Could've Thought Otherwise

Jonathan Weisberg
University of Toronto

Abstract. Evidence is univocal, not equivocal. Its implications don't depend on our beliefs or values, the evidence says what it says. But that doesn't mean there's no room for rational disagreement between people with the same evidence. Evaluating evidence is a lot like polling an electorate: getting an accurate reading requires a bit of luck, and even the best pollsters are bound to get slightly different results. So even though evidence is univocal, rationality's requirements are not "unique". Understanding this resolves several puzzles to do with uniqueness and disagreement.

You believe *P*, let's suppose. Could you have thought otherwise? More exactly: could you have reached a different conclusion from the same evidence, without being any less rational or justified?

On the one hand, it seems you could. Slight variations in opinion are commonplace and can seem quite reasonable, even once all the evidence and arguments are out on the table. When the evidence is complex and conflicted, significant variation can seem almost inevitable. On the other hand, there's something awkward about standing by one interpretation of the evidence while acknowledging that another is just as good. Why favour your interpretation if your fellow's is equally reasonable?

Some say the old Enlightenment ideal of evidence-based objectivity was overblown. Evidence is not the all-powerful beacon of rational light it's sometimes taken to be. Rather, much of what it says depends on help we supply articulating its implications. Interpretive assumptions we make, or values we hold, affect what conclusions the evidence supports.

The Jamesian pragmatist tradition sees evidence as open to different interpretations depending on one's epistemic values, for example (James, 1897; Levi, 1967, 1980). Subjective Bayesians hold that evidence says almost nothing about the probability of a hypothesis without help from our prior beliefs, which can vary wildly (de Finetti, 1937; Howson & Urbach, 2006; Jeffrey, 1983b; Ramsey, 1931). And some feminist

Thanks to Chris Meacham, Jennifer Nagel, Gurpreet Rattan, Miriam Schoenfield, Adam Sennet, Julia Smith, Sergio Tenenbaum, Mike Titelbaum, an anonymous referee for *Philosophers' Imprint*, and audiences at the University of Vienna, the University of Bristol, and the 9th Midwest Epistemology Workshop for their generous help and feedback on earlier drafts.

empiricists hold that auxiliary assumptions chosen based on ethical and political values are needed to interpret evidence, since evidence underdetermines theory (Anderson, 2004; Longino, 1990, 2002).

These are examples of *relativist* epistemologies. They make the correct interpretation of one's evidence relative to background beliefs, epistemic values, or political values. Whether body of evidence E supports P depends on a third relatum.

Contra relativism however, evidence is actually univocal. Its implications are absolute, not relative to a system of beliefs or values. The evidence says what it says, and it is our job to figure out what that is. At the same time though, there is more than one path to take in exploring what the evidence says. Thus there is more than one interpretation for reasonable people to arrive at after exploring the same body of total evidence.

In the lingo of recent debates, Feldman's (2007) 'uniqueness thesis' is false. But a closely related thesis I'll label 'univocity' is true.[1] Evidence speaks univocally, not equivocally. To think otherwise is to stand on unstable ground. Whichever interpretation of the evidence one embraces, a contrary interpretation must be deemed equally good, thereby undermining commitment to the first interpretation (White, 2005, 2013). Recent attempts to purge relativism of this instability fail, as we'll see in the first part of this paper (§§1–3).

Nevertheless, fleeing from relativism all the way to uniqueness goes too far. Uniqueness fails to account for the realities of how humans process evidence. Most interesting, perhaps: taking account of these realities sheds light on puzzles to do with variations in opinion and disagreements between peers, as we'll see in the paper's later parts (§§4–7).

1. The Instability Problem

Epistemologies that allow different reactions to the same total evidence have come to be called *permissive*. Subjective Bayesianism is one example of a permissive epistemology. Jamesian pragmatism and feminist empiricism are others, at least in certain versions.

These epistemologies are permissive because they are relativist. Because people with the same evidence can differ in their background beliefs, epistemic values, or political values, the same evidence can support different conclusions for different people. Although I reject relativism, I do embrace a different form of permissivism to be explained below (§4). Nevertheless, we'll focus on permissivism for now, since the

[1] Kopec & Titelbaum (2016) similarly distinguish between doxastic and propositional versions of the uniqueness thesis.

"instability" challenge I'll press against relativism is typically viewed as a challenge for permissive epistemologies in general.

I'll also focus on subjective Bayesianism as an example because of its popularity,[2] and because its formal character is useful for illustrating certain points. For example, it's fairly easy to see why subjective Bayesianism is permissive. According to it, a subject who acquires evidence E should set her new degree of belief in P to match her prior, conditional degree of belief $p(P \mid E)$.[3] But different people can have different prior degrees of belief. Any degrees of belief that obey the probability axioms provide a rational starting point for inquiry,[4] and the probability axioms are notoriously weak constraints. You can start with $p(P \mid E) = 1/10$, I can start with $p(P \mid E) = 9/10$. If we then acquire the same total evidence E, we'll end up with very different views about P. I'll think it's probably true, you'll think it's probably false.

Against such permissive views, White (2005, 2013) defends a thesis Feldman (2007) dubs Uniqueness.

Uniqueness: Given one's total evidence, there is a unique rational doxastic attitude that one can take to any proposition.

If you and I share total evidence E, we can't draw different conclusions about P, at least not without one of us making a mistake. If I'm right to conclude P, you are wrong to conclude $\neg P$. You can't even suspend judgment about P. The uniquely correct attitude is belief in P. Likewise for degrees of belief. If I'm right to be $9/10$ certain that P, and you have the same total evidence as me, then you would be wrong to be $1/10$ certain of P, or even $8/10$ certain.

Why think evidence is so restrictive? Why not allow that we can draw different conclusions from E, especially if our background beliefs about E's bearing on P are different? Permissive epistemologies tend towards a kind of internal instability.[5] They allow us to embrace one set of beliefs while simultaneously acknowledging that an alternative view is just as good. But acknowledging that a second perspective is equally legitimate threatens our commitment to the first. How can we favour one view over another if they are equally good?

---

[2] In particular, I am concerned to respond to authors like Douven (2009), Kelly (2013), Meacham (2014), and Schoenfield (2014), who either express sympathy for subjective Bayesianism or defend it outright.

[3] This is the rule of conditionalization (Talbott, 2001).

[4] Additional requirements are often imposed, like regularity (Shimony, 1955), reflection (van Fraassen, 1984), or the Principal Principle (Lewis, 1980). We'll focus on the naive version stated here for simplicty.

[5] More accurately, relativist epistemologies have this tendency. But we'll attend more carefully to this difference later.

Take subjective Bayesianism again as an example. Given your prior probabilities, observing a string of green emeralds makes it likely the next emerald observed will be green. But you must acknowledge, it would be just as rational to expect the next emerald to be blue, if one had "gruesome" prior probabilities instead. Although strange, gruesome probabilities are consistent with the laws of probability. If evidence can be so radically fickle though, what's it good for? Why even bother examining evidence if it can be interpreted any which way? Indeed, it seems you're at least as likely *a priori* to hit on a misleading interpretation as you are to be guided by the evidence to the truth. The next emerald could be orange, or purple, or yellow… you happen to favour an interpretation of the evidence that points to green. But with so many other, equally rational interpretations available, how can you favour just one with any confidence?

Note that this challenge assumes what Kopec & Titelbaum (2016) call "acknowledged permissive cases", where the agent realizes she is in a permissive case. Cohen (2013) defends a form of permissivism which denies such cases. However, we will focus on "acknowledged permissivism" here, since relativist views typically posit acknowledged cases. They claim to identify the third relatum evidential support is relative to, and that third variable is typically something we have at least some grip on: prior probabilities, epistemic values, or political values. So on these views, permissive cases will be acknowledged cases, at least sometimes.

White (2005, 2013) presses the instability problem further by extracting an especially absurd consequence: *belief toggling*. If factors besides the evidence affect what the evidence says, like background beliefs or values, then by manipulating these factors we can alter what beliefs our evidence supports. By embracing different prior probabilities, ones that favour grue over green, you can alter what conclusion your evidence supports. Instead of being highly confident the next emerald will be green, you can be highly confident the next emerald will be blue. And if the blue conclusion eventually bores you, you can switch back to the prior probabilities you had originally, and return to being confident the next emerald will be green.

Some Bayesians respond that you can't just change your prior probabilities (Douven, 2009; Meacham, 2014). Your prior degrees of belief are whatever they were before the string of green emeralds was observed. Degrees of belief can be updated when new evidence is acquired, but Bayesianism doesn't allow changing beliefs on a whim. It only allows changes in response to new evidence, as determined by your previous degree of belief $p(P \mid E)$.

Kelly (2013) responds in a similar vein. He notes that Bayesianism is only *interpersonally* permissive, not *intra*personally permissive. It permits different people

to reach different conclusions given the same evidence, but only when they have different beliefs prior to acquiring the evidence.

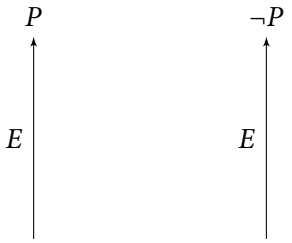

Figure 1. Interpersonal

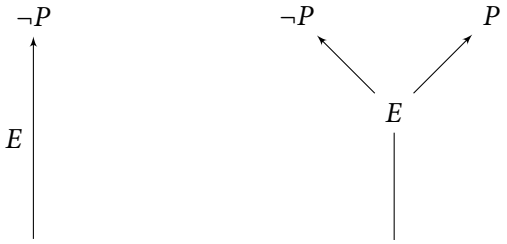

Figure 2. Intrapersonal

Bayesianism doesn't allow a single person to "choose" between two contrary responses to a piece of evidence. If your prior beliefs say P is probably true given E, then when you learn E, there is only one view you may take on P: it's probably true. You are then "locked in" to that judgment. You can't do anything to change the fact that, before learning E, you had a high conditional degree of belief $p(P \mid E)$. That's in the past.

Schoenfield (2014) offers a different response to the instability challenge, one independent of the Bayesian framework. According to Schoenfield, evidence has to be interpreted relative to one's *epistemic standards*. Roughly, your epistemic standards are the patterns of inference you think are truth-conducive.[6] Different people can have different epistemic standards, with full rationality. And these differences will oblige them to draw different conclusions from the same evidence.

What's to stop you switching epistemic standards, and thus toggling your beliefs? If you don't like what your evidence is telling you, why can't you just switch to standards that support a more pleasing interpretation? Because then you'll believe falsely, at least by your present lights. The standards you hold now tell you that competing sets of standards are misleading, less truth-conducive. Your current standards are the ones you think lead to the truth, by hypothesis. So if you switch to others, you will be setting yourself on the road to the false.

Neither of these responses to the instability problem proves satisfactory.[7] We'll start by uncovering the fault in the Bayesian reply of Douven and Meacham. That

[6]Bayesian probabilities can be interpreted as encodings of epistemic standards, rather than as degrees of belief. But Schoenfield's central claim isn't wedded to the Bayesian formal framework.

[7]Other responses merit consideration, though we can't engage with them here. For example, Ballantyne & Coffman (2011) note that externalist epistemologies can be permissive in a way, without being unstable. Since my main concern here is with epistemologies that make evidential support relative to internal factors, like background beliefs or values, I won't engage with this point here. Along with Douven (2009), they also note that inference to the best explanation might underwrite another kind of permissive, yet stable, epistemology. If P is the best explanation for E, one agent might conclude that P is true, while another who hasn't conceived of P does not. White would likely reply that these

will unearth a more general point, which guides us to the flaw in Schoenfield's appeal to epistemic standards. We'll then be in a position to identify what's really wrong with UNIQUENESS, and turn our attention to the correct view based on 'univocity'.

2. THE BAYESIAN REPLY

According to the Bayesian reply advocated by Douven (2009) and Meacham (2014), Bayesianism is only interpersonally permissive. People with different priors can draw contrary conclusions from the same body of evidence. But given whatever priors you have, there is only one conclusion to be drawn about P from evidence E. Your credence in P must match your conditional prior, $p(P \mid E)$. So you cannot belief-toggle.

The trouble for this response begins with a mundane observation: we can, and often do, revisit and re-evaluate evidence. On the orthodox version of subjective Bayesianism, evidence is only ever evaluated once. When a subject learns E, she adjusts her degree of belief in P to match her previous conditional degree of belief $p(P \mid E)$. After that, she never reconsiders what E tells her about P ever again. Yet in reality, we often go back and re-evaluate our evidence for P. Especially when we want to be sure we've evaluated that evidence correctly.

The orthodox model is well known to be unrealistic in this regard. But subjective Bayesianism is usually billed as a theory of ideally rational agents. So why does it matter if it makes no place for something we non-ideal agents do? Because making sense of what we non-ideal agents are doing has implications for ideal and non-ideal agents alike. Briefly, the reason is that we don't re-evaluate evidence using the probabilities dictated by our past degrees of belief. Instead, we use other probabilities. And choosing those other probabilities opens subjective Bayesianism up to the belief-toggling problem.

To begin, notice that our views about evidential support aren't always reflected in our current degrees of belief, as orthodox Bayesianism traditionally supposes. The reason is familiar from the literature on old evidence (Glymour, 1980). Once you have evidence E, your degrees of belief are such that $p(P \mid E) = p(P)$.[8] Looking at the conditional probability of P given E is useless now as a way of assessing E's bearing on P. Your degree of belief in P already reflects your previous evaluation

agents have different total evidence, since one knows that P is a possibility and the other doesn't. But assessing this dispute would take us too far afield, so I won't pursue it here.

[8]That's assuming $p(E) = 1$. What if instead $p(E)$ is merely close to 1, because we are fallible even with evidence? The problem would still arise for agents infallible about their evidence. And that would be enough for present purposes. But even for fallible agents who merely assign $p(E) \approx 1$, essentially the same problem arises, as Earman (1992) and Christensen (1999) show.

in light of E. If you want to revisit E's bearing on P, you must step back from your current perspective.

Step back to where? The natural answer is that you step back to your past perspective, from before you learned E. But notice, when you say that E is evidence for P, you are expressing a view you hold *now*. You aren't reporting an autobiographical fact, but rather articulating something about how you see things in the present. For example, you might be explaining that you believe P because you have good evidence for it, namely E. This present view you are expressing may align with the view you held in the past, before learning E. But it remains true that you are expressing a view you hold today, whether you held the same view yesterday or not.

So there seems to be some sense in which, for you, $p(P \mid E) > p(P)$ now. But the relevant probabilities here can't be your degrees of belief about P, E, and their Boolean combinations. Those degrees of belief are such that $p(P \mid E) = p(P)$.

Traditionally this point has been made in the literature on old evidence using examples of non-ideal agents, agents who fail to be logically omniscient.[9] This can make it appear as if the point is somehow limited to non-ideal agents, so that Bayesians concerned solely with ideal agents can safely ignore it. But that appearance is illusory.

Non-ideal examples do make it especially vivid that assessments of evidential support in the present are not autobiographical. Take a case where the agent conceives an explanation only after discovering the explanandum.[10] In such a case, $p(P \mid E)$ isn't well-defined until after $p(E) = 1$. So there is never a time when $p(P \mid E) > p(P)$, if probabilities are always degrees of belief.

Once we see that assessments of evidential support are not autobiographical, we can apply the point generally—to ideal Bayesian agents as well as to non-ideal agents. For the ideal agent, there may always be a past time when her degrees of belief were such that $p(P \mid E) > p(P)$. But her present view that E is evidence for P is no more autobiographical than the non-ideal agent's was. So that can't be what she is articulating when she says that E is evidence for P. There must be some other sense in which E probabilifies P for her *now*, even though her present degrees of belief are such that $p(P \mid E) = p(P)$.

---

[9] See Garber (1983), Jeffrey (1983a), Niiniluoto (1983), Eells (1990), and especially (Earman, 1992, 130-1) for excellent discussions of the issues here.

[10] Glymour's famous example of Einstein and Mercury's perihelion is one such case. Though it introduces complications, because Einstein failed to be logically omniscient in two different ways, what Earman (1992) labels LO1 and LO2. The cleaner kind of case, for our purposes, is one where the agent only fails logical omniscience in the second sense: she isn't aware of every possibility from the outset. I'm grateful to an anonymous referee for pressing me to clarify this point.

Many contemporary subjective Bayesians thus introduce the notion of 'hypothetical priors' (Meacham, 2008).[11,12] The rough idea is that the agent hypothetically deletes E from her total body of evidence and constructs a probability function different from her present credence function (and perhaps from her past credence functions too).[13] This probability function reflects "the way the agent her/himself views the problem" now, but with evidence E bracketed (Howson & Urbach, 2006, 300). And it's this probability function which underwrites her present view that E is evidence for P. Thus we have probabilistic commitments which aren't encoded in our current degrees of belief about P, E, and their Boolean combinations.

It's not clear from the Bayesian literature what hypothetical priors are in psychological terms, if they are not the agent's credences in P, E, etc. Are they beliefs with other contents, e.g. second-order beliefs about what your first-order beliefs ought to be, given this or that body of evidence? Are they instead first-order attitudes, but of a special kind or mode? Subjective Bayesians are tight-lipped here, so we can't say.

But whatever the psychological details, the problem for subjective Bayesians is to say why hypothetical priors shouldn't be toggled. When you step back from your belief about P, and re-evaluate E's import in light of your hypothetical priors, conditionalization fixes your credence in P. So there is no intrapersonal permission at this level. But, just as we can step back from our ordinary beliefs about P and E, we can step back from our hypothetical priors too. We can step back and ask whether we are obliged to project green instead of grue. And according to subjective Bayesianism, priors that project grue are just as rational as those that project green. So, when we step back to reconsider not only our "surface" beliefs, but also our hypothetical priors, we seem to be permitted to belief-toggle.

Ideal agents are not exempt here. Perhaps they need never step back from their beliefs in order to correct mistakes, since they don't make any. But even if they,

[11]Some other prominent examples include Bartha & Hitchcock (1999), Howson (1991), and Howson & Urbach (2006). The term 'ur-priors' is sometimes also used (Manley, forthcoming), though that term can have a different connotation. For example, Meacham (2010) uses 'ur-priors' for a subject's degrees of belief before she acquires any evidence, in order to formulate chance-credence principles like the Principal Principle (Lewis, 1980). Since this kind of clean-slate state is often viewed as a kind of fiction, 'ur-priors' might be conflated with the kind of 'hypothetical priors' needed to make sense of talk of evidential support (see e.g. Meacham 2008, fn. 7). But we needn't assume these two kinds of probability go together here.

[12]Some prefer to stick with probability as degree of belief, and rely instead on technical maneuvers, like insisting that evidence shouldn't have probability one (because no empirical proposition should), and that we should use a more sophisticated measure of evidential support than the simple difference between $p(P|E)$ and $p(E)$. Christensen (1999) and Hawthorne (2005) show what fates these maneuvers are doomed to. I won't rehearse the arguments here.

[13]Notice that it can't simply be the credence function she would have had if she hadn't learned E, for familiar reasons identified by Maher (1996) and Hawthorne (2005).

unlike us, are not obliged to re-evaluate their evidence, they can still be *permitted* to re-evaluate. It might be a waste of time and energy for them, and hence prudentially irrational (though given their powers, perhaps not). But that's no epistemic reason. So they will be epistemically permitted to step back from their hypothetical priors, and switch to others. Moreover, even if ideal agents aren't permitted to re-evaluate, it's bad enough if we, imperfect agents are allowed to belief-toggle on the hypothetical priors view. That's already a reductio. We are right to change our minds in light of new evidence or further reasoning, but not on a whim.

Here is another way to put the challenge for subjective Bayesianism—more contentious, but helpful for the discussion to follow.[14]

The trouble with orthodox subjective Bayesianism is that it fails to acknowledge the ability to take different perspectives. Having drawn various conclusions in light of E, we can of course continue to evaluate things from the point of view where E. But we can also step back to an E-neutral point of view, to revisit what E says. Yet orthodox subjective Bayesianism identifies probabilities with degrees of belief, and one's present degrees of belief reflect only one's current perspective, which assumes E.

The Bayesian reply to the belief-toggling problem only works by ignoring the E-neutral perspective. It works by pretending we can't re-evaluate E's bearing on P. Once that pretense is abandoned, some other kind of subjective probability has to be introduced. And that introduces the question why that kind of probability can't be toggled. Why must evaluations from the E-neutral perspective follow one coherent probability function rather than another?

It's natural to understand the hypothetical priors governing the E-neutral perspective along the lines offered by Schoenfield, as one's epistemic standards.[15] The probabilities in the E-neutral perspective represent deeper epistemic commitments than the ones being re-evaluated, i.e. deeper than our present credences in E, P, etc. And Schoenfield's view is expressly concerned to answer the question we've found ourselves left with here, namely why we shouldn't toggle our hypothetical priors/epistemic standards. So let's turn to the epistemic standards view now.

3. The Epistemic Standards Reply

Unlike the orthodox subjective Bayesian, Schoenfield (2014) does acknowledge the E-neutral perspective. From that perspective, we re-evaluate E's bearing on P by

[14] I'm especially grateful to Gurpreet Rattan for suggesting this way of framing the issue, greatly clarifying this argument and the one in the next section.

[15] Similarly, see Levi's (1980) notion of a 'confirmational commitment'.

applying our epistemic standards. Roughly, our epistemic standards are the inference rules we think most truth-conducive.

The reason we shouldn't change our epistemic standards is that any other standards besides our own would lead us astray. From the point of view of whatever standards one holds, any other standards are not as truth-conducive. If I hold set of standards S, I think S has the best shot at leading me from the evidence to the truth. Any other set S^* is less likely to lead me to the truth.[16] So I should stick to S for epistemic reasons, even if S^* promises a more pleasing stance on P from a personal, ethical, or aesthetic point of view.

What about the S-neutral perspective, though? We're not only capable of stepping back from "surface" beliefs like E and P. We can also step back from our standards and re-evaluate them. So there is an S-neutral perspective as well as an E-neutral perspective. And from the point of view of the S-neutral perspective, I have no reason to continue to hold epistemic standards S rather than some other rational set of standards S^*. So I can replace my present standards with S^*, and change my mind about P accordingly. That is, I can belief-toggle.

Schoenfield responds that this challenge is just the radical skeptical challenge. It demands a justification for endorsing standard S from a perspective where no standard is in place. So it asks the impossible. It's like the challenge of justifying the laws of logic without relying on logical inference. It can't be done, whether one is a permissivist or not. Even defenders of UNIQUENESS can't answer such a skeptical challenge, Schoenfield notes. So it isn't any fault of permissivism that it can't solve that problem either.

But there is a special problem for the permissivist here. By her own admission, there is a non-skeptical perspective from which we can evaluate competing standards S and S^*. This is exactly what a permissivist claims to be able to see: that our standards are not the only rational ones, others hold contrary standards that are rational too.[17] So from the permissivist's own perspective, we can see that S and S^* are equally promising as guides to the truth. So we would be rational to switch from S to S^*, and thus from believing P to believing $\neg P$. That is, we are permitted to belief-toggle according to Schoenfield's permissivism.[18]

[16] A family of theorems in probability theory lends some support to this idea in a Bayesian context (Greaves & Wallace, 2006; Joyce, 1998, 2009).

[17] And not just because we adopt a hopelessly agnostic perspective from which any old set of standards looks equally rational. According to "moderate" permissivists like Schoenfield, at least, only some standards are rational. Horowitz (2014) discusses challenges for moderate vs. immoderate permissivism in more detail.

[18] As an anonymous reviewer points out, Schoenfield might respond that there is no S-neutral perspective. Rather, there are many perspectives, each of which seems rational from its own point of

Proponents of Uniqueness aren't committed to this problematic perspective, however. Comparing S to S^* from an S-neutral perspective does not rate them on a par. Only S should be deemed rational, not S^*.[19] So even when we step back from our standards, we are forced to return to them and stand by the beliefs they generated. Endorsing Uniqueness thus affords a kind of internal stability that permissivism imperils.

To illustrate the point more concretely, consider a probabilistic formalization of Schoenfield's account.

A probability function p encodes a set of epistemic standards: rules for going from total bodies of evidence to doxastic states. $p(P \mid E)$ is the confidence you should have in P if your total evidence is E. The relativist says there are multiple rational sets of standards, different probability functions one might embrace. For the sake of illustration, let's follow those Bayesians who say that any probability function encodes a rational standard, because these functions, and only these functions, avoid being *accuracy dominated*. A function that violates the probability axioms will be further from the truth—no matter what the truth turns out to be—than a corresponding probability function (Joyce, 1998, 2009). But probability functions do not have this flaw.

Now suppose you've gone your whole life using a uniform probability function v, one that assigns equal probability to all possibilities.[20] Midway in your life's journey, you find yourself with a high credence in P, to your dismay. Can you get out of believing P somehow? You adopt an E-neutral perspective and go back to re-examine your evidence. No luck: $v(P \mid E)$ is indeed high. So you adopt a v-neutral perspective, i.e. you step back from the epistemic standards you've always held. You compare v to another probability function v^*, which is similar except that $v^*(P \mid E)$ is low. You think to yourself: "both v and v^* describe rational standards because both are equally good guides to the truth a priori. Neither one is accuracy dominated. And beyond that I have no reason to embrace one set of epistemic standards over

---

view. Given standards S, standards S look rational; given standards S^*, standards S^* look rational; etc. But this view seems too weak for the permissivist's usual purposes. Unless there is some perspective from which one can see that alternative standards really are rational—not just that they appear that way to those who hold them, but that they really are—then we cannot say that there are innocent disagreements between people who share their evidence. Nor can we say that people with different priors, or different epistemic values, or different political values, are rational to interpret the evidence differently. We can only say that they think they are. In truth, at most one of them is really right about that.

[19] Or, if we've stepped too far back to a hopelessly skeptical perspective, neither can be deemed rational or irrational.

[20] The possibilities will need to be parameterized first if there are continuum-many. That is, they'll need to be mapped onto a real interval one-to-one. Any arbitrary parameterization will do for now.

another." So you switch to v^* and modify your belief in P accordingly, becoming confident in $\neg P$ instead.

This can't happen on the objective approach to Bayesianism. Continuing our hypothetical illustration, suppose now that we follow instead those Bayesians who think the uniform probability function v is the *only* rational standard. Perhaps because it alone minimizes the risk of inaccuracy (Pettigrew, 2015), or because it alone minimizes information/bias (Jaynes, 1957), or because it alone achieves calibration (Hawthorne et al., 2015). Permissivists will object to all of these rationales. But let's pretend one of them is correct just for the moment. Under that pretense, even when you step back to a v-neutral perspective, you won't find yourself permitted to switch to some other probability function v^*. Even from the v-neutral perspective, v is the only rational standard to hold. It is the only probability function that minimizes the risk of inaccuracy, minimizes information/bias, or achieves calibration. These distinguishing features of v can be seen without embracing the epistemic standards encoded in v. They are theorems, deducible by mathematical means. So the v-neutral perspective need not be a hopelessly skeptical perspective.

The flaw in Schoenfield's defense is to assume that stepping back to evaluate our standards is always to adopt a hopelessly skeptical perspective. But permissivism by its very nature is a view that posits a not-hopelessly-skeptical perspective, where standards can be evaluated and compared. Permissivists claim to be able to identify multiple rational standards. So they must be able to see that those standards are equally truth-conducive from some more neutral perspective (presumably by some a priori means).[21] So there is a perspective from which we can see that some standards are rational to hold, others not. When we enter that perspective, switching standards becomes an option, and belief-toggling becomes a threat.

Schoenfield may object that switching standards from S to S^* isn't permitted, even from an S-neutral perspective. Just being able to see that S^* is rational in principle doesn't make it rationally permissible for you to adopt it. After all, you still hold standards S, so S^* should still look inferior to you as a guide to the truth.

But the whole point of stepping back is to bracket the commitment we're stepping back from, opening it up to revision. When we step back from a mundane belief like P, we do so precisely for the purposes of considering whether to change that belief. If we find that E actually supports $\neg P$, contrary to what we initially thought, we become *obligated* to change our view of P. Likewise, if we step back from standard S to find

[21] We illustrated what such a stance might look like in a Bayesian context. A priori, the mathematical, accuracy-dominance argument for probabilistic standards is available. But it doesn't pick out any one probability function as privileged.

that it's actually irrational, we would become obligated to reject it. By symmetry, if we find that S is rational but so is S^*, we become *permitted* to adopt either one.

The crux of the epistemic standards view is that some of our commitments are deeper than others, and thus take priority. When we step back to re-evaluate a more superficial commitment in light of a deeper one, the deeper one takes priority. If we step back from S to find that S^* is just as rational according to our deepest standards, those standards have priority. Since they say S and S^* are equally rational, we can embrace either one. But then we can belief-toggle.

4. Univocity, not Uniqueness

I've been criticizing permissivist responses to the instability challenge. But not because I think Uniqueness is true. It's false, and here is a first pass at the reason why.

Imagine I ask you how many of your friends wear glasses. Is it more than half? Suppose you answer yes: more than half. The belief you've just formed is presumably justified, rational. But with equally strong credentials it could have gone another way. Fewer of your spectacled friends might've come to mind. Or they might have come to mind but with much greater difficulty. In either case, you would have estimated the proportion of your friends who wear glasses lower, and rightly so. From the very same store of memories, you could have formed a contrary belief with equal justification.

The example highlights the role of *processing* in the move from evidence to belief. We can separate two stages leading up to the formation of beliefs about the world:

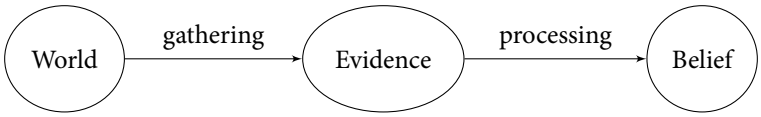

Variation emerging at the first stage is nothing novel. Pollsters can get very different samples even when studying the exact same population using the exact same methods. It takes a bit of luck to get a representative sample, and sometimes pollsters get unlucky.

Variation can emerge at the second stage too, though. If you estimate the proportion of your friends who wear glasses by sampling memories of your friends' faces, luck becomes a factor, much as it is for the pollster. Suppose you do a quick, random search of memories of friends' faces to see how many of those memories feature glasses. Most likely, your sample will be representative and you will reach an accurate estimate of the proportion of your friends who wear glasses. But there's a chance your sample will be biased and you will reach a different estimate.

The point is that processing evidence can be stochastic. There is a fact about how many of your memories of friends' faces include spectacles, and a fact about what conclusion that entire body of evidence supports. But the process by which you divine those facts may have a random element, like a pollster trying to divine the inclinations of an electorate. If there is such a random element to the process, it can lead to divergent conclusions given a single body of evidence. Assuming you are justified and rational in using such a process, you can justifiedly and rationally reach different conclusions from one body of evidence. So UNIQUENESS is false.

Several aspects of this argument need to be developed and defended. First there's the empirical question whether we ever really do use stochastic processes, like "memory polling", to evaluate our evidence. Second is a question about sameness of evidence. Do such processes really reach different conclusions from the same evidence? Or do they merely change what evidence we have? Finally there's the normative question whether using such processes generates justified or rational beliefs. We'll answer these concerns in the coming sections.

But first we should ask how this view avoids the instability challenge. The answer is: the same way a pollster does.

If you take a random sample, and you do your job well, you find yourself in a sort of "permissive" position. You conclude that the population resembles your sample, though you acknowledge that the very same method applied to the very same population could have yielded different results. Most likely it would have yielded only slightly different results, warranting only a slightly different conclusion about how the vote is likely to turn out. But there's also a small chance it would have yielded *very* different results. Yet there is no instability in this position. Your possible counterparts who get different results do not make any rational error or draw any unjustified conclusions. They are just less fortunate than you, dwelling in one of the less probable possible worlds that emerges from the sampling process.

Similarly, there isn't necessarily anything irrational about concluding P from the evidence while allowing that $\neg P$ could have been rationally concluded instead. You just have to think that someone who did reach that conclusion would have been unfortunately misled. Their survey of the evidence would have run afoul of some bad luck. This is exactly what happens in the glasses example. You realize that a quick sampling of your memories could have led to a different conclusion. But because you rightly take your sample to be representative (or near enough), you think that other conclusion would have been based on a misleading sample of the evidence. There is no perspective from which embracing your counterpart's conclusion is a rational option for you.

Notice that we are rejecting Uniqueness in a very different way than the other permissive views we've encountered so far. For one thing, we aren't just allowing interpersonal variation, like subjective Bayesianism. We are allowing *intra*personal variation too. A single person with total evidence E could conclude either P or $\neg P$, depending how their processing of E goes. It's just that they don't get to "choose" which way the process goes. Good, honest sampling means letting chance do its work.

But more importantly, our view is permissive for a very different reason. Views like subjective Bayesianism, Jamesian pragmatism, and feminist empiricism are relativist, as noted earlier. According to them, evidential support is relative to a set of beliefs or values. In my view, evidential support is not relative, but absolute. What a body of evidence says does not depend on what background beliefs or values we bring to the table. Call this claim Univocity.

Univocity: Facts about evidential support are not relative. Whether evidence E supports proposition P, and to what extent, does not depend on supplementary beliefs or values.

The idea behind Univocity is that what the evidence says about P is never ambiguous. Evidence is not open to multiple interpretations depending on the background beliefs or values it's supplemented with. Facts about evidential support are two-place rather than three-place. They are "saturated", rather than waiting to be supplemented.

One way to think about the difference between Uniqueness and Univocity is in terms of Firth's (1978) distinction between doxastic and propositional justification (Matheson, 2011). Suppose for a moment that evidence is all that matters for justification. Then Uniqueness is a claim about doxastic justification, while Univocity is a claim about propositional justification. Uniqueness says what attitudes a subject who possesses E as their total evidence is justified in holding. By contrast, Univocity speaks to the abstract, "logical" question whether E justifies the proposition P.

Another way to think about it is that discussions of Uniqueness tend to conflate two different questions (Kopec & Titelbaum, 2016; Senor, manuscript). One is a metaphysical question within epistemology, the question whether the *evidence for* relation is two-place or three-place. The other is a normative question within epistemology, the question whether the facts about what the evidence says (whether two-place or three-place) determine everything about what we are permitted to believe. By endorsing Univocity, I am answering "two-place" to the first, metaphysical question. By denying Uniqueness I am answering "no" to the second, normative question.

The stability challenge may not establish UNIQUENESS, but it does support UNIVOCITY. If evidential support were relative to whatever beliefs or values one happened to hold, then the question would arise why we should bother listening to the evidence. It could be made to say different things, just by embracing different background assumptions or values.[22] Or, as White develops the argument, we have the option to belief-toggle. We just need to know what third variable evidential support is relative to, and possess the ability to fiddle with it, whether via pills, persuasion, or politics.

Which thesis is really at stake in current debates, UNIQUENESS or UNIVOCITY? I suspect each has its direct motivations, which then inform one another. For example, UNIQUENESS is interesting because of its significance for the epistemology of disagreement. Whereas subjective Bayesianism's rejection of UNIVOCITY can be seen as arising naturally out of the project of confirmation theory, combined with the formal limitations presented by dominance arguments (Dutch books, accuracy) and problems of language-dependence (Bertrand's paradox, grue). But that route to subjective Bayesianism then leads naturally to the rejection of UNIQUENESS as well.

In any case, if UNIQUENESS and UNIVOCITY come apart, we ought to separate them. Especially if we want to understand messy, real-world questions of disagreement, as many discussants of UNIQUENESS do. In real-world cases, how people process evidence and arrive at their conclusions is likely to be every bit as important as abstract, logical questions about what the evidence says in principle. Indeed, we'll soon discover two ways the present view illuminates questions of disagreement and variation in opinion (§6).

But first, we need to see the empirical support for the central, descriptive claim at the heart of this section's argument against UNIQUENESS. Do we really process evidence stochastically?

5. THE PSYCHOLOGY OF EVIDENCE PROCESSING

Well, how do humans process evidence in general? Philosophers tend to talk in terms of natural deductions or Bayesian calculations. But psychologists nowadays prefer very different models. No single model enjoys universal support. Indeed, it seems unlikely that we form all our judgments using one, single kind of process. Different models will likely be appropriate to different types of reasoning. But the important thing is that many leading models share in common the feature we need: they are partly stochastic.

[22] Again, there are important externalist concerns we can't get into here, like those raised by Ballantyne & Coffman (2011). See Figure 7.

5.1. **The Self-Consistency Model (SCM).** The Self-Consistency Model (SCM) proposed by Asher Koriat is one example that is obviously stochastic (Koriat, 2011, 2012; Koriat & Adiv, 2015). According to this model, some judgments are produced by just the kind of internal sampling described in the glasses example of the previous section. Faced with a choice between two possible answers to a question, a subject samples internal "representations"—memories, percepts, associations, or other materials relevant to the question at hand. She then performs a statistical test on that sample to form a conclusion. For example, she might apply a significance test to see whether the sample rules out one of the answers. The larger the sample and the lower its variance (the eponymous "consistency"), the more confident she will be in her final judgment.

To illustrate, let's apply the model to the glasses example. There are two hypotheses, that more than half your friends wear glasses and that no more than half do. Let the null hypothesis be the first one, that more than half do. To test this hypothesis you might take a sample of memories of friends' faces and count how many of those memories include glasses. Suppose you sample ten memories, and in only two of them is the person wearing glasses. Given the null hypothesis, that more than half your friends wear glasses, this is a pretty improbable outcome, we may suppose.[23] Let's suppose it's improbable enough to reject the null hypothesis. So you conclude that no more than half your friends wear glasses. (Given the small sample you won't be too confident in your answer, though.)

We're not aware of executing any such statistical processing, of course. We might catch a conscious glimpse of some of the memories as they're sampled, but even the sampling can be unconscious (Koriat, 2012, 83). For the most part, we're only aware of the "gist" (Koriat, 2012, 82), a rough sense of how well the whole process supports the final answer.

It's also worth emphasizing that the sampling needn't be limited to just memories of friends' faces, or even just to memories. Other memories could be relevant, like a memory of a friend saying they wear glasses at night after they've taken their contact lenses out. Likewise, general background knowledge can be relevant: how common are spectacles in the general population? Are glasses currently fashionable? Etc. The sampling process can reach across multiple kinds of "representations", integrating sampled memories, items of knowledge, and more. So a fully realistic treatment of the glasses example would likely be much more complicated. But the complications needn't detain us here.

[23] For example, given a binomial distribution where each memory has probability 0.6 of containing glasses, the p-value here is below .013.

The important thing for us is that this model makes evidence-processing stochastic. Indeed, it makes it stochastic in just the way our pollster analogy suggests. Answering the question works much like a pollster polling a population of memories.[24] If we really do use something like SCM to form some of our judgments, then it's easy to see how one might conclude P while allowing that concluding $\neg P$ instead could have been reasonable, even using the same store of memories, background knowledge, etc.

5.2. **The Evidence Accumulation Model (EAM).** Another prominent model is the Evidence Accumulation Model, or EAM (Lee & Cummins, 2004; Newell et al., 2007; Newell & Lee, 2011). On this model, evidence is not sampled randomly but sequentially, starting with the most informative evidence first.[25] Less informative evidence is considered next, then even less informative evidence, and so on until until one answer emerges as sufficiently well supported for present purposes. Any remaining evidence is ignored.

Despite using sequential rather than random sampling, EAM is still subject to random variation. If two pieces of evidence are equally informative, it's a tossup which one the agent will consider first. And the order can make all the difference. Suppose E and E' are equally informative, and that E favours answer A while E' favours answer B. If the evidence examined so far has almost decided the question in favour of A, and our subject examines E first, she might never get to examine E'. E will decide the question in favour of A, and the process will terminate before E' is evaluated. If she first examines E' instead, however, the balance will start to tilt back towards B. And if the next few pieces of evidence also favour B, she might end up embracing B instead of A.

A realistic interpretation of EAM will admit randomness in other places too. Plausibly, it's somewhat random what evidence is recalled at any given moment, and thus what evidence even gets ordered by informativeness. Measuring or estimating

[24] Or perhaps memories, items of knowledge, and more.

[25] Formally, informativeness is a conditional probability. Let A and B be the two possible answers under consideration, and let $A \succ B$ mean that A is better supported by all the evidence on balance. Now consider various positive features each answer might have or lack, labeled F_i. For example, feature F_1 might be whether a quick, random search of memories of friends' faces jives well with the answer. F_2 be whether it comports with the frequency of glasses-wearers in the general population, and so on. Then a feature's informativeness is:

$$p(A \succ B \mid F_i(A) \wedge \neg F_i(B)) \tag{1}$$

That is, feature F_i is informative to the extent that an answer is more likely to be correct given that it has feature F_i and the other answer does not. A piece of evidence E_i specifies which of answers A and B has feature F_i and which does not.

Psychologists prefer to speak of 'validity' rather than 'informativeness'. Because philosophers usually reserve 'validity' for deductively valid arguments, I avoid that terminology here.

informativeness is likely to be subject to at least a bit of random noise too. So when two features are similarly informative, random fluctuations could easily cause them to be considered in different orders.

In addition to allowing random variation, EAM also admits what we might call *practical variation*. EAM terminates when one of the two answers has accumulated enough support. But when is "enough" enough? The model deliberately leaves this threshold as an independent parameter. If a lot hangs on getting the correct answer, the threshold for "enough" support can be set high. In that case, a good deal of evidence will be considered before arriving at a conclusion. If getting the right answer isn't so important, or if time is tight, then the threshold can be set low. In that case, less evidence will be considered, maybe even as few as one or two items.[26]

Even without any random variation, then, an agent who reasons by EAM ends up in a position similar to one who reasons by SCM. "I conclude P based on my examination of the evidence," she can say to herself, "but it's possible that in other circumstances I just might have concluded $\neg P$ instead." Had her situation been different, with different resources available or different priorities to attend to, she might have thought about the question more deeply, setting a higher evidential threshold. Or, she might have set the threshold lower, examining less evidence and reaching a different conclusion as a result.

Would she have been justified in her conclusion had she set her threshold lower, and processed less evidence? We will address such normative questions in §7.2.[27] For now we are concerned with the descriptive question how we process evidence, and how that processing can vary.

5.3. **The Bigger Picture.** We've seen how two prominent models support the claim that we process evidence stochastically. Importantly, these models are not idiosyncratic in this regard.

EAM was developed in response to the dominant approach of the 1980's and 1990's: the "adaptive toolbox" paradigm. According to this paradigm we have many ways of processing evidence, and we choose the best one for the circumstance at hand. Like the tools in a toolbox, some of them are better suited to some jobs, others to others. The idea is that people choose their tool adaptively—they choose the algorithm that does the job at hand well enough while minimizing costs like time and effort (Gigerenzer, 2001; Payne et al., 1993; Payne & Bettman, 2004).

[26] The exact details are left open by EAM, and we can leave things vague for present purposes. See Hausmann & Läge (2008) for some discussion. See Payne et al. (1993), Payne & Bettman (2004), and Shah & Oppenheimer (2008) on the same issue in related frameworks.

[27] But, to preview, the short answer is yes, she might still have been justified

EAM aims to unify this motley assortment of tools, replacing it with a single tool that can be adjusted to the needs at hand. Instead of using different tools for different jobs, EAM always uses the same method: process evidence, in order of informativeness, until one answer passes the required threshold of evidential support. What varies is how high that threshold is set. The lower the threshold, the less evidence will be examined. The higher the threshold, the more evidence will be examined.

Whether EAM succeeds at unifying and replacing the adaptive toolbox is a matter of current investigation and debate (Bröder & Newell, 2008; Newell & Lee, 2011; Söllner et al., 2014). For our purposes it doesn't matter how that debate turns out. What matters is what the adaptive toolbox paradigm has in common with EAM and SCM.

Like EAM, the toolbox paradigm introduces both "random" and "practical" variation into the processing of evidence. And for essentially the same reasons EAM and SCM do. Some of the tools in the box use random sampling of the evidence like SCM, others use sequential sampling like EAM (Gigerenzer, 2001). Tools that use random sampling are stochastic on their face, the same as SCM. Those that use sequential sampling will, like EAM, show stochastic and practical variation for subtler reasons.

It's important to appreciate the roles models like SCM and EAM play in psychological theorizing. First, they are not necessarily competitors. It might well be that some of our judgments are formed by SCM, others by EAM, and still others according to some further means.[28] How various models fit together is hardly clear at present. But it's quite possible that something like SCM and EAM are both important parts of our overall cognitive architecture, along with other sorts of mechanisms as well. Second, neither SCM nor EAM is meant to be entirely realistic. For example, SCM gives equal weight to each item it samples. A more realistic, refined version would presumably be more intelligent, weighting some pieces of evidence more than others (Koriat, 2012, 110). It might also be only quasi-random, sampling more salient evidence with greater frequency. Just as pollsters use likely-voter models and other techniques to improve their projections, a refined, realistic version of SCM could be more accurate.

Luckily, we don't need to wait on a grand unified theory wherein all these details are worked out. The main point for us here is that there is very broad empirical support for the descriptive hypothesis introduced in §4. Several influential models posit stochastic evidence processing, including those models that have dominated

[28]Indeed, SCM comes out of research that has tended to focus on "Type I" processing, i.e. processing that is largely unconscious and automatic. Whereas EAM comes from research that has tended to focus more on "Type II" processing, which is conscious and deliberate. See Sloman (1996) and Evans & Frankish (2009) for background on the Type I/Type II distinction.

research on human judgment for the last three decades. Whatever the details, it looks quite likely that we process evidence stochastically, at least to some extent.

What about the careful, considered reasoning philosophers tend to be interested in, though? Especially when discussing peer disagreement, philosophers tend to be interested in persistent disagreements between experts—about the existence of God or the causes of climate change, for example. Can the kinds of considerations introduced here really bear on the sophisticated and thorough reasoning at work in such disagreements?

It can. For one thing, philosophers tackling disagreement often begin with simple cases of everyday reasoning (e.g. Elga 2007 and Christensen 2007). So philosophical theorizing is partly built on the same kinds of examples these psychologists study and model. But more than that, there is an important parallel between sampling memories and aggregating complex arguments and bodies of evidence. We can't hold all the arguments for and against the existence of God before our minds at once. Instead we have to work through arguments and responses, making mental notes and drawing tentative conclusions along the way. At any given moment, our opinion might be based on a sample of those mental notes and remembered conclusions, à la SCM. Or maybe it's more like an accumulated tally, à la EAM. Whatever the details, our cognitive limitations are limiting here if anywhere. Knowing how we cope with those limitations in general can only illuminate how we cope when things are especially difficult, as they are in philosophy and other complex areas.

There are other concerns still to be addressed, and we'll take them up in §7. But first let's enjoy some of the fruits of our labours. What virtues does the present view possess, and what lessons does it have to teach us?

6. Two Virtues

I claim that evidence is univocal: facts about evidential support are absolute, not relative. Yet gauging evidential support is a stochastic process, similar to measuring voter support with a poll. This view illuminates several issues in current debates about Uniqueness and disagreement. Here are two.

6.1. **Moderation.** In many cases where permissivism is intuitively appealing, it's a "moderate" form of permissivism that appeals (in the terminology of Horowitz 2014). For example, Kelly (2013, 299–300) notes that given the available evidence about the upcoming presidential election in the U.S., you and I might agree that the Republican is the favourite, yet differ a bit in our respective confidence levels. And in such a case, we might well regard one another as equally rational. But if I were to conclude from that evidence that the Green Party candidate is a shoo-in, you would

rightly think me unreasonable. The more opinions differ, the more pressure there is to say one of them must be irrational.

Why are small differences in opinion more likely to be seen as reasonable, and larger ones less likely?

Because small variations are a normal result of the stochastic process of assessing the evidence. If you and I survey the same body of evidence using a random sampling process like SCM, we will probably reach similar assessments. But if we reach wildly different conclusions, it's unlikely such an extreme difference is the result of random variation. Our difference in opinion probably isn't merely a result of getting different samples of the evidence. It's more likely I'm using some badly misguided method to evaluate the evidence, or that I'm just grossly biased. In general, the kind of random variation introduced by methods like SCM and EAM will be normally distributed with fairly tight clustering around the mean. So the further my opinion is from yours, the less likely it is to reflect a mere difference in my sampling of the shared evidence. More likely, it reflects a biased or otherwise irrational assessment.

Relativists will offer their own explanations for the phenomenon Kelly observes. For example, a Jamesian pragmatist might say that people tend to have similar epistemic values, and thus they tend to make similar assessments of the evidence. When their assessments differ wildly, it's more likely that bias or another error is the cause, rather than a difference in values. Likewise, subjective Bayesians might say that people tend to have similar priors. So a wild difference of opinion is more likely to reflect an error than a radical difference in priors.

The virtue of the present view is that it can explain the phenomenon without going relativist and risking instability. We can hold that evidential support is absolute, yet reject the claim that rationality's requirements are unique. And our reason for rejecting UNIQUENESS—that assessing evidence is partly a stochastic process—makes exactly the moderately permissive predictions Kelly observes. Reasonable differences of opinion will typically be small differences.

6.2. **Near Agreement.** Like UNIQUENESS, UNIVOCITY has implications for the epistemology of disagreement. Given UNIVOCITY, agreeing to disagree will generally be unacceptable for epistemic "peers", equally capable agents who share all their evidence. One of the disagreeing parties must be out of line with what the evidence supports, since they share the same evidence and it renders a univocal verdict. But neither party has any reason to think their peer is the one out of step with the evidence, rather than them.

Univocity thus lends support to the *equal weight view*, that each party should give the same weight to their peer's opinion as they give their own (Elga, 2007; Christensen, 2007). And, in practical terms, the equal weight view seems to recommend resolving the disagreement by compromise, where the two parties "split the difference". If one is 60% confident in *P* and the other 40%, they should meet in the middle at 50%.

A challenge for the equal weight view comes from cases of *near agreement*. To adapt an example from Christensen (2009),[29] suppose you're a doctor trying to choose the right dosage for a patient's condition: 5 mg vs. 10 mg. After re-reading their chart and giving it a think, you settle on 5 mg with 97% certainty. Now suppose you learn that your colleague arrived at the same dosage, but with slightly less certainty than you: 96%. Impressed that you both came to the same conclusion with very nearly the same certainty, you increase your confidence that the correct dosage is 5 mg. You shouldn't split the difference in this case. In fact you should go in the opposite direction from your peer, becoming more confident rather than less. But why?

As authors like Christensen note, learning that your peer has a different opinion provides you with additional, relevant evidence. But it needn't be evidence that either of you was irrational in your original judgment. It can be just the opposite: evidence that your original judgment was right on target. Your peer's opinion gives you evidence about what the evidence says—specifically, evidence that it says just what you took it to say: 5 mg is the right dosage. To see why, let's return to the polling analogy.

Suppose you and I conduct separate polls of the U.S. electorate using the same methodology. You find a solid majority supporting the Republican candidate, 65% to the Democrat's 30%. Let's suppose you come away 97% confident in a Republican victory as a result. I come back reporting 96% confidence in a Republican victory, though I don't tell you what proportions I found in my sample. Will you split the difference with me? Of course not: you'll become even *more* confident that the Republican will win. And I'll become more confident too. What each of us has learned is that another sample, besides the one we each took ourselves, also has a

[29]Christensen makes the example one where the relevant reasoning is deductive. This introduces complications I want to avoid in the main text. But very briefly, my account extends to such cases as follows. It's generally accepted that *fluency* is a significant factor affecting our confidence in our judgments (van Overschelde, 2008). The more quickly and easily a calculation goes, the more inclined you'll be to trust it. But fluency is subject to some random variation. So suppose you and your peer calculate the same result, but with slightly different confidence levels. What you learn is that another, equally competent agent came to the same conclusion, just a bit more slowly or with a smidge more difficulty. (Maybe they were momentarily distracted.) So you become more confident in your answer, because your peer essentially serves as a double-check, assuring you that you did the calculation right the first time.

solidly Republican majority. We've each effectively doubled our sample size.[30,31] And thus we've each bolstered the certainty of our original findings.

But notice that we get the opposite result if we focus on a different question. Suppose instead of asking which candidate will win, we ask how likely a randomly selected voter is to support the Republican candidate. Before consulting me, you'll think it about 65% likely. But after consulting me, you'll become *less* certain a randomly selected voter will favour the Republican, not more. What you've learned from my 96% confidence is that my sample, while staunchly Republican, is still slightly less Republican than yours. Maybe you can infer that my sample was 63% Republican. So now you effectively have a total, aggregated sample that is 64% Republican. You thus become 64% certain that a randomly selected voter supports the Republican candidate, down from 65%.

To apply these observations to the puzzle about near agreement, we just put all this "in the head".

When our two doctors consult one another, they discover something about the body of evidence they share: that it almost certainly favours a 5 mg dosage. You've found that not only does your sampling of the evidence lean towards 5 mg, but mine does too. So it's even more likely now that the whole body favours that dosage. And thus it becomes even more likely that dosage is correct.[32]

If we change the question though, we get a "split the difference" result instead. Suppose that, before you consult your colleague, you're asked how likely it is a randomly selected line on the patient's chart will indicate a 5 mg dose rather than 10 mg. Whatever answer you give, after consulting your colleague you'll give a lower estimate. If you said 65% initially, you'll say something like 64% after. For another thing you've learned from your colleague is that the total evidence supporting a 5 mg dosage may not be quite as decisive as it seemed at first. You become *more* confident that the total evidence supports 5 mg on balance, but slightly *less* confident that that balance tilts as heavily in favour of 5 mg as you initially thought.

A second virtue of the present view, then, is the flexibility to explain different responses to disagreement. It explains why we should split the difference sometimes, yet mutually boost others. In fact, discovering the very same point of disagreement can simultaneously cause you to split the difference on one proposition while mutually

[30] Assuming for simplicity that we're unlikely to get overlapping samples in such a large population.

[31] We've also partly verified the soundness of our methodology, which could be another factor at play in cases of near agreement.

[32] Glynn et al. (forthcoming) analyze and explain such "synergy" effects using a Bayesian analysis. They propose a rule for updating on the credences of others that can be computationally simpler than Bayesian updating, while having the same effect in a range of circumstances.

boosting on another. By attending to the character of the evidence in the shared pool, and to the way it bears on different questions, we acquire an elegant explanation for the puzzling, contrary-looking phenomena of difference-splitting and mutual-boosting.

7. Two Challenges

With the positive case for the present view on the table, and some of its virtues explored, let's turn to answer some challenges.

7.1. **Sameness of Evidence.** Grant that we form judgments using mechanisms something like SCM and EAM. When these mechanisms can generate either of two contrary judgments given a body of evidence, would those judgments ultimately be based on the same evidence in either case, or on different evidence?

There is an obvious reason to think they would rest on different evidence. Processing involves computation, which alters the contents of memory, by definition. And differences in memory create differences in evidence. Some philosophers hold that your total evidence just is the sum total of your experiential and memorial states (Lewis, 1996). If so, changes in memory are necessarily changes in evidence. Others hold propositional accounts, where your evidence is the set of propositions you know (Williamson, 2000), or that you know directly on the basis of experience (Maher, 1996), or something similar. On these propositional views, differences in memory are not necessarily differences in evidence. But they often will be, because the changes in memory are noticeable, or have noticeable effects.

Some defenders of Uniqueness won't themselves be open to this line of thinking. For example, whatever criterion for "same evidence" White has in mind, it doesn't seem to vindicate this objection. White (2005, 2013) repeatedly separates what evidence one has from the effects of "examining" that evidence. His arguments for Uniqueness focus on the state of the evidence prior to its examination. So either the evidence isn't affected by the examination process, or the changes made by the examination process aren't relevant to Uniqueness, as he intends it.[33]

But other defenders of Uniqueness, like Feldman (2007), would likely be more open to this line. Conee & Feldman (2004) defend a mentalist view of evidence

[33]For example, White summarizes his central argument in the following passage:

> [...] it is incoherent to suppose that a whole body of evidence could count both for and against a hypothesis. So then it is impossible that my examination of the evidence makes it rational for me to believe that Smith is guilty but also rational to believe instead that he is innocent. (White, 2005, 447)

If the evidence were different before and after the examination process, White would have to intend Uniqueness as a restricted claim, concerned only with the pre-examination evidence.

on which many subtle mental differences make a difference to what evidence one has. Nevertheless, I think there are good reasons to follow White in individuating evidence somewhat coarsely, at least for the purposes of discussing UNIQUENESS.

One reason is a threat of trivialization. The more fine-grained we get about evidence, the less substantive is the claim that only one response to that evidence is rational, since fewer and fewer responses become possible. To illustrate, if *every* mental difference is a difference in evidence, then two agents with the same evidence can only reach different conclusions to the extent that mental processing can be influenced by non-mental factors. Only a sub-mental or external difference could cause them to reach different conclusions. But, presumably, it's not these differences that UNIQUENESS is meant to address. Rather, differences in motivation or personal taste are meant to be ruled out as irrelevant to what conclusion one should draw from the evidence.[34]

Of course, there is a more moderate view one might take here. One might hold that only some mental differences make for a difference in evidence. For example, we might say that when two agents get different samples of memories using SCM, they have different evidence. But if they get the same sample, and differ only in the algorithms they use to test hypotheses against that sample, then they do not have different evidence. Suppose for example that one of them applies a significance test at a level of .05, while the other uses a level of .01. Then they can reach different conclusions from the same evidential state. So it's a substantive claim that only the latter's conclusions are rational.[35]

The main challenge I see here is that we need some principled reason for drawing the line where we have. Why is one of these mental differences a difference in evidence, but not the other? I have no argument that a principled rationale cannot be offered here. But absent any obvious rationale, we have a prima facie reason for drawing the line earlier. Our agents have the same evidence because they are working with the same store of memories, regardless of how they then go on to process it.

Another point to keep in mind is that evidence processing is often unconscious. As we saw earlier, Koriat is explicit about this in the case of SCM. But it's also generally acknowledged that much of our reasoning relies on unconscious, automatic processes (Sloman, 1996; Evans & Frankish, 2009). So propositional views of evidence should acknowledge that at least *some* processing doesn't alter our evidence. We don't normally know about the memories being sampled by SCM, or the estimates of informativeness being made by EAM. We don't even have beliefs about them. If we

[34]Thanks to Sergio Tenenbaum and Mike Titelbaum for their help clarifying this point.

[35]Thanks to Mike Titelbaum for pressing me to appreciate this point.

had some inkling about these processes going on within us, psychologists wouldn't have to work so hard to reverse-engineer them.

What about non-propositional views of evidence, like the view that your evidence is the sum total of your experiential and memorial states? There are memories and then there is "memory". The kind of memory used by unconscious processes like SCM might be aptly called *subpersonal* memory, because it does not store memories of the person, but of subprocesses within the person.[36] This kind of memory is less plausible as a candidate for inclusion in one's evidential state than consciously accessible memories. While your auditory or visual experiences may be part of your evidence, the computational states along the way to producing those experiences are not. The product of two matrices computed by the visual system in the course of edge-detection during low-level visual processing is not a part of your evidence, even if the experience it ultimately contributes to is. Likewise for the subpersonal memory states of a process like SCM, as it does the work of sampling, calculating likelihoods, and performing statistical tests. The "gist" of the whole process that consciousness catches a glimpse of might be considered part of your evidence, but not the nuts and bolts. In the relevant sense of 'memory', you can be in the same total memory state as your stochastic counterpart, who reaches a different conclusion via SCM because she draws a different sample of memories.[37]

In the end, though, it may not matter too much if defenders of Uniqueness still insist that differences in processing create differences in evidence. For the purposes of the applications pursued in §6, the picture I prefer:

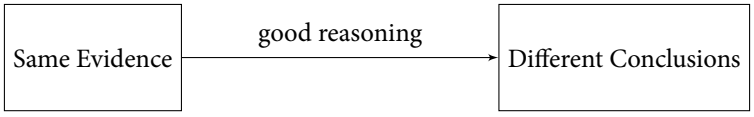

[36]The personal/subpersonal distinction comes from Dennett (1969). See Drayson (2014) for an overview.

[37]What about the consciously accessible gist, though? Could that be a relevant difference in your evidence, which explains why you and she are justified in drawing different conclusions?

According to the dominant view (Nelson & Narens, 1990; Dunlosky & Bjork, 2008), the consciously accessible feelings of plausibility and confidence that emerge from processes like SCM can be used by higher-level processing to control, correct, and redirect first-order processing. But the first-order processing is not based on those consciously accessible feelings. Your judgment that more than half your friends wear glasses is based on sampled memories of friends' faces, background knowledge about the frequency and fashionability of wearing spectacles, etc. The felt plausibility of that judgment is not the evidence on which the judgment is based, at least not initially. Instead, that felt plausibility helps determine whether you subsequently experience higher-order doubts about your initial judgment, whether you go back to reconsider the question, etc.

Much more would need to be said to fully satisfactorily address all the issues that arise here, but doing so would take us too far afield. Readers who remain unsatisfied can still benefit from much of the analysis to follow, however, especially the applications of §6.

may be as good as the alternative:

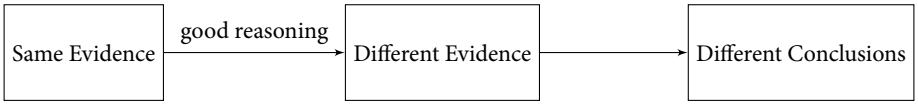

What's important is that the difference in evidence in the second picture derives solely from the reasoning process. If defenders of UNIQUENESS want to say the reasoning process itself provides new evidence of a kind, that may be fine. We agree that reasoning based on the same shared pool of evidence received "from the outside" can still rationally diverge.

7.2. **Justification and Rationality.** Let's turn to a second challenge. Are methods like SCM and EAM rational? Are beliefs formed by such methods justified?

There is an obvious reason to think they are not. SCM, EAM, and their like only "examine" some of your evidence. So they will often run afoul of Carnap's (1947) famous total evidence requirement. The degrees of belief generated by these methods won't generally match the quantity of support supplied by your total evidence.[38,39]

But satisfying the requirement of total evidence is largely hopeless for limited agents like us. Even many permissivists will acknowledge that our statistical evidence often dictates a fairly narrow range of degrees of belief. And given our limited abilities, we'll rarely hit that target exactly. We are pretty much never capable of examining and correctly responding to all the relevant evidence in our possession. For example, if I walk into a Las Vegas casino and pull up a chair at a blackjack table, I'll have credence 1/13 that the next card dealt will be an ace. But that's almost certainly not the exact degree to which my total evidence supports that proposition. I know it's against the casino's interest to deal aces to punters. Yet I also know that casino's are heavily regulated in Nevada, and that there are safeguards in place to prevent them stacking decks. On the other hand, I have heard of cases where casinos broke the rules. On still another hand, regulators have imposed additional constraints since then. And so on.

Maybe in this case my total evidence really does support credence 1/13 exactly. But if it does, it's just a fluke. When we do hit the target exactly, it's going to be a rare matter of luck. Presumably, the best we can hope for is to come close to what our

[38]It may be somewhat anachronistic to interpret Carnap's requirement as governing degrees of belief. As Carnap originally formulated the requirement, it governed "applications of inductive logic" and "degrees of confirmation". However, Carnap was concerned with applying inductive logic to settle one's betting odds, which traditionally have been closely associated with degrees of belief. Moreover, author's since have taken Carnap's requirement to govern degrees of belief: see Good (1967) and Neta (2008), for example.

[39]Thanks to an anonymous referee and Mike Titelbaum for their help clarifying several aspects of the discussion in this section.

evidence supports a lot of the time. So if the total evidence requirement really is a requirement of rationality, we are irrational a great deal of the time.

The natural thing to say, of course, is that the total evidence requirement is a requirement of ideal rationality, not of the kind of rationality we ordinarily attain and attribute to one another: what Titelbaum (2013) calls "everyday" rationality. I am skeptical of the notion of ideal rationality at play here. But for simplicity, let me grant for the sake of argument that the total evidence requirement is a requirement of ideal rationality, though only of ideal rationality. Then the question that remains is whether there is any residual reason to think methods like SCM and EAM fail to deliver rational or justified beliefs in the everyday sense, the one we are actually capable of attaining on a regular basis.

One line of thought is that we should process as much evidence as is available, so that our beliefs have the best shot at being accurate. The theorems of Good (1967) and Oddie (1997) might be taken to support this line. But for limited agents who can't process all their evidence, much less process it all correctly, the best way of maximizing accuracy may be to process as much evidence as they can process effectively. For example, Dallmann (manuscript) shows that there is a clear sense in which agents with limited memory best serve the aim of accuracy by processing just some of their evidence. (See Staffel manuscript for a different sort of result bearing on the gap between ideal and non-ideal agents, and for a broader discussion of the ideal/non-ideal gap in epistemology.)

In fact, in an important sense it can be clearly *ir*rational to do what ideal rationality supposedly requires. To borrow a particularly extreme example from Talbott (2005), consider the question whether the trillionth digit of π is odd or even. Given our limited processing abilities, the rational response to this question is to have credence 1/2 in each answer. But standard theories of ideal rationality require credence one in the even answer instead—the trillionth digit of π is a 2, as it turns out, an a priori truth deducible from our evidence. The only way you could have actually arrived at that answer, though, would have been to take an arbitrary guess. And if you had done so, you would have been unjustified in your belief, irrational.

The example illustrates another important distinction, supplemental to the ideal/everday distinction. Simon (1976) famously distinguishes between *substantive* rationality and *procedural* rationality. Considering only the substance of the question about π's trillionth digit, your evidence logically entails that even is the correct answer, so credence one in even is the rational attitude. But considering the *procedure* used to arrive at the judgment, that belief would not be rational in your case. Arbitrary

guessing is not a rational procedure. Indeed, given the procedures you have available to you, the only rational attitude in the procedural sense is credence 1/2.

The substantive/procedural distinction is important because it isolates a sense in which judgments can be rational even when they depart quite a bit from the ideal. Even when your credence is 1/2 instead of 1, it is still rational in the procedural sense. Indeed, it is more rational than a credence of 9/10 would be, even though 9/10 would be closer to the ideal, in the substantive sense. Likewise, even when two agents with a shared body of total evidence reach radically divergent conclusions via a method like SCM, one concluding P and the other concluding $\neg P$, they can be equally rational, in the procedural sense. Each has proceeded exactly as she should have, and holds the belief she ought to hold as a result.[40]

Some defenders of UNIQUENESS may respond that they are concerned solely with ideal rationality. Indeed, White (2013, 312) clarifies his (2005) statement of UNIQUENESS in terms of "full" rationality. But consider why questions of ideal rationality are interesting.

One reason for contemplating ideal rationality is to use it as a proxy for something else. We might contemplate the beliefs of ideally rational agents because they will track what we're really interested in, namely evidential support. If what we're really after is whether there is always a unique, objective fact about whether E supports P, then we might ask what an ideally rational agent would make of E's bearing on P.

In that case however, what we're really interested in is UNIVOCITY. And we've already seen that UNIVOCITY is well-supported by the instability argument. The question whether UNIQUENESS is also true is distinct, as are the implications for peer disagreement and related issues. As we saw in §6, we benefit our understanding by distinguishing these questions, and by attending to the ways they can come apart for everyday rational agents.

A second reason one might be interested in ideal rationality is as a guide to everyday rationality. For example, Titelbaum (2013) defends the project of studying idealized Bayesian models of rationality as "an important first step" in figuring out what everyday rationality requires. If we can first sort out what ideal rationality requires, then we can start going about figuring out how to come as close to that ideal as possible, given our limitations.

Yet here again, the present project proves important. Ideal agents align their credences with the support of their total evidence, we are supposing. But given our limited memory, time, and computational powers, we can't do that exactly and in every case. So how can we come close? Psychologists tell us that we have methods

[40]Relatedly, see Titelbaum's (2016) distinction between *prescriptive* and *evaluative* rationality.

like SCM and EAM at our disposal. And we've seen that our use of such methods illuminates normative phenomena like moderation and near agreement (§6). So even if Uniqueness is true for ideally rational agents, we should be interested in the ways it fails for everyday rationality. Doing so helps us understand the everyday phenomena, so that we can come closer to the rational ideal.

8. Conclusion

There is a difficult tension in the way we think about rationality and evidence. On the one hand, to hold a belief is to reject alternatives, and thus to treat one's own stance as superior. At the same time, we recognize that opinions vary and that it's often difficult—even impossible—to say who went wrong where. The former line of thought seems to push us towards an uncomfortably stern-looking objectivism, the latter towards an uncomfortably unstable-looking subjectivism.

A way out of this dilemma emerges when we embrace objectivity in the form of Univocity rather than Uniqueness, and subjectivity in the form of permissivism but not relativism. There are objective, non-relative facts about what the evidence says. But they are complex and difficult to discern, and our best methods for discerning them are variable. So it's no surprise that we're often unsure exactly which take on the evidence is right, yet confident that wildly different interpretations must be wrong.

In our pursuit of beliefs that align with the evidence, some variation is inevitable, and permitted. But the pursuit of perfect alignment is still obligatory. These subjective and objective considerations create the illusion that pulls us in two directions at once. The stable solution is to see each element for what it is, so that we can embrace both.

References

Anderson, Elizabeth. 2004. Uses of Value Judgments in Science: A General Argument, with Lessons from a Case Study of Feminist Research on Divorce. *Hypatia* 19(1): 1–24.

Ballantyne, Nathan & E. J. Coffman. 2011. Uniqueness, Evidence, and Rationality. *Philosophers' Imprint* .

Bartha, Paul & Christopher Hitchcock. 1999. No One Knows the Date or the Hour: An Unorthodox Application of Rev. Bayes' Theorem. *Philosophy of Science* 66(S): S339–S353.

Bröder, Arndt & Ben R. Newell. 2008. Challenging Some Common Beliefs: Empirical Work within the Adaptive Toolbox Metaphor. *Judgment and Decision Making* 3(3): 205–214.

Carnap, Rudolf. 1947. On the Application of Inductive Logic. *Philosophy and Phenomenological Research* 8(1): 133–148.

Christensen, David. 1999. Measuring Confirmation. *The Journal of Philosophy* 96(9): 437–461.

Christensen, David. 2007. Epistemology of Disagreement: The Good News. *The Philosophical Review* 116(2): 187–217.

Christensen, David. 2009. Disagreement as Evidence: The Epistemology of Controversy. *Philosophy Compass* 4(5): 756–767.

Cohen, Stewart. 2013. A Defense of the (Almost) Equal Weight View. In David Christensen & Jennifer Lackey (eds.), *The Epistemology of Disagreement: New Essays*, 98–120. Oxford: Oxford University Press.

Conee, Earl & Richard Feldman. 2004. *Evidentialism*. Oxford, England: Oxford University Press: Oxford University Press.

Dallmann, Justin. manuscript. When Obstinacy is a Better (Cognitive) Policy .

Dennett, Daniel C. 1969. *Content and Consciousness*. London: Routledge & Kegan Paul.

Douven, Igor. 2009. Uniqueness Revisited. *American Philosophical Quarterly* 46(4): 347–361.

Drayson, Zoe. 2014. The Personal/Subpersonal Distinction. *Philosophy Compass* 9(5): 338–346.

Dunlosky, John & Robert A. Bjork (eds.). 2008. *Handbook of Metamemory and Memory*. New York: Psychology Press.

Earman, John. 1992. *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. Cambridge: The MIT Press.

Eells, Ellery. 1990. *Bayesian problems of old evidence* 205–223. Minneapolis: University of Minnesota Press.

Elga, Adam. 2007. Reflection and Disagreement. *Noûs* 41(3): 478–502.

Evans, Jonathan St. B. T. & Keith Frankish (eds.). 2009. *In Two Minds: Dual Processes and Beyond*. New York: Oxford University Press.

Feldman, Richard. 2007. Reasonable Religious Disagreements. In Louise Antony (ed.), *Philosophers without gods: Meditations on atheism and the secular*, 194–214. New York.

de Finetti, Bruno. 1937. La Prévision: Ses Lois Logiques, Ses Sources Subjectives. *Annales de l'institut Henri Poincaré* 7(1): 1–68.

Firth, Roderick. 1978. Are Epistemic Concepts Reducible to Ethical Concepts? In Alvin I. Goldman & Jaegwon Kim (eds.), *Values and morals*, 215–229. Dordrecht: Springer Netherlands.

van Fraassen, Bas C. 1984. Belief and the Will. *The Journal of Philosophy* 81(5): 235–256.

Garber, Daniel. 1983. Old Evidence and Logical Omniscience in Bayesian Confirmation Theory. In *Testing scientific theories*, 99–131. Minneapolis, MN: University of Minnesota Press.

Gigerenzer, Gerd. 2001. The Adaptive Toolbox. In Gerd Gigerenzer & Reinhard Selten (eds.), *Bounded rationality: The adaptive toolbox*, 37–50. Cambridge.

Glymour, Clark. 1980. *Theory and Evidence*. Princeton, NJ: Princeton University Press.

Glynn, Luke, Kenny Easwaran, Christopher Hitchcock & Joel Velasco. forthcoming. Updating on the Credences of Others: Disagreement, Agreement, and Synergy. *Philosophers' Imprint* .

Good, I. J. 1967. On the Principle of Total Evidence. *Philosophy of Science* 17(4): 319–321.

Greaves, Hilary & David Wallace. 2006. Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility. *Mind* 115(459): 607–632.

Hausmann, Daniel & Damian Läge. 2008. Sequential Evidence Accumulation in Decision Making: The Individual Desired Level of Confidence Can Explain the Extent of Information Acquisitio. *Judgment and Decision Making* 3(3): 229–243.

Hawthorne, James. 2005. Degree-of-Belief and Degree-of-Support: Why Bayesians Need Both Notions. *Mind* 114(454): 277–320.

Hawthorne, James, Jürgen Landes, Christian Wallmann & Jon Williamson. 2015. The Principal Principle Implies the Principle of Indifference. *The British Journal for the Philosophy of Science* .

Horowitz, Sophie. 2014. Immoderately rational. *Philosophical Studies* 167(1): 41–56.

Howson, Colin. 1991. The 'Old Evidence' Problem. *The British Journal for the Philosophy of Science* 42(4): 547–555.

Howson, Colin & Peter Urbach. 2006. *Scientific Reasoning: The Bayesian Approach*. Peru, IL: Open Court.

James, William. 1897. The Will to Believe. In *The will to believe and other essays in popular philosophy*, 1–31. New York, London, and Bombay: Harvard University Press.

Jaynes, Edwin T. 1957. Information Theory and Statistical Mechanics. *Physical review* 106(4): 620–630.

Jeffrey, Richard C. 1983a. Bayesianism with a Human Face. In John Earman (ed.), *Testing scientific theories*, 133–156. Minneapolis, MN: University of Minnesota Press.

Jeffrey, Richard C. 1983b. *The Logic of Decision*. Chicago, IL: University of Chicago Press: The University of Chicago Press 2nd edn.

Joyce, James M. 1998. A Nonpragmatic Vindication of Probabilism. *Philosophy of Science* 63(4): 575–603.

Joyce, James M. 2009. Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief, 263–297. Dordrecht, Netherlands: Springer.

Kelly, Thomas. 2013. Evidence Can Be Permissive. In Matthias Steup, John Turri & Ernest Sosa (eds.), *Contemporary debates in epistemology*, 298–311. Blackwell.

Kopec, Matthew & Michael G. Titelbaum. 2016. The Uniqueness Thesis. *Philosophy Compass* 11(4): 189–200.

Koriat, Asher. 2011. Subjective confidence in perceptual judgments: A test of the self-consistency model. *Journal of Experimental Psychology: General* 140(1): 117–139.

Koriat, Asher. 2012. The Self-consistency Model of Subjective Confidence. *Psychological Review* 119(1): 80–113.

Koriat, Asher & Shiri Adiv. 2015. The Self-Consistency Theory of Subjective Confidence. In John Dunlosky & Sarah K. Tauber (eds.), *The oxford handbook of metamory*, Oxford University Press.

Lee, Michael D. & Tarrant D. R. Cummins. 2004. Evidence Accumulation in Decision Making: Unifying The "Take the Best" and The "Rational" Models. *Psychonomic Bulletin & Review* 11(2): 343–352.

Levi, Isaac. 1967. *Gambling With Truth*. New York, NY: Alfred A. Knopf.

Levi, Isaac. 1980. *The Enterprise of Knowledge: An Essay on Knowledge, Credal Probability, and Chance*. Cambridge, MA: The MIT Press.

Lewis, David. 1980. A Subjectivist's Guide to Objective Chance. In Richard C. Jeffrey (ed.), *Studies in inductive logic and probability*, 263–294. Berkeley, CA: University of California Press.

Lewis, David. 1996. Elusive Knowledge. *Australasian Journal of Philosophy* 74(4): 549–567.

Longino, Helen E. 1990. *Science as Social Knowledge* Values and Objectivity in Scientific Inquiry. Princeton University Press.

Longino, Helen E. 2002. *The Fate of Knowledge*. Princeton University Press.

Maher, Patrick. 1996. Subjective and Objective Confirmation. *Philosophy of Science* 63(2): 149–174.

Manley, David. forthcoming. On Being a Random Sample.

Matheson, Jonathan. 2011. The Case for Rational Uniqueness. *Logos & Episteme* 2(3): 359–373.

Meacham, Christopher J. G. 2008. Sleeping Beauty and the Dynamics of De Se Beliefs. *Philosophical Studies* 138(2): 245–269.

Meacham, Christopher J. G. 2010. Two Mistakes Regarding the Principal Principle. *The British Journal for the Philosophy of Science* 61(2): 407–431.

Meacham, Christopher J. G. 2014. Impermissive Bayesianism. *Erkenntnis* 79(6): 1185–1217.

Nelson, Thomas O. & Louis Narens. 1990. Metamemory: A Theoretical Framework and New Findings. In Gordon H. Bower (ed.), *The psychology of learning and motivation*, 125–174. San Diego, CA: Academic Press.

Neta, Ram. 2008. What Evidence Do You Have? *The British Journal for the Philosophy of Science* 59(1): 89–119.

Newell, Ben R., Patrick Collins & Michael D. Lee. 2007. Adjusting the Spanner: Testing an Evidence Accumulation Model of Decision Making. *roceedings of the 29th Annual Conference of the Cognitive Science Society* 533–538.

Newell, Ben R. & Michael D. Lee. 2011. The Right Tool for the Job? Comparing an Evidence Accumulation and a Naive Strategy Selection Model of Decision Making. *Journal of Behavioral Decision Making* 24(5): 456–481.

Niiniluoto, Ilkka. 1983. Novel Facts and Bayesianism. *British Journal for the Philosophy of Science* 34(4): 375–379.

Oddie, Graham. 1997. Conditionalization, Cogency, and Cognitive Value. *British Journal for the Philosophy of Sciencei* 48(4): 533–541.

van Overschelde, James P. 2008. Metacognition: Knowing About Knowing. In John Dunlosky & Robert A. Bjork (eds.), *Handbook of metamemory and memory*, New York, NY: Psychology Press.

Payne, John W. & James R. Bettman. 2004. Walking with the Scarecrow: The Information-processing Approach to Decision Research. In Derek J. Koehler & Nigel Harvey (eds.), *Blackwell handbook of judgment and decision making*, 110–131. Malden, MA: Blackwell.

Payne, John W., James R. Bettman & Eric J. Johnson. 1993. *The Adaptive Decision Maker*. New York, NY: Cambridge University Press.

Pettigrew, Richard. 2015. Accuracy, Risk, and the Principle of Indifference. *Philosophy and Phenomenological Research* .

Ramsey, Frank P. 1931. Truth and Probability, 156–198. London, England: Kegan, Paul, Trench, Trubner & Co.

Schoenfield, Miriam. 2014. Permission to Believe: Why Permissivism Is True and What It Tells Us About Irrelevant Influences on Belief. *Noûs* 48(2): 193–218.

Senor, Thomas D. manuscript. Disagreement, Evidence, and Perception .

Shah, Anuj K & Daniel M. Oppenheimer. 2008. Heuristics Made Easy: An Effort-Reduction Framework. *Psychological Bulletin* 134(2): 207–222.

Shimony, Abner. 1955. Coherence and the Axioms of Confirmation. *The Journal of Symbolic Logic* 20(1): 1–28.

Simon, Herbert A. 1976. From Substantive to Procedural Rationality. In T. J. Kastelein, S. K. Kuipers, W. A. Nijenhuis & G. R. Wagenaar (eds.), *25 Years of Economic Theory: Retrospect and Prospect*, 65–86. Leiden: Springer.

Sloman, Steven A. 1996. The Empirical Case for Two Systems of Reasoning. *Psychological Bulletin* 119(1): 3–22.

Söllner, Anke, Arndt Bröder, Andreas Glöckner & Tilmann Betsch. 2014. Single-Process Versus Multiple-Strategy Models of Decision Making. *Acta Psychologica* 146: 84–96.

Staffel, Julia. manuscript. Should I Pretend I'm Perfect? .

Talbott, William. 2001. Bayesian Epistemology. `http://plato.stanford.edu/entries/epistemology-bayesian/`.

Talbott, William J. 2005. Review of David Christensen's Putting Logic in Its Place. `https://ndpr.nd.edu/news/24874-putting-logic-in-its-place-formal-constraints-on-rational-belief/`.

Titelbaum, Michael G. 2013. *Quitting Certainties*. Oxford: Oxford University Press.

Titelbaum, Michael G. 2016. Continuing On. *Canadian Journal of Philosophy* 45(5–6): 670–691.

White, Roger. 2005. Epistemic Permissiveness. *Philosophical Perspectives* 19(1): 445–459.

White, Roger. 2013. Evidence Cannot Be Permissive. In Matthias Steup & John Turri (eds.), *Contemporary debates in epistemology*, 312–323. Blackwell.

Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford, England: Oxford University Press.