\bullet \mid E) = p(\bullet) / p(E)$ when you become certain that E. This makes our plausibilities diverge from the standard probabilistic definitions of conditional probability.

Is that something we can live with?

17. It does raise the threat that we might be vulnerable to diachronic Dutch Books.

It's not clear to me how bad that would be, if it were true. The philosophical significance of diachronic Dutch Books is much contested.[19]

---

had *some* amount of dogmatic bias for PERC and a *competing* amount of bias for U. This is a useful part of the formalism, but I'm not sure there are any intuitive applications for it.

19 Howson and Urbach 1989; Christensen 1996; and G. Hellman, "Bayes and beyond" *Philosophy of Science Vol 64* (1997), 191-221 try to give diachronic Dutch Books a non-pragmatic interpretation. More skeptical: P. Maher, "Depragmatized Dutch Book arguments" *Philosophy of Science Vol 64* (1997), 291-305. Others?

Still, I'd be more comfortable if my theory *didn't* recommend that subjects expose themselves to diachronic Dutch Books. And as it turns out, it might not.

How can that be? Don't we know from the existing literature that one will be Dutch Bookable just in case one's conditional probability p(PERC | E) diverges from the ratio p(PERC & E) / p(E)? And doesn't my theory propose that conditional probabilities *should* so diverge, exactly when E dogmatically supports PERC (and you're not yet certain that E)?

As we've seen, the answer to the second question is yes. But let's be careful about the first. What the existing literature tells us is that one will be Dutch Bookable just in case one's decisions are determined by expected utilities derived from *a single* probability function whose conditional values p(• | E) diverge from p(• & E) / p(E). But recall that our plausibility functions are generated by the combination of one's mass distributions and *a range* of inheritance plans. This will determine *a range* of plausibility functions. The literature may teach us that *individual members* of that range recommend decisions that are diachronically Dutch Bookable. But it's not clear that it follows that decisions guided by *the range itself* must be Dutch Bookable too.

Here's what I have in mind.

The Bookie puts you in a scenario like this. First, your eyes are masked and you're asked to bet on certain questions having to do with whether you'll soon have visual experiences as of a hamster; and if you do, whether they'll constitute *perceptions* of a hamster. Then your eyes are unbound and you settle whether you do have those visual experiences. If you do, the Bookie may then transact further bets about whether your experiences constitute perceptions. After any such further transactions, you get to reach out and settle, by feeling, whether there is a hamster there. (We suppose there is no possibility of your tactile perceptions being mistaken. We suppose also that, in the scenario, you'll visually perceive a hamster iff there's a hamster there. There is no possibility of your having veridical hallucinations.)

What the Bookie will try to do is this:

(i) He offers you a bet that pays \$1 if U—that is, if you have experiences as of a hamster but aren't thereby perceiving (because there is no hamster). You will pay up to \$ p(U) for this bet.

(ii) He offers you a bet that pays \$ p(U)/p(E) if not-E—that is, if you don't have experiences as of a hamster. You will pay up to \$ p(U)/p(E) * p(not-E) for this bet.

(iii) He offers you a bet that pays \$ p(U)/p(E) – p(U | E) if E. You will pay up to \$ (p(U)/p(E) – p(U | E)) * p(E) for this bet.

At this point, you will have spent up to: \$ p(U) + p(U)/p(E) – p(U | E) p(E). Now you get to open your eyes.

If you don't have experiences as of a hamster, your bets (i) and (iii) lose, and bet (ii) pays you \$ p(U)/p(E). At this point, you'll be down by as much as \$ p(U) – p(U | E) p(E). If p(U | E) < p(U)/p(E), this will be a positive amount.

If on the other hand, you *do* have experiences as of a hamster, then bet (ii) loses, bet (iii) pays \$ p(U)/p(E) – p(U | E), and bet (i) is still open. The Bookie then continues by:

(iv) Offering to buy from you a bet that pays him \$1 if U. Since you've now become certain that E, you'll be willing to sell him this bet for as low as \$ p(U | E).

At this point, you'll be down by up to \$ p(U) – p(U | E) p(E). Now you get to reach out and touch and see whether you really are perceiving a hamster. If there is a hamster there, he'll bite you (this doesn't enter into your utility calculations). Bet (i) will then lose, but at the same time you don't have to pay on bet (iv). If on the other hand, there is no hamster there, then you win \$1 on bet (i), but have to pay out \$1 on bet (iv). In either case, you're left down \$ p(U) – p(U | E) p(E). If $p(U \mid E) < p(U)/p(E)$, this will be a positive amount.

So goes a standard diachronic Dutch Book.

However, let's now consider what happens if your betting decisions are determined by a *range* of plausibility functions. When offered bet (i), you'll be willing to pay up to \$ $p_{low}(U)$—the lower bound of your expected utility for bet (i). When the Bookie offers to buy bet (iv), you'll be willing to sell it for as low as \$ $p_{high}(U \mid E)$. And so on. We have to keep track of whether the relevant plausibilities are your upper bounds, or your lower bounds, or some value in between. So far as I've been able to determine, you will be Dutch Bookable only if you violate one of these two inequalities:

(a) $$\frac{p_{high}(X \,\&\, E)}{p_{low}(E)} \geq p_{low}(X \mid E)$$

(b) $$p_{high}(X \mid E) \geq \frac{p_{low}(X \,\&\, E)}{p_{high}(E)}$$

And whether my model will permit violations of those inequalities turns out to depend on the size of one's dogmatic bias *K* and what range of inheritance plans one is working with. This problem is too hard to solve analytically, but it can be mathematically modeled. With the help of Branden Fitelson, I modeled it for some 3-atom hypothesis spaces and the α-based inheritance plans. It appears to be the case that: if a subject uses *the entire available range* of α-based inheritance plans, then arbitrary large dogmatic biases *K* are possible, without violating (a) or (b). If a subject uses smaller ranges of α-based inheritance plans, then there will be constraints on how large *K* can be, before it permits violations of (a) or (b) for some possible mass distributions. However, as long as the range of inheritance plans has positive measure, there will still be positive values of *K* that don't expose the subject to any possible violations of (a) or (b).

I wish that this matter were more clearly settled. I'd like to know whether diachronic Dutch Books do in fact always require violations of (a) or (b). I'd like to have an analytic solution to the question how high *K* can be, before exposing a subject to violations of (a) or (b). And I'd like to have a solution for γ-based inheritance plans as well as α-based ones. Still, the evidence we do have suggests that my model might *not* expose subjects to being diachronically Dutch Booked.

18. Summarize how this model solves the difficulties we had making probabilistic sense of dogmatism:

- How it models the difference between direct and mediate undermining
- The sense it gives to the idea that your experiences can be dogmatically biased towards PERC rather than U, even though having experiences η will raise the plausibility of U too—and given some kinds of background evidence (God told you $E \supset U$)—may raise p(U) much more than it raises p(PERC).

- On the dogmatist model, p(PERC | E) need not always be < p(not-U). So Theorem 2 from §4 can be violated. However, one’s evidence may be dogmatically biased towards PERC even in cases where that inequality is not violated.