

SOLVING THE PROBLEM OF LOGICAL OMNISCIENCE

Sinan Dogramaci
The University of Texas at Austin

Abstract

This paper looks at three ways of addressing probabilism's implausible requirement of logical omniscience. The first and most common strategy says it's okay to require an *ideally rational* person to be logically omniscient. I argue that this view is indefensible on any interpretation of 'ideally rational'. The second strategy says probabilism should be formulated not in terms of logically possible worlds but in terms of *doxastically* possible worlds, ways you think the world might be. I argue that, on the interpretation of this approach that lifts the requirement of certainty in all logical truths, the view becomes vacuous, issuing no requirements on rational believers at all. Finally, I develop and endorse a new solution to the problem. This view proposes dynamic norms for reasoning with credences. The solution is based on an old proposal of Ian Hacking's that says you're required to be sensitive to logical facts only when you know they are logical facts.

I. The Problem

What is the problem of logical omniscience? The problem as I'll be concerned with it here is a problem for the view known as *probabilism*.

Probabilism is a necessary condition on rational credence. It says that, necessarily, a person's credences are rational only if they obey the axioms (and consequent laws) of probability theory.

We'll formulate the axioms in the following standard way:

(A1 Logicality): the probability of any logical truth is 1.

(A2 Non-negativity): every proposition has a non-negative probability.

(A3 Additivity): the probability of a disjunction with logically exclusive disjuncts is the sum of the probabilities of the disjuncts.

Some notable laws that follow from this include:

(L-Entailment): if q is a logical consequence of p , the probability of q is no lower than the probability of p .

(L-Partition): propositions in a logical partition receive probabilities summing to 1.

Probabilism thus requires any rational person to be extremely sensitive to logical truth and consequence. Just by the first axiom, probabilism has it that the rational credence in any logical truth is *always* 1, the maximal credence, normally interpreted as *certainty*. This is the problem of logical omniscience. It just doesn't seem true that any rational person has to always be certain of every logical truth, or as confident of a consequence as she is of any proposition that entails it. Was it irrational to be less than certain of Fermat's Theorem in the past, or even to have less confidence in it than in the Peano Axioms? Is it irrational even now? (The first "proof" Andrew Wiles publicly presented had an error in it.) Other manifestations of the problem are the requirement, via the L-Entailment law, to give logically equivalent propositions exactly the same credence. Was it irrational to give the Axiom of Choice and the Well-Ordering Theorem and Zorn's lemma distinct credences before we proved their equivalence (given ZF)? Intuitively, it is *possible* for *someone* at *some time* to have *rationally* given *some* pair of equivalent propositions distinct credences.

While my intuitions tell strongly against the view that rationality requires such "logical omniscience", some other people's intuitions favor the requirement. Some simply intuit that you must accept any logical consequence of what you know. A more modest sounding intuition that many more people have is that you at least must accept a consequence that follows by a "simple" deductive step, such as a step by Modus Ponens. (See, e.g., Steinberger (forthcoming, pp.10–1) and his references there to Broome (2013) and MacFarlane (2004).) Though that latter intuition sounds more modest, it still entails that anyone who fails to accept all logical consequences of what they know is committing some failure of rationality. This is because, as we know, we can fashion a deduction system that licenses only simple steps and that is "complete", i.e. every logical consequence (I mean of first-order logic) can be reached by a sequence of such steps. So, if it's intuitive that any "simple step" is required, then the perfectly rational person (one who is guilty of no failure of rationality) must be logically omniscient, and so probabilism's requirement on credences ends up supported by intuition, even the more "modest" intuition.

There are not only intuitions but also arguments in favor of logical omniscience as a rational requirement. In particular, there is an argument

in favor of probabilism that's very popular right now, the so-called "accuracy" argument (Joyce (1998), Pettigrew (2016)). This argument aims to show that anyone who violates probabilism has less accurate attitudes than someone who obeys probabilism. Alongside the accuracy argument, another very popular view among Bayesians is that logical omniscience is a correct requirement, so long as we understand that it concerns the *ideally rational* person (Talbot (2008), Christensen (2004), Smithies (2015)).

Some people who appreciate the problem with probabilism's requirement of logical omniscience have looked for better theories that remove the logical omniscience requirement while, hopefully, saving what's good in probabilism.

In a recent paper, Robbie Williams (forthcoming) motivates and develops such a view, a view he calls doxastic probabilism. The idea of this view, which I'll explain more below, is to replace logical possibilities with a person's "doxastic" possibilities in the formulation of probabilism.

And much earlier Ian Hacking (1967) offered a suggestion for fixing up the logical omniscience requirement of probabilism. He suggested that you should obey axiom A1 only when you *know* that the proposition is a logical truth, and you should obey A3 when you *know* that the two disjuncts are mutually exclusive.

My plan for this paper:

In section II, I'll say why the accuracy argument doesn't support probabilism.

In section III, I'll criticize interpretations of probabilism as a condition only on *ideal* rationality, and in doing so I'll critically discuss the "modest" intuition I mentioned some people have that seems to support probabilism.

In section IV, I'll present Williams's doxastic probabilism and its motivations, and I'll explain why I don't favor taking that approach.

In section V, I'll sketch a new solution to the problem that builds on Hacking's view. I'll incorporate Hacking's idea, together with another important idea I'll take from Gilbert Harman (1986), into a larger framework that says credences in logical truths must obey "dynamic" norms of rationality, that is, norms that concern how we revise our credences by reasoning over time. So, for example, while I reject probabilism's requirement that we *always* be certain of every logical truth, I will endorse, among other things, some norms that tell us to revise our credences as we acquire proofs that give us knowledge that particular propositions are logical truths.

II. The Accuracy Argument

The argument that any violator of probabilism is less accurate than some follower of probabilism has valuable contributions to offer epistemology. But, on its own, it just doesn't support probabilism. It especially

clearly doesn't support probabilism's logical omniscience requirement on rationality.

The disconnect between the argument's conclusion and probabilism is due to a simple and familiar point: a rational attitude is a different thing from an accurate attitude or even a provably accurate attitude. The point is familiar and clear when it's made about full belief in traditional epistemology. It's just implausible that, necessarily, a rational person must believe, much less know, every logical truth. If you are unaware of any proof of Fermat's theorem, believing it is *not* rational. (And, as the next section will discuss, it's not a requirement on being rational that you *produce the proof*.) For anyone who feels it's objectionable for a theory of rationality to require full belief in all logical truths, it will be no reassurance at all to hear that being logically omniscient gives you provably accurate beliefs. That was not unclear. Likewise, for anyone who feels it's a problem for a theory of rationality to require full *credence* in every logical truth, it's simply no reassurance to hear the accuracy argument for probabilism.

That's my dismissal of the accuracy argument as a positive argument for probabilism and especially for logical omniscience. It's brief, I know. But what can be said against it? If the accuracy argument and probabilism, with its requirement of logical omniscience, are to be defended against my dismissal, the defenders must say more. They will, I expect, turn to the notion of ideal rationality. So let's examine that now.

III. Probabilism and Ideal Rationality

Some people say that *ideal rationality* requires a person to have both full beliefs and maximal credences in all logical truths. The idea is widely popular—Talbot (2008) says in his Stanford entry on Bayesian Epistemology that “most Bayesians maintain the assumption of logical omniscience and treat it as an ideal to which human beings can only more or less approximate” (sec. 6.1.A)—but I will focus on the two extensive recent defenses by David Christensen (2004) and Declan Smithies (2015).

What is “ideal rationality”? There are two readings we can distinguish.

On a first reading, ideal rationality is the subject matter of a theory that *idealizes* in the sense of abstracting away from, or ignoring, certain real world constraints. It's the sense of “idealize” in which we sometimes ignore friction in a physics calculation. What might the epistemologist similarly ignore? Candidates include the limits of human memory, the limited amount of time we have to think and to come up with proofs and arguments or (relatedly) the limited speed at which we think or perform calculations, or the limited stock of concepts we possess or (relatedly) the limited range of hypotheses we can take under consideration at any time.

In defense of idealizing in this way, I see only one possible point: it has the advantage of making our work easier. But that's only a pragmatic defense. It does nothing to show, for any absurd feature that might turn up in the theory that results from such easy work, that the absurdity should be accepted as part of the truth. So the first reading doesn't really get anywhere with the problem of logical omniscience. (Smithies also distinguishes and sets aside this first reading, p.2774.)

The second, more promising, reading of the notion of ideal rationality is the notion of the *ideally*, or *perfectly*, rational person. It's the notion of a person who obeys every instance of every norm of rationality, without a single failure or without falling short in any way. If any norm calls on us to do something to the "maximum", the ideally rational person does that thing to that maximum. (Or, even if there is no "maximum" (like there is no largest real number less than 2), if a norm says that the more you do something the better, then the notion of the ideally rational person here may be like a limit case of doing more and more of that thing.)

Is there any reason to think the ideally rational thinker, ideal in our *second* sense, is ideal in that *first* sense, that is, has no limits on things like memory, speed, computational power, or conceptual repertoire? No, I see no reason to believe that the thinker who follows the norms of rationality to the fullest, without any error, will have unlimited memory, speed, computing power, concepts, or any such thing. This claim (that the second sense of 'ideal' entails the first sense) would require some argument that there is a norm of *rationality* that prescribes having a better memory or prescribes reasoning faster or crunching through more proofs or calculations. But any norm that says something like "Calculate *faster!*" is not a norm of rationality. It leads to a theory of the ideally *fast* thinker or calculator, not the ideally *rational* thinker. (I'm not denying there can be some kind of ideal that requires your being as fast as possible at some job. Brian Knab points out to me that it might be true of *the ideal runner* that they're "maximally" fast. I'm just saying nothing like that is true of *rationality*.)

I have no *argument* in defense of my claim about what rationality is (not) about; that is, I don't rest that claim on firmer premises. I'm just appealing to your understanding of our concept of rationality. But I think, once pointed out, it's an intuitive claim. So, we must avoid, on pain of changing the subject, assuming that the ideally rational thinker, in our present second sense, has unlimited memory, time or speed, computational power, or any such thing.

In fact, I can't see how there is *any* plausible norm that tells us we ought, in *any* sense of 'ought', to do every calculation or computation possible. It might be true that the ideally fast and thorough calculator produces and performs every calculation super fast, but that doesn't mean that anyone or anything *ought to* do that. Consider the ideally fast dancer or singer, someone who dances or sings "maximally" fast. That's not something anyone *ought to do!* Or consider the ideally fast, accurate, and relentless counter of the

exact number of jellybeans in every jellybean jar. That's again not something anyone *ought* to do. So, I think that unwisely assuming the ideally rational thinker in the second sense is ideal in the first sense would actually change the subject not just away from rationality, but away from normativity entirely. (While I'm sure that being rational doesn't require producing every proof, I'm a bit less confident in this paragraph's further claim that it's false that we, in any sense, ought produce every proof. That claim seems related to the claim that it's false we ought to have a belief in (much less knowledge of) every truth including every bit of trivia. I believe the latter claim too, but the issue's controversial.)

Which of these two notions of ideality do Bayesians actually have in mind when they defend probabilism, and logical omniscience in particular, as a norm of "ideally rational belief"? Smithies explicitly opts for the second (p.2774). Christensen doesn't explicitly draw the distinction we've just drawn, but I'm confident the correct view to assign him is this second interpretation. He often talks of "the standard of absolute rational perfection" (p.153), and this second reading often makes the best sense of his defense of probabilism.

So, now, why think that an ideally rational thinker (in the second sense) will be logically omniscient? Some people (among them Christensen cites Ian Hacking, Richard Foley, and Philip Kitcher) worry that requiring logical omniscience is as implausible as saying the ideally rational person is omniscient period, including "factually" omniscient (i.e. knows all non-logical truths). Christensen responds to that worry with an example (pp.154–5) designed to show that rationality does require us to believe logical consequences (of what we rationally accept) in a way it does not require us to believe all non-logical truths. In his example one person, Cherry, is rationally confident that (P1) anyone near a Grizzly cub in the wild is in danger, while another person, Kelly, is likewise rationally confident of P1 but also of something Cherry has no confidence in, that (P2) Kelly is herself near a Grizzly cub in the wild. Cherry and Kelly both fail to be confident of the logical consequence of P1 and P2 that Kelly is in danger. But, Christensen points out, obviously this is a serious rational failure on Kelly's part and no rational failure at all on Cherry's part.

Christensen doesn't just use the example to *defend* against the worry (that requiring logical omniscience is as bad as requiring omniscience period). He suggests it *positively supports* the claim that logical omniscience is a rational requirement. After presenting the case of Kelly's mismanaged response to her evidence, he says that "logical omniscience emerges naturally as the limiting case of one of the basic ingredients of good thinking. . . . Eliminating this sort of mistake [failing to respect logical relationships as Kelly did] yields, in the limit, logical omniscience." (p.156) Smithies endorses Christensen's use of the Kelly example and draws the same conclusion that Christensen does (sec. 2). The thought Christensen and Smithies have here seems to be an elaboration of the "modest" intuition that I brought up at the start of this

paper, an intuition some people have and that can be leveraged to support logical omniscience as a rational requirement: according to this intuition that many people have, we're required to make any simple logical inference, and since every consequence (at least of first-order logic) follows by some sequence of simple steps, we are required to be logically omniscient.

I think this is not right. I think Christensen and Smithies generalize from Kelly's case to the requirement for full logical omniscience too quickly. There are epistemically significant differences between the particular logical consequence we intuitively feel Kelly must accept and other logical consequences like Fermat's theorem.

The important difference, I suggest, is that Kelly is, or ought to be, *interested* in whether she's in danger, while (we can suppose) she currently doesn't have, and it's not the case she now ought to have, any interest in Fermat's theorem. In general, what you ought to believe depends on what you are, or ought to be, interested in. The first underlined 'ought' there expresses epistemic rationality, the traditional necessary condition of knowledge. I won't argue over the correct interpretation of the second underlined 'ought', since a pragmatic or prudential reading offers us at least one attractive option and we don't have to pick here.

The proposal I want to endorse is due to Gilbert Harman (1986). He famously observed that there is a negative norm on good thinking: *avoid clutter!* (p.12) Don't infer logical consequences when it would only clutter your mind (e.g. inferring disjunctions like that birds fly or the moon is cheese). The clutter avoidance norm explains why, as I said, ideality in the present second sense doesn't entail ideality in the earlier first sense. Don't waste time producing proofs or calculations of clutter either! Furthermore, there's an important caveat Harman immediately added after he introduced the clutter avoidance norm, a caveat that's often missed or forgotten but is important and relevant to the Kelly example: if you are, or ought to be, interested in something you recognize is a consequence (of what you believe), then you are *required* to believe it. (I'm interpreting—correctly I believe—remarks that Harman makes on p.15, crediting Robert Stalnaker; see also p.55. Stewart Cohen (forthcoming) also defends the view that *attention*—which I take to be the same as interest here—is what converts an epistemic permission into a requirement, though Cohen doesn't also make the claim for things you *ought* to attend to.) With this caveat included, Harman's full proposal offers a more plausible view than the view that an ideally rational person, one who obeys every instance of every rational norm, is logically omniscient. Kelly ought to be interested in the proposition that she's in danger, and her acceptance of P1 and P2 means she recognizes that that proposition follows from things she knows. Since Harman's view fully explains the Kelly example, the example is no help to the case for, or in defense of, logical omniscience. (Friedman (forthcoming) gives a valuable recent examination of Harman's clutter avoidance principle. She doesn't reject the principle, but she raises

important critical points about it by demonstrating some surprisingly significant implications it carries. I'm unable now to address the points she raises.)

The Harman proposal is attractive also because it makes sense of that intuition some of us have that logical omniscience and "factual" omniscience ought to be treated analogously in the theory of rationality. The Harman proposal, on my favored interpretation, says that there are things you are in a position to and *merely permitted* to know or rationally believe, but there is no obligation, and there is some reason (namely, clutter avoidance) to not form the belief. But this is equally true for beliefs in logical consequences and for beliefs you are in a position to perceptually or introspectively justifiably adopt. I'm in a position to form many beliefs about my environment or my mind by meditating on the evidence constituted by my current experience, but I don't have to form those beliefs to be rational. But if I do, or if I ought to, take an interest in certain experiences, maybe while walking along a busy sidewalk, then I'm in a situation like Kelly's—then it is irrational to not believe these things I was previously merely permitted to (was in a position to justifiably) believe.

To make my position clear, let me put down an official statement of the proposal from Harman as I interpret it and intend to endorse it.

(Harman's Interest Proposal): If you have no interest, and ought have no interest, in p , then you're not rationally required to believe p (though you may be rationally permitted). If you do or ought to have an interest in p , and p is a *simple* consequence or a *recognized* consequence of your (undefeated) knowledge, then you are required to believe p .

I'll leave 'simple consequence' undefined. The intuitive example is something that follows by a single application of Modus Ponens; a full definition might simply enumerate a list of such familiar rules.

We can now also criticize another example Christensen brings up to support his claim that ideal rationality "in the limit" requires logical omniscience. He suggests that ideal rationality is analogous to ideal chess strategy. And he argues that we should concede that ideal chess play is beyond any living human's abilities. We would of course hesitate to assert that Garry Kasparov is not a good chess player, but we can admit that he is not an *ideally* good chess player. Christensen rightly points out that we easily understand the sense in which he, or any real person, falls short of the ideal. These claims about ideal chess play sound reasonable, but I claim they are not a good analogy to ideal rationality, especially once we consider things in light of Harman's proposal. A chess player is *forced* to make some move on each turn, but a believer is not similarly forced, so there is room to say an ideally rational believer just doesn't "make any move", doesn't form any attitude, in many cases, especially when it would result in clutter. So, the chess

analogy doesn't support the case for logical omniscience. (Cohen (forthcoming) suggests a similar diagnosis, though he is not discussing Christensen's chess example in particular.)

Smithies adds a further argument for logical omniscience as a norm on ideal rationality. Smithies's case for requiring logical omniscience can be responded to in a way similar to how I've responded to Christensen's already.

In order to make his case for logical omniscience as a requirement of rationality, Smithies devotes most of his effort to supporting the idea that we all have justification, specifically what he calls *propositional justification*, for all logical truths. (His defense involves arguing that the logical facts themselves, not any experience or mental state of ours, constitute our propositional justification for the logical truths, and therefore constitute propositional justification for logical omniscience.) From there, he then goes on to make the claim that ideal agents hold *doxastically justified* beliefs in all those logical truths that we enjoy propositional justification for. (Ordinary humans, though, are of course incapable of believing all the infinitely many things they have propositional justification to believe, and Smithies also argues that, even for many of the logical truths we can believe, we non-ideal humans are incapable of enjoying doxastic justification for them.)

My response to Smithies is to question his move from "we have propositional justification to believe . . ." to "an ideally rational person will believe . . .". This is a critical step he takes, but does not argue for.¹ The step seems to me to be a mistake, and avoiding it avoids the problem of logical omniscience. Again, we should say there are many things we have mere permission, without obligation, to believe. This should be the situation with most of the logical truths Smithies says we have propositional justification to believe (the ones we have and ought have no interest in).

Do we have propositional justification to believe every logical truth? I think it's a terminological question, since 'propositional justification' is a term of art. We might define 'propositional justification' in ordinary terms as "what we have sufficient support for rationally believing, whether or not we do rationally believe it". That doesn't leave me 100% comfortable saying we have propositional justification for logical omniscience, but I'm not closed to the view. Schoenfield (2012) usefully distinguishes (a) what we have support for (given the evidence), and (b) what we rationally ought to believe, i.e. are required to believe. I'm not 100% comfortable saying that (a) applies to all logical truths and consequences, even while (b) doesn't—but I'll leave the position open.

In any case, it doesn't matter. The strategy of saying that we have propositional justification for logical omniscience does not save probabilism, which says an (ideally) rational person *will be* logically omniscient. This is explicit in the formulations of probabilism by Bradley (2015, sec. 3.1), Briggs (2009, p.60), Christensen (2004, pp.viii, 106–7), Joyce (1998, pp.576, 580), Hájek (2008, p.229), Meacham (2014, p.1186), Pettigrew (2016, p.17), Titelbaum

(ms, sec. 2.2), Weisberg (2011, sec. 3), and every formulation of probabilism I've ever seen. I have never seen a Bayesian endorse probabilism or logical omniscience by putting it in terms of a *permission*. However we want to define technical terms like 'propositional' and 'doxastic justification', the thing that's important to me is just this: *if* we have permission to be logically omniscient, it is a *mere* permission. And that makes probabilism false. Wherever you want to fit 'propositional justification' into this position is fine with me.

I'll incorporate Harman's proposal (that we have mere permissions that our interests can convert into requirements) as an important part of the positive proposal I'll endorse in the final section of this paper. But first, the next section looks at a different approach to the problem of logical omniscience.

IV. Replacing Logical Possibility with Doxastic Possibility

In a valuable recent paper, Robbie Williams (forthcoming) motivates and presents a new view to replace traditional probabilism. His motivation is an instance of the more general problem of logical omniscience. He's concerned with instances of probabilism's implausible demand that our credences in all logical truths always be 1. But the instances that Williams focuses on are interestingly different from the ones I started this paper with. He does not have in mind remote, hard-to-prove logical truths or consequences like Fermat's theorem. Rather, to motivate his project, he argues against probabilism's requirement that we must be certain even of *elementary* logical truths, like instances of $P \vee \sim P$. With this as his motivation, he proposes a view he calls *doxastic probabilism*. I want to examine doxastic probabilism because the view very naturally looks like it promises to overcome the whole problem of logical omniscience, including our worries about the demand to be certain of "remote" logical truths like Fermat's theorem.² I'll present Williams's motivation, present his view, and raise a worry about the view.

Williams's motivation is this. Logical paradoxes, like the Liar paradox, have generated different reasonable, or at least non-crazy, responses from philosophers. One such response declares that propositions of the form $P \vee \sim P$ are not all logical truths and that some paradox-generating instances of it should receive a credence of 0 (e.g. Field (2003)). Other responses leave the status of all instances of $P \vee \sim P$ in tact as logical truths and give them a credence of 1. Williams says that it could be perfectly reasonable to be agnostic about which of these philosophical approaches is on the right track. Just as a probabilistic framework can recommend a middling credence in the face of uncertainty over some conventionally contingent matter like a proposition stating the location of, say, a stolen car (has it left the state?), so

likewise, Williams says, a sensible probabilistic framework should be able to also recommend a middling credence in the face of uncertainty over where the logical truth lies. Since this isn't allowed by traditional probabilism, which Williams calls *logical probabilism*, we need a new view.

The new view that Williams proposes, doxastic probabilism, is achieved by replacing the role of *logical possibilities* in the standard view with something he calls *doxastic possibilities*. (See his sec. 7.) To understand the resulting view, it's thus obviously central to understand what this new technical notion of a doxastic possibility is. But Williams actually doesn't define the notion. He says he will "[t]ake as a working primitive the notion of an agent's *doxastic space*, the set of worlds that are *doxastically possible* for her." However, he does put a constraint on the notion. He says he will "[a]ssume the following: For [person] z , (it must be that p) iff (p is true at each [world] w doxastically possible for z)." (p.6, underlining, square bracketed words, and punctuation, added for readability here.) So, Williams gives us a biconditional linking doxastic possibilities with certain epistemic modals, 'must' claims. I take it we can also restate the constraint using 'might' claims. That is, I take it the following biconditional holds too: for a person, (it might be that p) iff (p is true at some world that is a doxastic possibility for her). I also take it we could re-organize things to isolate the notion of doxastic possibility on one side of the biconditional: a world w is a doxastic possibility for a person iff it might be true for her that w is the actual world.

It will be important to my worry that we correctly understand the laxity of the conditions for correctly attributing attitudes using these epistemic modals. For some people, at some times, what might be true includes logical impossibilities and inconsistent combinations of logical possibilities (e.g., P together with $\sim P$; or P , $\sim Q$, and if- P -then- Q all together). There are no apparent prohibitions on what might be true for someone, certainly no apparent logical restrictions.

We can now ask exactly what requirements doxastic probabilism puts on a rational person's credences. What we'd like to know is how to convert the traditional axioms of probability into norms stated using doxastic worlds, since doxastic probabilism will then just be the requirement that a rational person's credences obey those norms (and all their consequent requirements).

But what exactly do the re-written axioms look like? There's no need to revise the non-negativity axiom, but the others will need to be revised.

The original logicity axiom was this (with some emphasis added):

(A1 Logicity): the probability of *any logical truth* is 1.

What is the new rule for doxastic probabilism? It will be this:

(D1 Doxastic logicity): the rational credence in *anything that must be true for you* is 1.

Here is my worry about this. I don't see how anyone could clearly violate this rule. Is anyone rationally required to be certain of anything on this view? Imagine some possible examples. Suppose I've invited my student to consider Descartes's proof: you are a thinking and doubting thing, and therefore you exist. And suppose my student is certain of the premise, but remains uncertain of the conclusion. This student is off the hook, so long as *for him*, it *might be* that he doesn't exist, even while he remains certain he is thinking and doubting (that he exists!). Or suppose we're in logic class, and we've proved a completeness theorem for propositional logic, but a student is withholding her certainty. On doxastic probabilism, there is nothing irrational about being uncertain even after fully understanding a correct proof and being certain of all its premises, because the student can always say, and can rightly say, for her, it *might be* that the conclusion is false. We all always have the potential to *rightly* say this because doxastic worlds are constrained only by these epistemic modals, and it is up to each person what might be true for them. By 'up to' I don't mean we choose (maybe we can, I don't know). I mean that for different people, there can be different things that might be true for them, and we've been given no rational constraints on what might be true for you. Certainly, Williams suggests no such rational constraints, and I don't see how there could be any if our ultimate *motivation* is to rationally permit violations of the "correct" logic. Williams's own original motivation was an instance of excluded middle, and Williams says he thinks such an instance *might be* false even if classical logic is the "correct" logic. So, any proposition, as far as I can see, could be such that it might be false for you or me or whoever, and we have to say there is nothing irrational about it.

Next, we want to know how the additivity axiom goes under doxastic probabilism. The simplest suggestion would seem to be this:

(D3* a first pass): if, for you, no doxastic world makes true both disjuncts of a disjunction, the rational credence in the disjunction is the sum of your rational credences in its disjuncts.

But, as Williams observes in another paper (Williams 2016), for some familiar ways of rejecting classical logic, or rejecting classical semantics, D3* is not an attractive principle. There is the supervaluationist who thinks the disjunction, "Patchy is red or not red", is true while each disjunct is not true, because Patchy is a borderline shade of red. Such a supervaluationist will want to reject D3* as well as anything too closely resembling traditional additivity. (They will instead want a sort of principle often called 'subadditivity'.) So, D3* doesn't help us see what principle a doxastic probabilist should propose.

A better option for the doxastic probabilist might be to reformulate D3* so that it isn't explicitly about disjunctions. An idea for a way to do that can be drawn from Williams (2016). Instead of talking of disjunctions as such, we can talk about propositions that are partitioned by other propositions, where

the partitioning is not classical logical partitioning. Here is a proposal—I'm drawing this from a more general technical discussion in Williams (2016, sec. 3)—for a rule for how doxastic probabilists might, in effect, handle disjunctions. Think of r as serving as the disjunction of p and q in the following rule.

(D3 Doxastic additivity): if, supposing r is true then one of p or q must be true for you, and, (still supposing r is true) it's not the case that both of p and q might be true for you, then, if you're rational, your credence in r is the sum of your credences in p and in q .

My worry, though, is that, as with D1 earlier, D3 cannot be violated. To see this, let's consider when it is supposed to be violable.

We'll get cases where D3 and A3 issue differing permissions in cases where two "disjuncts", p and q , are both true in no logically possible world but are true in a doxastically possible world (for you). While A3 would then require you to give the disjunction the sum of the disjuncts' credences, D3 will permit you to give a lower credence. (This is one way of letting $P \vee \sim P$ have a credence less than 1.) My worry, though, is that D3 is permissive here because it is *always* permissive. Any case where it would initially *appear* someone is violating it will be ultimately better described, and *correctly* described, as a case where they are not violating it. Suppose I initially appear to violate D3 by saying I think there's half a chance of rain tomorrow and half a chance of snow and I think there's three quarters of a chance of either rain or snow. Any such case will be better described as one where I turn out to think it might both rain and snow, and thus there are simply more doxastic possibilities (dreamt of in my philosophy) than it first appeared when it appeared I was violating D3. So, I worry D3 will *never* require anyone to give a disjunction a credence as high as the sum of its disjuncts.

The source of the general problem here is this. The only grounds we've been given for attributions of doxastic possibilities are attributions using epistemic modals. Any *apparent* violation of this axiom (by giving the disjunction less credence than the sum of the disjuncts) will itself be sufficient reason for an attribution using an epistemic modal, a 'might' claim that the disjuncts are both true. And that attributed 'might' claim will amount to an attribution of doxastic possibilities where both disjuncts are true. The problem arises here because of a feature of doxastic probabilism that Williams endorses when he says: "the true rational constraints, I contend, govern the agent's partial beliefs together with their modal attitudes—what is possible, and what is impossible, by their lights." (p.7) What makes this problematic, to my mind, is that between these two moving pieces that the rational constraints always jointly govern, it will always be an option, and always make most sense, and thus always be *correct*, to attribute attitudes expressed by modals such that the partial beliefs *don't* violate the probabilistic constraints. I don't have an

argument for this claim that my objection rests on, but I think it is highly plausible: it is never correct to attribute irrationality when the alternative option carries no cost at all but the attribution of a doxastic possibility.

Reasons to deny a view are, like all reasons, defeasible, so my objection comes with a constructive suggestion for those who disagree with me. My constructive suggestion would be: if you wanted to stick with doxastic probabilism in the face of my worry, I suggest the way to defend it would have to be to develop some plausible constraints on the rationality of what doxastic possibilities a person has, where these constraints are wholly independent of probabilism and its motivations. My view is that this isn't the most promising approach to salvaging probabilism, but of course this depends on a comparison with the alternatives, including the one I want to advocate below.³

I talked about whether you could violate D3 by giving a disjunction too *low* a credence, but, to indicate the generality of my worry, let me also quickly mention the alternative. Does doxastic probabilism issue a requirement that we not give a disjunction a credence *higher* than the sum of its disjuncts? No, D3 doesn't prohibit this either. Williams wants doxastic probabilism to let us give credence to any solution to the logical paradoxes. It is designed to accommodate even someone like the supervaluationist already mentioned earlier, someone who suspects it might be that $P \vee \sim P$ is true even though each disjunct receives zero credence: for the disjunction could simply be one of my doxastic possibilities while neither disjunct is. In general, the motivations for doxastic probabilism leave no restriction on what violations of even elementary logic it's reasonable to entertain.

(For a very different sort of objection to a doxastic probabilistic view (a view suggested by Kenny Easwaran earlier, but very similar to Williams's view), see Elga & Rayo (ms.), fn. 22.)

V. Developing a Dynamics for Rational Credence

Williams proposes that probabilistic norms be anchored to the attitudes we express and attribute using the epistemic modals 'might' and 'must', and my worry has been that this anchor doesn't catch any firm ground. A similar but crucially different approach anchors probabilistic norms to *knowledge* instead of to those 'might' and 'must' attitudes. (Williams is explicit that this "knowledge" view is not his view. On p.9, last two paragraphs, he says what "must" be for a person is not the same as what they know.) Such a knowledge-anchored view was first proposed by Ian Hacking in an important paper on the problem of logical omniscience from 1967. My experience is that some formal epistemologists will say in conversation that they like Hacking's proposal, but no one to my knowledge has tried to develop it or has even endorsed the view in print. I'll describe Hacking's proposal and

then I'll propose and try to support a new view, a more general "dynamic" framework (I'll explain what I mean by 'static'/'dynamic'), one that preserves and incorporates Hacking's insights. (In an article that is a companion paper to this one, Dogramaci (2018), I develop at a bit more length the positive view that I will begin to outline here. There I also do a bit to discuss this view in the context of some empirical psychology.)

I'm going to combine Hacking's idea (to tie credal norms to knowledge) with another important idea, the one we already heard from by Gilbert Harman: you only acquire rational requirements if you are, or ought to be, *interested in* the propositions in question. Or, other formulations I take to say the same thing: . . . only if you do, or ought to, *consider* the propositions, or turn your *attention* to them. (Harman uses "considering" when he first brings up the point (p.15) but then switches to talking of when you "have or ought to have an interest" in the propositions. Steinberger (forthcoming) who defends and applies Harman's idea, sticks with 'considers'. Broome (2013) puts Harman's idea in the equally reasonable terms of 'caring', pp.157–8. As mentioned, Cohen (forthcoming) uses 'attention'.) Harman's "considering"/"interest" restrictions make a very important positive difference to the plausibility of the theory that results. (For brevity and readability, below I'll just say 'you consider' instead of 'you do or ought consider'.)

I'll also amend Hacking's proposal in the following way: I'll say the correct norms depend not on what the agent knows, but on what they are *in a position to* know, or for brevity I'll talk of what we 'can' know. This notion of what you are in a position to, or can, know is widely used in epistemology, but there's no obvious analysis, so I'll rely on familiarity with it through its theoretical utility elsewhere. (See e.g. Williamson (2000), esp. p.95.) An important assumption I make about the notion is that we *can't* know remote consequences, like Fermat's theorem, *without a proof*, though we can know things that follow by simple deductive steps (e.g. Kelly can know she's in danger). We need to use this notion, 'can know', in our proposed norms rather than just 'know', otherwise by simply refusing to believe (and thus know) something you could get off the hook too easily.

So amended, and otherwise paraphrased just slightly for clarity, the Hacking proposal (or perhaps better the "Hacking-Harman" proposal) for the norms is as follows, written here as if they replace A1 and A3 — A2, Non-negativity, is left alone:

- (H1): if you have considered p , and you can know p , then the rational credence in it is 1.
- (H3): if you have considered p and considered q , and you can know that they are not both true, i.e. $\sim(p \& q)$, then the rational credence to give their disjunction is the sum of your credences in the disjuncts.

(Again, I mean “you considered or ought to have considered” throughout here.)

Hacking just gave those two. I think there are many more such (static) norms, just as basic as the above two. (I call a norm ‘basic’ to mean it is not derivable from other norms.) We can, for example, propose plausible (static) norms that correspond to the other laws of probabilism and which incorporate Hacking’s knowledge constraints:

(H-Entailment): if you have considered p and considered q , and you can know that it’s not the case that $p \& \sim q$, then the rational credence to give q is at least the rational credence you give to p .

(H-Partition): if you have considered a set of propositions, and you can know that exactly one of them is true, then the rational credences to have in those propositions add up to 1.

Below I’m going to argue that there are also basic norms of a different sort (dynamic norms), in the full explanatory theory of rationality, but at this point we can already see that the above (static) norms accommodate the motivation Williams had for his own view. When a paradox leaves you in rational doubt about whether some instance of $P \vee \sim P$ is a logical truth, you will also be in rational doubt about whether $P \vee \sim P$ is a truth at all, and thus you will not be in a position to know it, and so, there is no violation of rationality.

Why, when I wrote out these H norms, didn’t I mention necessity or entailment? Why did I say “you can know that it’s not the case that $p \& \sim q$ ” instead of something closer to the traditional idea, like “you can know q is a consequence of p ”? Similarly, why not talk explicitly of exclusion and partitions? I did, after all, keep those suggestive terms in the names of the two rules ‘H-Entailment’ and ‘H-Partition’. Why drop all talk of logic and modality? (Some might say these clauses I used are a species of “material” modality, but I’m not sure that’s a helpful observation.) The reason I didn’t put the norms in modal terms is that I don’t think the most plausible basic norms here require knowledge of any modality, not logical, metaphysical or epistemic modality. But it is true that *if* you knew any kind of such necessity, that would trigger the above H norms, since knowledge of such necessities would rationally commit you to the various conditions I used in my clauses. For example, if you know that certain propositions are, say, metaphysically or nomologically exclusive, or form a partition of the metaphysical or nomological possibilities, then H-Partition will kick in for you. Even if the modality is merely epistemic—the kind Williams used—the norms kick in. For example, if you can know that $p \& q$ must be false, in the same epistemic sense of ‘must’ that Williams was discussing, this triggers H3. The critical difference with Williams is that you need to *know*

this thing that we use an epistemic modal to express or attribute. (There's a million interesting and difficult philosophical questions about knowledge of epistemic modals, but we don't need to settle them before we endorse the present view. All endorsements are defeasible anyway!)

Hacking's proposed norms are *static* norms. They are (as amended) plausible, but, as I said, I think they give us an incomplete picture.⁴ They're only one part of the theory of rational credence. We need *dynamic* norms, as well as static norms. Let me now clarify this distinction. When I talk about "dynamic" norms, I'm talking about principles that say how a rational person—a person with rational attitudes—may, or must, revise their attitudes by engaging in *reasoning*. And when I talk about reasoning, I'm talking about *basing*; I'm talking about the adoption, or rejection, or re-establishing, of an attitude on the *basis* of other attitudes (or perhaps other mental states, e.g. maybe experiences). Probabilism makes no mention of reasoning or basing. So, we can call probabilism or any such norm a "static" norm. It tells you how your attitudes, at any given time, should have ended up, but doesn't comment on how you reasoned your way there.

The proposal I want to advance here is motivated by this general idea: any true static norms will be entailed and explained by the true dynamic norms. I won't give a full argument for this general idea, but I will rely on it as a reasonable motivation for the proposal I want to develop, which conforms to the general idea. John Broome's book, *Rationality Through Reasoning* (2013), makes a strong case for the general idea. I'll give one motivating illustration.

To see that we need dynamic norms, and to see what they may look like and what their explanatory role is, let's briefly turn away from credence and look at the rationality of full beliefs for a moment.

Full beliefs are also subject to static norms. For example, here is one plausible static norm concerning consequence, though there is admittedly controversy (found in the contemporary debate over so-called "closure" norms):

(F-S-Entailment): if you are considering the propositions p and q , and you can know that q is a consequence of p , then it's irrational to believe p and disbelieve q .

But the theory of rational full belief must also include something that explains how a rational person reasons. There are, in other words, dynamic norms. Again, any specific proposal of a dynamic norm will be controversial (in this case due to the ongoing debate in contemporary epistemology over cases that Crispin Wright has named 'transmission failure'). But just for the sake of giving a concrete illustration, let me give one plausible suggestion for a dynamic norm here (underlining is to help readability):

(F-D-Entailment): if you're considering the propositions p and q , you have and rationally maintain belief p , you can know that q is a consequence of p , and you lack any reasons to disbelieve q , then you're rationally required to infer q (or, if you already believed it, to re-establish q) on the basis of your belief that p .

There are also some related dynamic norms including one requiring you to infer $\sim p$ when you *do* have sufficient reason to believe $\sim q$. I won't write it all out. These dynamic norms for reasoning are what the logical proof rules of Modus Ponens and Tollens mirror, but since we are discussing *reasoning*, not logic, we need these clauses about your lacking defeating reasons and the bases being known or rationally believed. (I don't know how to give principles that fully explain when you have these defeating reasons I'm mentioning. But the phenomenon is intuitive.)

In this example, the static norm F-S-Entailment is entailed and explained by the dynamic norms F-D-Entailment and its "tollens" counterpart.⁵ It's plausible that the situation is analogous for credences. This analogy is my main reason for believing that there are dynamic norms for credences, and that they entail and explain the static norms.

What might the dynamic norms for credence be? Here are some suggestions that I hope will look plausible. Even if they're not exactly correct, I hope they at least illustrate the general project here in a way that supports thinking it's on the right track.

(HD-1): if you are considering p as well as evidence in virtue of which you can know p , then you are rationally required to adopt (or re-establish) a credence of 1 in p on the basis of that evidence.

(HD-3): if you are considering p and considering q , and you can know that $p \& q$ is false, then you're rationally required to, on the basis of that knowledge, revise (or re-establish) your credences in the disjunction and the disjuncts so that your credence in the disjunction is the sum of your credences in the disjuncts.

(HD-Entailment): if you're considering p and considering q , and you can know $\sim(p \& \sim q)$, and your credence in q is lower than your credence in p , then you're rationally required to, on the basis of that knowledge, raise your credence in q or lower your credence in p so that the former is no less than the latter.

(HD-Partition): if you're considering a set of propositions and you can know exactly one of them is true, you're rationally required to, on the basis of that knowledge, revise your

credences in those propositions so that your credences add up to 1.

It would be nice to also have principles that tell us some “only if”, i.e. necessary, conditions on when to adopt certain credences, but these are a bit hard to find. For example, you can have rational credence 1 in p even when you don’t know it. One case is when $\sim p$ is one of infinitely many equally likely outcomes, for example when God is picking a random real number and p is the proposition that they’ll pick 42; you might then have credence 1 that they’ll not pick 42. Another problem: I think I should have credence 1 that I have hands because I’m in a position to know that (and Clarke (2013) and Greco (2015) convinced me credence 1 is not problematic), but then, as Declan Smithies reminded me, I should also recommend credence 1 for the poor handless brain in a vat, even though they’re not in a position to know they have hands. One option is to propose principles concerning not what we’re in a position to know but what we’re in a position to rationally believe, but I’ll not speculate further here about the right necessary conditions on having these rational credences.

It may seem that, by gesturing at so many candidates for dynamic norms on credal reasoning, I’m only making it less convincing that these are *basic* norms (i.e. are among the norms that don’t just derive from other norms) that constitute part of the correct theory, a theory that replaces probabilism with its very short and elegant list of axioms, A1–A3. But that would be a mistake; it would be a mistake to expect the correct theory of rational reasoning to feature such a short and elegant list of basic principles. This mistake naturally arises if you confuse being a good candidate for a mathematical axiom and being a good candidate for a basic epistemic norm. What makes a principle a suitable candidate for being an axiom is something to do with its mathematical interest or practical use in mathematics, and certain such interests or practical uses pressure mathematicians to seek out the shortest and most elegant sets of axioms that generate the desired theory. But the basic norms governing rational reasoning are often inelegant, for example they may include logical or mathematical redundancies. HD-Partition, for example, is (in a sense) logically “redundant” given the other norms, and the others are likewise “redundant” given HD-1 and HD-3, but I take this as no reason to think HD-Partition is any less fundamental, any less likely to feature as a basic norm. (In calling some norms ‘redundant’ here, I don’t mean that they are non-basic, i.e. that you can derive them from the others; you can’t. I mean that, while obeying some but not all of the norms, you could rationally reach all the same conclusions just by “considering” the right propositions and getting yourself in a position to know the right proofs.)

I think HD-Partition is a basic norm because natural examples of rational reasoning with credences seem to be nicely explained by it, with the explanation *not* running through any other norms like HD-1 and HD-3. For

example, suppose the detective is wondering what the chance is that the butler committed the murder and the only other possible suspect is the maid. If the detective knows that just one of them did it, i.e. these are two “exclusive and exhaustive” possibilities, then the detective’s credences in the two should add up to one.

(What if it’s possible that the butler and the maid were in on it together, or possible that the death wasn’t a murder at all? What if you or I know these things are possible? In that case, we’d have to agree, the detective did not *know* that the two possibilities of the butler’s being the murderer and maid’s being the murderer form a partition, that exactly one of these possibilities is true.)

The next topic to discuss would be dynamic norms for conditional credence. Space constraints prevent me from discussing it here, so I’ll refer interested readers either to a slightly longer version of this paper that adds a bit on conditional credence at the end here (posted on my website, www.sinandogramaci.net), or to the related discussion of the topic in Dogramaci (2018).

Acknowledgements

For helpful comments, I’m very grateful to Brian Knab, Karl Schafer, Josh Schechter, Declan Smithies, and Robbie Williams.

Notes

1. Smithies makes the undefended step when he says: “My working assumption is that facts about rationality are explained by facts about reasons or justification. More specifically, I assume that rationality **requires one to believe** a proposition just when and because one has sufficient reason or justification to believe that proposition.” (p.2775, bold added) (On p.2782, Smithies makes clear that the above language concerns propositional justification, “Propositional justification is a matter of having reasons or justification to believe a proposition... .”) Smithies also again explicitly says that an ideally rational agent has a doxastically justified belief in every proposition she has propositional justification to believe on p.2789.
2. After kindly reading a draft, Williams said that he’s wary of claiming his view will solve the *whole* problem of logical omniscience, where the “whole” problem would include both his own motivation about rational low credence in “elementary” truths and our motivation about rational low credence in “remote” truths. (See the next footnote for a bit more on why Williams is wary of claiming his view solves the whole problem.) I should repeat, to be clear, that in his paper Williams didn’t take our problem of rational low credence in “remote” truths as his motivation. I just say it would be natural to think his view promises to solve the whole logical omniscience problem, and Williams also agrees that many people could naturally think or hope so, even though he is wary of this himself.

3. In his helpful comments on a draft, Williams agreed that it's a challenge for him to say what further facts constrain a persons' doxastic possibilities in such a way that her credences may violate doxastic probabilism. But he's optimistic about meeting the challenge, and in his comments he suggested some possible resources for meeting it.

For a first example, Williams suggested that one factor that could help constrain what doxastic possibilities a person *has* are the facts about what doxastic possibilities are *justified* for her by the genuine *reasons* she has. (I understand the suggestion to be reminiscent of one component of David Lewis's views on what determines what attitudes a person has.)

For a second example, Williams suggested that a person may *accept* (in a semi-technical sense) a particular logic, where what logic a person currently accepts is not settled by her current beliefs or credences, and her acceptance of a logic means that her doxastic possibilities contain only the possibilities of that logic. Thus, someone who accepts classical logic won't have any doxastic possibilities in which $P \ \& \ \sim P$ is true, and this may be so even if, at the moment, she is sincerely asserting that $P \ \& \ \sim P$ might be true. Her assertion will then be mistaken, Williams suggested.

As mentioned in an earlier footnote, Williams still doesn't take this to help solve the whole problem of logical omniscience. Someone who "accepts" classical logic and is confident of the Peano Axioms will now face a counter-intuitive requirement to be confident of Fermat's theorem. As Williams fairly describes the predicament: either he can decline to meet the challenge (of constraining a person's doxastic possibilities), and his view will lack any bite in the way that I worry in the main text, or else he can take up the challenge as he wishes to, but that then prevents the resulting view from solving our original problem of logical omniscience for remote truths.

4. A rare discussion of Hacking's view is Elga & Rayo (ms). They raise an objection to Hacking's view, and propose to replace it with a "fragmented" decision theory, which is a theory that attributes different credence functions to agents depending on what choice is before them in a given context. I actually find Elga & Rayo's objection and proposal sensible, but incorporating it into the present paper would make things complicated. I'll trust that others will be able to see how the view I'll propose and the view Elga & Rayo propose can be taken as fragments that should be integrated in the correct overall view.
5. Declan Smithies, in helpful comments, pushed me on why it cannot be that the static norms are fundamental and explain the dynamic norms. On this suggestion, the dynamic norms are true *because* there must be some way of reasoning that has you fulfill the static norms, and the dynamic norms just tell you to reason in some such way. I reject this view because I think the complete specification of all the dynamic norms will require you to reason in fairly specific ways to fulfill the static norms; the mere fact that you must reason in some way to fulfill a static norm doesn't entail what specific ways you must reason. For example, F-S-Entailment could be fulfilled by reasoning in some weird and complex way that doesn't look rational at all but still has you end up obeying that static norm. But you have to obey F-S-Entailment by reasoning in certain ways, like the Modus Ponens/Tollens

way that F-D-Entailment indicates. So, the more specific dynamic norms entail the more general static norms, but not vice versa.

Bibliography

- Bradley, Darren, 2015, *A Critical Introduction to Formal Epistemology*, Bloomsbury.
- Briggs, R. A., 2009, "Distorted Reflection", *The Philosophical Review*, 118, 59–85.
- Broome, John, 2013, *Rationality Through Reasoning*, Wiley-Blackwell.
- Christensen, David, 2004, *Putting Logic in its Place*, Oxford.
- Clarke, Roger, 2013, "Belief Is Credence One (In Context)", *Philosophers' Imprint*, 13, 1–18.
- Cohen, Stewart, forthcoming, "Reasons to Believe and Reasons to Act", *Episteme*.
- Dogramaci, Sinan, 2018, "Rational Credence Through Reasoning", *Philosophers' Imprint*, 18, 1–25.
- Elga, Adam, and Agustín Rayo, ms., "Fragmented Decision Theory", (version of Sept 2, 2016).
- Field, Hartry, 2003, "No Fact of the Matter", *Australasian Journal of Philosophy*, 81, 457–480.
- Friedman, Jane, forthcoming, "Junk Beliefs and Interest Driven Epistemology", *Philosophy and Phenomenological Research*.
- Greco, Daniel, 2015, "How I Learned to Stop Worrying and Love Probability 1", *Philosophical Perspectives*, 29, 179–201.
- Hacking, Ian, 1967, "Slightly More Realistic Personal Probability", *Philosophy of Science*, 34, 311–325.
- Hájek, Alan, 2008, "Arguments for—or against—Probabilism?", *British Journal for the Philosophy of Science*, 59, 793–819.
- Harman, Gilbert, 1986, *Change in View*, MIT.
- Joyce, James, 1998, "A Nonpragmatic Vindication of Probabilism", *Philosophy of Science*, 65, 575–603.
- MacFarlane, John, 2004, "In What Sense (If Any) Is Logic Normative for Thought?", available at www.johnmacfarlane.net.
- Meacham, Christopher, 2014, "Impermissible Bayesianism", *Erkenntnis*, 79, 1185–1217.
- Pettigrew, Richard, 2016, *Accuracy and the Laws of Credence*, Oxford.
- Schoenfield, Miriam, 2012, "Chilling Out on Epistemic Rationality", *Philosophical Studies*, 158, 197–219.
- Smithies, Declan, 2015, "Ideal Rationality and Logical Omniscience", *Synthese*, 192, 2769–2793.
- Steinberger, Florian, forthcoming, "Consequence and Normative Guidance", *Philosophy and Phenomenological Research*.
- Talbott, William, 2008, "Bayesian Epistemology", *The Stanford Encyclopedia of Philosophy*.
- Titelbaum, Michael, ms., *Fundamentals of Bayesian Epistemology*.
- Weisberg, Jonathan, 2011, "Varieties of Bayesianism", *Handbook of the History of Logic*, vol. 10, eds. Gabbay, Hartmann, and Woods, Elsevier.
- Williams, Robert, 2016, "Probability and Nonclassical Logic", *Oxford Handbook of Probability and Philosophy*, eds. Hajek and Hitchcock, Oxford.
- Williams, Robert, forthcoming, "Rational Illogicality", *Australasian Journal of Philosophy*.
- Williamson, Timothy, 2000, *Knowledge and Its Limits*, Oxford.