

Changes in Attitude*

Daniel Drucker

I Introduction

While it's hard to judge proportions here, I'd say most epistemic requirements philosophers propose relate doxastic attitudes of the same given type. Here are two picked somewhat at random (note that I don't necessarily endorse them).

MODUS PONENS REQUIREMENT. Rationality requires of N that, if N believes at t that p , and N believes at t that if p then q , and if N cares at t whether q , then N believes at t that q .¹

DIACHRONIC CONDITIONALIZATION (*narrow scope*). If you have a credence function c at t_0 , and between t_0 and t_1 you learn E and nothing more, then at t_1 you should adopt $c(-|E)$ as your credence function.²

My question is: are all genuine (i.e., correct) epistemic norms like that? Or are there requirements that take us from one type of doxastic attitude to another type? I'll explore the following answer:

CONSERVATION OF ATTITUDES. No one is epistemically required to have a doxastic attitude of type D merely because they have tokens of a different type of doxastic attitude D' .

*Thanks to Sinan Dogramaci, Dmitri Gallow, Miriam Schoenfeld, Roy Sorensen, Rohan Sud, Matthew Vermaire, and an audience at the University of Texas at Austin for very helpful comments.

¹ See Broome (2013, page 157).

² See Briggs and Pettigrew (2020, page 167) for this particular example (they don't themselves quite seem to accept it but do think it's pretty representative).

What this says should be pretty intuitively clear: I can't, e.g., be epistemically required to have a credence of .89 that it'll rain today because I believe the sky is dark. Something else must be involved in the explanation of why I am so required, if I am, for example some other credences I might have. Beyond being interesting in its own right—focusing on an under-appreciated dimension of potential epistemic requirements—it would also have important consequences for some prominent debates, among which whether we ought to have credences or even imprecise credences.

Still, a number of parts of this principle need explaining. Start with 'doxastic'. I'll say that a type of attitude is doxastic if tokens of that type can be evaluated positively or negatively as accurate or inaccurate, respectively. This is pretty similar to the idea from Anscombe (2000) and Searle (1970) that beliefs have "mind-to-world" direction of fit. Beliefs, of course, are like this, as are credences. Desires aren't (probably³), nor is hatred or admiration. Notice caring is non-doxastic, too, which means that MODUS PONENS REQUIREMENT does not conflict with CONSERVATION OF ATTITUDES.

The 'because' is normative-explanatory. It's roughly the same 'because' as what appears in statements like:

- (1) According to subjective consequentialists, if I ought to do something, I ought to do it because (i.e., that reason why I ought to do it is that) it maximizes expected global well-being.

Perhaps this is a kind of grounding,⁴ but perhaps not. Nothing I say will depend on the specific theory of normative explanation we work with.

Most problematic is 'different type'; as is familiar from the generality problem for reliabilism, giving identity criteria can be very hard. Officially I leave 'different type of doxastic attitude' as primitive, since there are relatively clear cases I'll focus on—belief, precise credence, and imprecise credence—and clear non-cases, like belief that p and belief that q , $p \neq q$, and credence x and credence y that p , $x \neq y$. There are also borderline cases (*de dicto* vs *de se* belief? etc.).

Unofficially, here's a heuristic to use that will help decide a lot of cases. I said that a doxastic attitude-type is one where tokens can be evaluated positively or

³ See, e.g., Lewis (1988) for an argument against identifying the desire that p with the belief that p would be good.

⁴ See especially Berker (2018) for discussion.

negatively as accurate or inaccurate. But evaluating the accuracy of a *belief* is different from evaluating the accuracy of a (precise) credence: in the one case, for example, we can give it 0 or 1, depending on whether the belief is false or true, respectively,⁵ while in the second we must have some non-trivial notion of *distance* from the truth.⁶ And for imprecise credences, we need yet another way of measuring accuracy, some kind of generalization of the way we measure precise ones.⁷ But the way in which we assess the belief that *p* and the belief that *q* for accuracy won't fundamentally differ: we look at whether they're true or false, and score accordingly.⁸ This is just heuristic, though, because someone might insist that the accuracy of the belief that *p* and of the belief that *q* *must* be measured differently because for one you need to check the truth of *p* and for the other you need to check the truth of *q*. Such nitpicking can, I think, be blocked with a careful-enough description of *methods* of measuring doxastic accuracy, but that would take me too far afield. That said, the heuristic is still useful, I think: at the typical level of generality epistemologists operate at, the methods of measuring belief accuracy obviously will apply to both the belief that *p* and that *q* but not credence .4 that *p* (e.g.), or the imprecise credence from .3 to .5 that *q*. So epistemologists, I think, are operating with the same intuitive sense of difference of type that I am.

In this paper, I will explore arguments for and against CONSERVATION OF ATTITUDES with a specifically epistemic flavor. While an interesting thesis in and of itself, whether it's right or not also has implications for other questions philosophers have tried to answer: for example, ought we to have imprecise credences? And indeed, looking at some of the arguments they give *for* (or, even, against) the “yes, we ought to” answer will help us to see some possible routes for *denying* CONSERVATION OF ATTITUDES. But the arguments for the principle may also help see the flaws arguments that we ought to have (imprecise) credences, if

⁵ Or we may weight being wrong as being worse than being right is good. See Easwaran (2016) for good discussion.

⁶ For good, now classic discussions of scoring rules, see, e.g., Joyce (2009) and Pettigrew (2016).

⁷ See, e.g., Seidenfeld et al. (2012), Konek (forthcoming), and Schoenfeld (2017).

⁸ There's some complexity here. Perhaps accuracy measures shouldn't be “extensional”, i.e., distance from the truth for certain propositions, say about certain subject-matters, should perhaps be measured differently. But then this will still be a truth about how *all* beliefs ought to be measured—given they're about *X*, etc.—which couldn't apply to precise credences.

there are any such flaws, or at least places in which the standard arguments require supplementation. For if CONSERVATION OF ATTITUDES is true, arguments that we ought to have imprecise credences, e.g., of a certain kind need to add the descriptive premise that we in fact do already have some (non-trivial?) imprecise credences (or other kinds).

Beyond that, whether the principle is true seems most directly connected to competing pictures of epistemic normativity: specifically, ought we to *improve ourselves* in fundamental ways, or ought we to make the best of how we already are? There are different ways to make these competing pictures precise. One way of getting at it has been to ask whether epistemic norms might require us not just to do our best with the evidence we have, but to get *new evidence*.⁹ I take the question of whether CONSERVATION OF ATTITUDES is true to be a question like that.

Finally, note that CONSERVATION OF ATTITUDES is weaker than it might have been. Essentially it rules out *narrow-scope* requirements *to* have a given type of doxastic attitude *because* one has another one. It does not, without extra premises anyway, rule out wide-scope coherence constraints to, e.g., not both have some token of *D* and not have some token of *D'*. You may think allowing such requirements would violate the spirit of conservatism. Perhaps that's right, but I'm not sure. I leave discussion of such requirements in light of conservatism for another occasion.

This paper is exploratory: my aim will be to consider the best cases I can for conservatism, and the best cases I can against it. I will focus on a specific argument for it that appeals to reasoning, though in the conclusion I will also much more briefly sketch one appealing to a kind of instrumentalism about epistemology. But though this paper is exploratory, since my hope for this paper is that people take conservatism seriously, much of what I will say will be biased in its defense. Be assured I have no idea whether conservatism is right.

2 Reasoning to (Totally) New Attitudes

One natural thought, hard to make precise and persuasive, is somehow that we *cannot* acquire new attitudes, in some sense relevant to 'ought'-implies-'can' type

⁹ The classic discussion is Kornblith (1983). Kelly (2003, page 627) also discusses this issue.

principles; and so if we can't in that sense, then there are no genuine and correct epistemic requirements to have them. In this section, I'll formulate this argument in more detail, arguing that the best version of it appeals to the thought that we can't acquire new types of attitudes *by reasoning*. If the only genuine and correct epistemic requirements are ones we can satisfy by reasoning, then there will be no genuine and correct epistemic requirements that we ought to have new types of attitudes.

We can all recognize at least paradigm instances of reasoning, and though I will offer a characterization of it, I will to a greater extent rely on the ability to know and understand clear instances when we see them (or don't). Thus, when someone thinks to themselves:

- (2) A president of the United States can only have two, four-year terms in the United States. So, if Biden wins his second term in this next election, he will no longer be president by 2030.

that person may have acquired the belief that Biden will not be President in 2030.

What reasoning's nature is is controversial. Here's a minimal characterization I find useful:¹⁰

1. *Basing*. Reasoning has a basing component, whereby the attitude change that results from reasoning is based on the premises (and perhaps also a rule of inference, etc.).
2. *Evaluativity*. In virtue of that feature, the resulting attitude changes can be evaluated as well-grounded, justified, etc., or not.
3. *Responsibility*. In virtue of the first two features, the reasoning person can be regarded as having reasoned well or badly.

Thus the person who engages in the reasoning they might express with (2) bases their new belief about Trump on their old beliefs about presidential term limits (and perhaps other suppressed premises, along with a rule of inference or two). This is a causal process. The resulting belief can be evaluated as well-grounded,

¹⁰ See [Boghossian \(2018\)](#) for this characterization.

justified, etc. on the basis of those premises and inferences rules, and thereby the person themselves held responsible for reasoning well or badly. This characterization is compatible with a wide variety of alternative conceptions of what reasoning really is. Does it involve rule-following, for example?¹¹ Maybe yes, maybe no—certainly I don't want to commit either way. Does it involve taking the conclusion to be supported by the premises, or something like that? Again, I won't commit, and nothing I argue will depend on how someone comes down on these questions.

(2), as I understand it, begins with some antecedent beliefs and takes the person to a new or newly-based belief, one expressed by the conclusion of the reasoning. The typical examples of reasoning that people give are like this. Thus:

(3) If it's raining, the streets are wet. It's raining. So, the streets are wet.

But must all reasoning work this way?

I agree with philosophers who argue that we can reason to degrees of belief.¹² I think examples like the following show that we can do this:

(4) It's fifty-fifty whether she'll get a corgi or a cat. So it's fifty-fifty whether she'll get a dog.

I think someone who reasons in this way reasons to the credence of .5 that the relevant person will get a dog. We'd expect someone who reasons this way to have such a credence by the end of it, for the credence to be based on the other beliefs and credence she expresses, and for her belief and she herself to be evaluable in the ways the minimal characterization requires. But notice that the speaker in (4) better already have a .5 credence, namely that she'll get a corgi. So she'll start with credences anyway and they'll be part of why she has the credence she has at the end.

In my discussion of (4), I assumed that expressions like 'it's fifty-fifty that ϕ ', when uttered sincerely, are used to express speakers' credences. Other examples might be 'possibly ϕ ', 'it's 75% likely that ϕ ', etc.¹³ In assuming that I'm

¹¹ For yes answers, see, e.g., [Boghossian \(2014\)](#); and for no answers, see, e.g., [Dancy \(2018\)](#) and [Valaris \(2017\)](#). (There are many other people who've written on this issue, of course.)

¹² See, e.g., [Staffel \(2013\)](#) and [Dogramaci \(2018\)](#).

¹³ See, e.g., [Yalcin \(2007, 2011\)](#).

not assuming that this is *all* they do, or that they don't have compositionally-determined semantic values that somehow suit them for this work.

Can a person reason from non-degred belief to degred belief? Some philosophers seem to think “yes”. Dogramaci (2018, page 14) has the following “rule of reasoning”, e.g.:¹⁴

(R1) If you know [or, if you grasp a proof] that p , then you are defeasibly permitted to adopt, on that basis, maximal credence in p .

Such a rule, could we follow it, would take us from *knowledge* (or rather knowledge-constituting belief) of a proposition to a particular (maximal) credence in that proposition. Can we follow it in reasoning, even without other credences? More importantly, can we follow it in *good* reasoning without other credences? (To anticipate: it wouldn't help for the existence of epistemic requirements to have new types of attitude if the only way we could reason to them was badly.)

One way to convince ourselves that we can is by producing a compelling instance of good reasoning where the rule seems to be followed. Such reasoning, at least when displayable in writing, typically ends with a sentence that would be used by the person going through the reasoning to express the new state arrived at, which is what we saw with (3) and (4). With that in mind, consider:

(5) ???She'll get a corgi. So, it's absolutely certain that she'll get a corgi.

The first sentence seems to express the knowledge—or belief the speaker ordinarily *takes* to be knowledge—and the second sentence their maximal degree of belief. But this is intuitively bad reasoning, at least insofar as I can understand it at all. Intuitively, having the mental state expressed by the premise just doesn't support having the attitude expressed by the conclusion. Belief, or even knowledge, does not support absolute certainty.

Let's go through this a bit more carefully, though, since there's much more to say. Now, it might be that belief *just is* credence above a certain possibly vague, possibly context-sensitive threshold. There are reasons for thinking any threshold below 1 itself will not work, chief among them (I think) the lottery paradox: for arbitrarily large n , it seems possible for someone have credence $1/n$ that their

¹⁴ He bases this rule on an analogue from Hacking (1967).

ticket will win, while still failing to believe that it won't.¹⁵ And anyway, suppose our threshold *could* be less than 1. In that case, (5) would be very obviously bad reasoning: it would take us from high confidence to maximal confidence! So it seems that our only option here is to say that belief is credence 1, i.e., maximally high credence.¹⁶ But on *this* line, (5) sounds bad for a clear reason: one cannot *reason from* believing that p to being certain that p , since one would already be certain that p just in believing that p .

So the option left for the anti-conservationist is one on which belief *requires* (in some way) maximally high credence without analytically or metaphysically entailing it. Perhaps this is possible. But there are a couple reasons to doubt that this option will really work. First, it would be unclear why (5) sounds bad. Shouldn't that just be normal, good reasoning like (3) is? Second, there is some reason to think belief is "weak", i.e., compatible with pretty low credence. For example, if three horses A, B, and C are in the race and you ask me which one I think will win, I might say B even though I'm only .45 confident it'll be her (because I'm, say, only .3 and .25 confident of A and C).¹⁷ If such philosophers are right, either (6) won't be good reasoning, or a person can coherently accept that probably, p even while having less than .5 credence that p . Finally, it's also just hard to understand *why* belief would require maximal credence without belief's just being maximal credence.

So I doubt (5) can be what we're looking for.¹⁸ But we only need to produce one instance, so perhaps we should simply be less ambitious than we were in trying to go from knowledge to absolute certainty. The trouble is that it's tricky to find examples where a particular credence is picked out. After all, we could try examples like:

(6) She'll get a corgi. So, it's at least likely that she'll get a corgi.

There's reason to doubt that even *this* would be good reasoning of the kind we're after. First, remember, some philosophers think belief is weak, in that it needn't

¹⁵ For this argument, see, e.g., Kaplan (1981) and Clarke (2013).

¹⁶ For a defense of this view, and a context dynamics to make it more plausible than many have taken it to be, see Clarke (2013).

¹⁷ See, among others, Hawthorne et al. (2016), and Moss (2019) for an important recent response to such cases.

¹⁸ Thanks to Sinan Dogramaci for discussion of (5) and how the anti-conservationist might respond.

go with high or even not-low credence. If such philosophers are right, either (6) won't be good reasoning, or a person can coherently accept that probably, p even while having less than .5 credence that p . But even if they're not right so that maybe (6) could express good reasoning, it wouldn't pick out any particular credence: there's no credence that is *simply* greater than .5 without anything more specific underlying it. So in order for the speaker to express a credence at all with the conclusion of (6), that more specific credence has to come from somewhere other than the reasoning expressed by (6).¹⁹

Here's one more try.²⁰ Lewis (1980)'s PRINCIPAL PRINCIPLE, as typically formulated, relates credences to credences. In its usual guise it is a coherence constraint, something like (where ' Ch ' is the objective chance function, suppressing reference to time):

PRINCIPAL PRINCIPLE. Everyone is epistemically required to be such that their credence function Cr satisfies the following constraint: $Cr(p \mid Ch(p) = x) = x$, unless the person has "inadmissible" information.²¹

But we can imagine it as somehow underlying good reasoning. After all, doesn't the following sound pretty good, even reasonable?

- (7) The objective chance of the atom decaying in the next five minutes is .2. So, it's .2 likely that it'll decay in the next five minutes.

In my opinion, this is the anti-conservationist's best shot, along with other sorts of reasoning involving deference, e.g. with testimony, REFLECTION,²² or RATIONAL REFLECTION.²³

But does (7) really express good reasoning? I think likely it doesn't. The underlying rule one might be following, I take it, is something like: from the belief

¹⁹ At this point, you may wonder what, then, \lceil it's likely that ϕ \rceil *does* express. In my own view, it expresses a commitment, namely the commitment roughly to maintaining a credence of greater than .5 that ϕ , no matter how one fills out the rest of one's credences (without getting any new information, however that's to be understood). This is, perhaps, like a constraint; see, e.g., Swanson (2016) for an approach that uses constraints to model the semantics of epistemic modals.

²⁰ Thanks to Dmitri Gallow for suggesting this possibility.

²¹ There are some variations on this principle aimed at solving problems with this formulation. See, e.g., Hall (1994) and Ismael (2008). The differences in formulation won't matter.

²² See van Fraassen (1984) and Briggs (2009).

²³ See, e.g., Christensen (2010).

that the objective chance that p is x , (one may) infer that p is x likely. (Remember, I don't insist that one be following rules when one reasons. So maybe this is just the underlying *pattern*.) Notice that such reasoning is only licensed by the PRINCIPAL PRINCIPLE itself when the agent is maximally confident that the objective chance is x . In almost every other situation, the PRINCIPAL PRINCIPLE does not just license but *forbids* forming that credence. That is, assuming belief that p doesn't entail credence 1 that p , as I think we should think it doesn't, by far the more normal situation is for credence x that p to be *impermissible*. So, either we must also assume the relevant person already has credence 1 that the objective chance of the atom decaying in the next 5 minutes is .2, or we have to think the reasoning is simply bad.

It turns out that it's surprisingly hard to display examples of reasoning from binary on-off belief to credences. That gives some "inductive" reason to believe that it doesn't happen. The basic difficulty so far, it seems to me, is that good reasoning with one type of attitude won't have the right guarantee of corresponding to good reasoning with another type of attitude. Intuitively, belief is very coarse, and so there are circumstances in which it is perfectly permissible to have a given belief in which it is thereby permissible to have a given specific credence and circumstances in which it is not; the very coarseness of the attitude means that no specific credence will be permissible in the way that reasoning must make it if it is good. There are too many circumstances in which a given belief is permissible, and they won't always line up properly with when a specific credence is permissible. This explanation, of course, depends on the different types of attitudes being comparable in terms of their intuitive coarseness, and so won't generalize very readily to differences between types of attitudes other than differences in coarseness.

That explanation does suggest one place to look for an example of good reasoning like this, namely in going from the specific to the coarse. Now, granted, this isn't completely trivial to do. After all, for most n between 0 and 100, the following sounds bad:

(8) It's $n\%$ likely that she'll get a corgi. So, she'll get a corgi.

(9) It's 95% likely that she'll get a corgi. So she'll get a corgi.

Maybe if n gets high enough, the reasoning isn't obviously bad. I suspect that'll be pretty controversial, turning as it does on very difficult issues made vivid by, e.g., the lottery paradox. But let's skip straight ahead to the extreme case:

(10) It's absolutely certain that she'll get a corgi. So, she'll get a corgi.

This doesn't, at least, sound like *bad* reasoning, not in the way (9) does. Someone would not exhibit a lack of caution if, somehow, they reasoned in the way they would seem to there.

Yet something does sound off about it. The reason is that it is hard to imagine someone reasoning as in (10) so that they *end up* with the belief that the relevant person will get a corgi. It is not at all clear to me that someone can have a maximally high credence that p even though they don't believe that p .²⁴ The challenge would be to find a functional role for belief that p to have that being maximally confident that p won't itself play.²⁵ That is, if for the premise in (10) to really express the person's attitude they *already* have to have belief that there's a corgi, there's no reasoning *from* credences to beliefs: the episode of reasoning won't give rise to the belief. This is a different problem, notice, from before. Rather than the reasoning being bad, it is really not reasoning at all, at least if reasoning must produce the attitudes expressed by the conclusion in at least some cases. And moreover, notice the more extremal our n gets in (8), the more likely this is true of that kind of reasoning, too; whereas the less extremal it gets, the more the reasoning looks straightforwardly bad.

This response to (10) does not require that every belief reduces to credence above a certain threshold, even given a context-sensitive and/or vague threshold. For everything I've said it might be that, in general, belief and credence can come apart. Rather, all I assume is that maximal confidence entail belief. I think any reasonable functional specification of certainty and belief will support that claim, but I won't go too into that here, since the functional roles of each are controversial.

²⁴ Perhaps an infinite dartboard case is possible, one in which a person has credence 1 that an infinitely thin dart did not hit the exact bullseye without believing that it didn't. But then there's good reason for thinking "credence 1" wouldn't really be the attitude expressed by "it's absolutely certain that ϕ ". See Williamson (2007) for important discussion of dartboard-like cases.

²⁵ Perhaps there is more to being a belief than functional role, e.g., some characteristic phenomenology. See, e.g., Schwitzgebel (2002). But I very much doubt that phenomenology can fail to be present despite being maximally confident.

Given that (10) won't do what we're hunting for, it looks like nothing will: that really seems our best shot, at least so far. Still, you may have noticed that I haven't really discussed *imprecise* credences and how we may reason to or from them. For my own part, that's in part because I'm not at all confident I can produce *any* reasoning the conclusion of which clearly expresses an imprecise credence. The other reason I haven't focused on them is because I expect all the same concerns I've voiced about transitions between beliefs and precise credences (and *vice versa*) will apply to imprecise credences, too, even if we clearly can reason to them in some cases.

That said, and to illustrate those concerns, the following seems worth checking on:

(11) It is absolutely certain that the objective chance that this coin will land heads when flipped is exactly between .4 and .6. So, it's between .4 and .6 likely the coin will land heads when flipped.

The problem with (11), or anyway what makes me uncertain about its expressive properties, is that the following *doesn't* express any nontrivial imprecise credence:

(12) If something is .5 likely, it's between .4 and .6 likely. So, if it is absolutely certain that the objective chance the coin will land heads when flipped is .5, then it is .5 likely to land heads when flipped. It is absolutely certain that the objective chance that this coin will land heads when flipped is exactly between .4 and .6. So, it's between .4 and .6 likely the coin will land heads when flipped.

If this reasoning is good (and I'm not positive it is, but it seems to be), then the conclusion of (11) needn't express an imprecise credence; a person may reason as in (12) without having anything like an imprecise credence.

So even in the case of (11), we haven't found a case of reasoning that clearly concludes in the expression of an imprecise credence. With both belief and precise credence, we have conventional ways of expressing the credences we're interested in. With imprecise credence, though, we seem not to. That doesn't mean there aren't any such cases of reasoning, but minimally it does make them very hard to produce. Perhaps if we did have such expressions, (11) could somehow be developed into a concrete and plausible case of reasoning to an imprecise credence.

Maybe this is just a question of creating the right expressions with the right conventional expressive powers. So let's invent a predicate, $\lceil \text{likely}^*(\phi, x, y) \rceil$, for conventionally expressing imprecise credences that ϕ that form an interval $[x, y]$. (I'll ignore the more general case where the imprecise credences don't form intervals.) Then our analogue of (11) is:

- (13) It is absolutely certain that the objective chance that this coin will land heads when flipped is exactly between .4 and .6. So, it's likely*(the coin will land heads when flipped, .4, .6).

The question now arises whether (13) would express good reasoning. It seems to me likely it does not. Notice, again, that one can accept the premise of (13) while rejecting the imprecise credence, because one has more information about the objective chance than one says. What one needs, for (13) to be good reasoning, is a justification for treating every value in the interval on a par, as imprecise credences do.²⁶ One way this could happen is if one has imprecise credences represented by a set of (precise) credence functions each of which is maximally certain that the objective chance is x , for some value $x \in [.4, .6]$. Then upon conforming to the PRINCIPAL PRINCIPLE, the person could update their probabilities function-by-function and end up with the relevant imprecise credence functions. But *that* kind of reasoning is reasoning from imprecise credences to imprecise credences, and so not what we're after. And notice that if some of those credence functions are *less* than maximally certain that the objective chance of the coin's coming up heads is some particular $x \in [.4, .6]$, the reasoning in (13) would *conflict* with the PRINCIPAL PRINCIPLE, for the reasons we saw in the discussion of (7). So, unless one does start out that way—with imprecise credences—the person seems to be going beyond what is licensed by their other attitudes. The problem is pretty similar to the problem with using something like the PRINCIPAL PRINCIPLE to go from belief to precise credence: the reasoning, to be good, assumes the existence of tokens of the new doxastic type, and thus cannot be counterexamples to CONSERVATION OF ATTITUDES.²⁷

²⁶ For an explicit model of imprecise credences that allows for a “degree of membership”, see, e.g., Gärdenfors and Sahlin (1982), Lyon (2017, section 4), and Carr (2020, section 2.2).

²⁷ Perhaps one could use “principle of indifference”-type reasoning here; but then it would seem arbitrary to use that reasoning at the point of deciding on imprecise credences and not at the point of determining a precise credence, rather than reasonable. Anyhow, I won't explore this option further.

There's a lot more to say about many of these examples, and I'm sure many other examples it would eventually be important to consider.²⁸ Despite that, I'll move on; the conclusion that we cannot reason to different types of attitude at least looks more plausible now. But that is only the first step in the argument I'm developing in this section. The next step has to do with the connection between reasoning and requirements.

The first part of this connection, one I won't defend in detail, is that for someone to be epistemically required to φ , they have to be able to φ rationally, since we generally have to be able to satisfy our obligations in for them to be real obligations. That is, whether they φ must be under their rational control.²⁹ Of course, 'ought'-implies-'can' principles are in general not uncontroversial, but I will simply assume they are ultimately correct, and that this one in particular is correct.³⁰ The second part is that, when it comes to the formation of new attitudes, it's at least *prima facie* plausible that the way to do so would be *via* reasoning: control involves *basing* on relevant inputs, and in light of that basing, the new states and agent themselves ought to be evaluable as well-founded (etc.) and responsible, respectively. Perhaps there are other sorts of intelligible rational control than this, but since I cannot myself think of them, I will—perhaps more tentatively—assume that this part is right, too.

That completes the argument. To state it explicitly:

PI. If there's an epistemic requirement to have a *D*-type attitude *because* one has some *D'*-type attitudes, then one can conform to the requirement by

²⁸ Here are some other possibilities, very briefly. First, it seems I can reason from the supposition that p and the derivability of q therefrom to that if p , q . But supposition isn't by my reckoning a kind of doxastic attitude, since suppositions aren't evaluable positively or negatively as accurate. Second, perhaps I can reason to suspension of judgment, maybe with something like: 'there's strong evidence that p , strong evidence against p , so who knows what's right?'. (It'd take work to come up with a compelling concrete example.) Once again, though, I doubt suspension of judgment is doxastic in my intended sense: it is not, it seems to me, accurate *or* inaccurate, which is different from, e.g., middling credences. (See Friedman (2013) for an argument that suspension is different from middling credence.) Thanks to Sinan Dogramaci for suggesting I look at these cases. And finally, you might wonder about what we might call *thoroughly* vague credences, perhaps another way of understanding someone's view that it's likely that their friend will get a corgi, while having no more specific attitude. Ultimately I wonder if this will collapse into some other variety of credence, and I also need to see a better-developed model of how to understand such attitudes. But I want to mark it here as at least a further possibility.

²⁹ For this idea and a good defense, see Wedgwood (2013).

³⁰ For some objections to the idea that 'ought' implies 'can', see, e.g., Sinnott-Armstrong (1984) and Graham (2011).

reasoning well from the D' -type attitudes to the D -type attitude.

P2. When $D \neq D'$, one cannot reason well from D' -type attitudes to D -type attitudes.

P3. So, there are no epistemic requirement requirements to have a D -type attitude *because* one has some D' -type attitudes, where $D \neq D'$.

I've given at least initial defenses of P1 and P2. As a result of this, I think this argument is worth taking very seriously. And its conclusion C1 is CONSERVATION OF ATTITUDES. In other words, we now have pretty good reason to think that thesis might well be right. As I said, of course, much of what I've said is apt to be quite contentious, and there are likely many examples yet to consider. My aim is just to present a plausible enough argument with an important and illuminating structure.

And yet perhaps it would be somewhat surprising were CONSERVATION OF ATTITUDES right, since epistemologists have indeed *argued* that we ought to have new types of doxastic attitudes than, perhaps, ones we start with. In light of that fact, in the next section I'll look a bit at these arguments in order to see whether they point to a weak spot in the argument I developed in this section, or rather whether it's the other way around and the argument in this section can point to problems with these other arguments.

3 *Arguments for New Doxastic Attitude-Types*

As I said, in this section I'll look somewhat briefly at broadly epistemological arguments that we ought to have new types of doxastic attitudes. I'll start by considering an argument-type that we ought to have credences rather than just all-out beliefs, and then I'll consider a prominent argument that we ought to have imprecise credences rather than (just?) precise ones.

The preface case is by now well-known: a historian writes a book, coming to believe her main claims, C_1, \dots, C_{100} ; but on the basis of past experience of her own and others' work, she also believes at least one of her main claims is wrong, i.e., $\neg(C_1 \wedge \dots \wedge C_{100})$. But it's easy to derive a contradiction from all these beliefs. But she doesn't seem irrational. So, it's either not irrational to have beliefs that entail a contradiction—which is unpalatable—or there is something

wrong in her believing as she does. I want to focus on the reaction that says that an agent is, in such cases, epistemically required to have not binary beliefs but *credences*.³¹ For it is perfectly coherent to have high credences in all of C_1, \dots, C_{100} with low credence in their conjunction. On this reading of the argument from the case, we are recommending to the historian that she switch from belief to credences (in this sort of case, at least). That is:

- (★). In the preface case, the historian is epistemically required to switch from belief in the various relevant propositions to given high credences.

Next, let's turn to arguments for imprecise credences. One of the most standard arguments calls our attention to cases in which our evidence is radically impoverished with respect to a certain proposition. Elga (2010, page 1), e.g., has the following case: "A stranger approaches you on the street and starts pulling out objects from a bag. The first three objects he pulls out are a regular-sized tube of toothpaste, a live jellyfish, and a travel-sized tube of toothpaste. To what degree should you believe that the next object he pulls out will be another tube of toothpaste?". Any specific numerical degree of confidence would seem in such a case to go beyond your evidential situation. The idea is, roughly, that unspecific evidence calls for unspecific doxastic attitudes. *Which* unspecific non-doxastic attitudes unspecific evidence calls for is perhaps controversial, but one leading contender is of course imprecise credences, attitudes representable as *sets* of precise credence functions (for particular propositional arguments).³² Our attitudes would respect our radical evidential impoverishment in virtue of how large a range of values elements of that set take for the given proposition. Of course, for this argument to have a chance of succeeding, we need some argument that this is the *best* way to handle evidential unspecificity, but all I want to do right now with this argument is show the general form the overall argument tends to take.³³

³¹ This is how I read Christensen (2004), or at least how I am inclined to read him.

³² For these kinds of arguments, see, e.g., Levi (1985) Joyce (2005), and Sturgeon (2008), among others. There are, it bears emphasizing, other epistemic arguments for the requirement to have imprecise credences; see, e.g., Schoenfield (2012).

³³ For his part, Sturgeon argues that using imprecise credences (so conceived) is *not* the best way to handle unspecific evidence, but rather even fuzzier variants. See also Carr (2020) for an alternative proposal to handle cases like these, using precise credences about what the evidence supports or what is rational.

Such arguments, let's suppose, aim to motivate the following:

- (†). In cases of evidential impoverishment, we are epistemically required to have (non-trivially) imprecise credences.

So we have two claims, (★) and (†), of similar forms. Are they right? Well, suppose they *are* right. Then by P1, we should be able to reason to bring ourselves in compliance with them. Can we do that?

Note that even if we *could* reason to that new belief, to the belief that we are epistemically required to have imprecise credences, that doesn't mean that we can reason *to* the new imprecise credences themselves. Reasoning to the conclusion that one ought to do something, or is required to be some way, is very different from reasoning to doing that thing or being that way. Thus:

- (14) I ought to believe anything that follows from beliefs I hold strongly *via* an application of *modus ponens*. That *q* follows from my strongly held beliefs that *p* and that if *p*, then *q*. So I ought to believe that *q*.

To actually *reason* to *q*, the person might actually use *modus ponens* in reasoning; (14) doesn't do the trick.

This is an important point. Arguments to the effect that we ought to have a given type of doxastic attitude—reasonably enough, of course—tend to conclude with the claim that, well, we *ought* to have them, that we're epistemically required to have them, etc. Someone who is convinced by these arguments and goes through the reasoning itself, though, would not *ipso facto* form the attitudes they think they ought to have.

But *couldn't* they? You might think we should be able to, since anti-akratic principles (roughly, principles forbidding believing that one ought to be some way without being that way) are really plausible.³⁴ But consider how one might reason:

- (15) I ought to believe that she'll get a corgi. So, she'll get a corgi.

³⁴ Thanks to Miriam Schoenfeld for pushing me to say more on this point. It's worth noting that not *everyone* finds them plausible. For some doubts, see, e.g., Kolodny (2005), Lasonen-Aarnio (2020), and Weatherston (2019, chapter 10).

This, it seems to me, is bad reasoning. p in no way follows from that one ought to believe that p , either deductively or (it seems) inductively, either. It is not like using *modus ponens* or any other familiar rule of reasoning. One may add these in as premises, of course. That is, the following might be good reasoning:

(16) I ought to believe that she'll get a corgi. If I ought to believe that she'll get a corgi, then she will get a corgi. So, she'll get a corgi.

I doubt the conditional is very easy to motivate in cases like these, but that's not important. Imagine how the analogous reasoning would have to go between beliefs and credences:

(17) I ought to assign .98 credence to C_{38} . If I ought to assign .98 credence to C_{38} , it's .98 likely that C_{38} .

(17) will not help, though, because most likely the middle premise expresses a conditional credence of .98 in C_{38} given that the person ought to assign .98 credence to C_{38} . In other words, we have not gone only from beliefs to credence, but beliefs and (conditional) credences to credences.

So I'm not optimistic this response will work. So someone cannot, it seems to me, use (\star) or (\dagger) to comply with (\star) and (\dagger) . There's a gap between the conclusion that one *ought* to be a certain way, or that one is epistemically required to be that way, and the conclusion the accepting of which would amount to making oneself that way by reasoning. Perhaps there is some way to do it, but it seems to me that the kinds of arguments I presented in order to motivate (\star) and (\dagger) don't themselves allow us to reason in ways that allow us to comply with them.

Of course, all of this should be worrying to the people who have put these arguments forward only to the extent that P1 from the previous section actually is right. But assuming for the moment that it *is* right, we have arrived at an interesting distinction between two kinds of arguments that attempt to motivate rational requirements for our doxastic attitudes. The first kind of argument is one that a person themselves can *use* in reasoning to conform to the putative rational requirement. Call this a *facilitating* argument. The second kind is one that a person *cannot* use in reasoning to conform to the putative rational requirement. Call that a *non-facilitating* argument. Other examples of non-facilitating arguments include:

- Dutch book arguments for synchronic and diachronic coherence.
- Accuracy-dominance and expected inaccuracy arguments for various epistemic norms.
- Money pump arguments for (e.g.) transitivity of preferences.

It's important to see, though, that even if P_I is correct, non-facilitating arguments needn't be failures: we might have *other* ways of conforming to the requirements such arguments purport to establish. To run an argument like the one from the previous section, we need reason to think in these cases we won't be able to reason to comply with those requirements. And often, we don't have any such reason. Still, when we *do* have a non-facilitating argument, we should wonder about the possibility of a P_I - C_I -type argument, and we should always try to assure ourselves that the requirements can be complied with by reasoning with *other* arguments. Take, for example, MODUS PONENS REQUIREMENT, the requirement with which I began. Such a requirement *can* be satisfied because individuals can use *modus ponens* directly in good reasoning. So an argument that we ought to conform to that requirement can work, because of the existence of good reasoning by *modus ponens*.

Because of the ambition of these conclusions, though, it's worth looking further into what advocates of these arguments (that we ought to have precise or imprecise credences) might say in response to the threat of the P_I - C_I -type argument. In the next section, I'll look at what is to my mind, the most promising response.

4 Picking New Attitudes

A brief summary of the dialectic so far should be helpful. In section 2, I presented an argument, P_I to C_I , that if successful showed that there is no reasoning from one doxastic attitude-type to another, e.g., from beliefs to precise credences, precise to imprecise credences, etc. But then in the previous section, I presented some *prima facie* compelling arguments for (\star) and (\dagger) , that is, arguments that purport to show that are epistemically required to form some of these new attitude-types, given that we have the old attitude-types. Yet those arguments don't have an obvious response to the P_I - C_I -type argument, and so

(\star) and (\dagger) look suspect. In this section, I want to explore a kind of response to the challenge the P_I–C_I-type argument presents.

According to this response, I have been too limited in considering the kinds of reasoning that might do the work of bridging us from one type of doxastic attitude to another. Or perhaps a different way to think about it is: we can actively conform to rational requirements in ways other than reasoning directly to states that comply with those requirements. And we can *specifically* do that when the rational requirements are existential in nature, as (\star) and (\dagger) are. To see this, consider a standard “Buridan’s ass”-type case. Remember that in this case a person sees two goods A and B that are indiscernibly as good as one another, that is, good in every respect the agent cares about, and to the same degree. Further, the agent knows that they can only take at most one of A and B. (In the standard example, we have an ass facing a choice between two relevantly qualitatively identical bales of hay placed equidistantly from the ass,³⁵ but we can imagine picking between two identical boxes of cereal at a supermarket.) Manifestly the agent *would* take one of them. The philosophical challenge is to make intelligible how this might happen.

Perhaps we use some device, including the vagaries of our attention, to pick roughly randomly. One of the options “strikes us” the right way, and that’s the one we choose. But this leads to regress: how do we choose *striking us* as the criterion, as opposed to some other criterion, when all the criteria are indiscriminably good ones? And how do we choose when something *has* struck us?³⁶ A different possibility is that we *simply* pick. (‘Plump for’ seems to me to come to roughly (?) the same thing.³⁷) Of course, that’s just a name for a solution without further analysis.³⁸ The idea, then, is that in Buridan’s ass-type cases, we can pick in that way, even if we don’t have reason to prefer one to the other. To the extent that we are convinced that they work *somehow* we can offer them to

³⁵ Apparently Al-Ghazali’s original example involved a man choosing between two qualitatively identical dates, in his *Tabafut Al-Falasifah*.

³⁶ For the idea of randomization, see, e.g., Rescher (1959), and for regress problems for randomization, see, e.g., Ullmann-Margalit and Morgenbesser (1977) and Chislenko (2016), among others; apparently the idea of a regress here comes from Leibniz.

³⁷ See, e.g., Blackburn (2010).

³⁸ One possibility, not the only one, is that when we pick an option, we decide or intend to nonintentionally perform one of the relevant actions. For this analysis, see Chislenko (2016). This sort of decision is familiar from when we decide to doodle or pace “randomly” while talking on the phone.

the anti-conservationist as a way of responding to the P_I–C_I-type argument.

So perhaps we can imagine the following response on behalf of the defenders of (★) and (†). We can, using their arguments, conclude with (★) and (†). But we have seen that we can't get ourselves to a point where we can't reason to any *specific* credence and no *specific* imprecise credence. Still, we can at that point *pick* a credence, precise or imprecise as the case may be, that is, decide to nonintentionally arrive at a relevant credence. Thus, return to (6), e.g.:

(6) She'll get a corgi. So, it's at least likely that she'll get a corgi.

The main problem with (6), remember, is that there's no credence that is simply greater than .5 without anything more specific underlying it. But perhaps a person could just *pick* one in the way described, coming to a .9 credence (say) that the relevant person will get a corgi. Then just as we may say that the person in a Buridan's ass case selects rationally, so we may say that the person in that case selects rationally as well.

Indeed, this even seems to capture the phenomenology of a lot of cases that are a bit less typologically interesting than this one, in that we're not going from no credences to some credences. Thus you ask me to say how confident I am that such-and-such candidate will win the election, and every number within a range looks roughly good enough to me, but I end up saying “.6”, having picked it in a way that's not so different from the way I may pick a box of cereal. If we can do that in these cases, shouldn't we be able to do it in these other cases, too?

Of course, one can worry about the permissivism this reply seems to be assuming. (Permissivism is roughly the view that, for any given individual at a time with given total evidence, there are multiple overall belief states rational for that agent to adopt.) But in fact permissivism need never be assumed. Picking can be rational when two options seem indiscriminably good *to the agent*; they might still nevertheless in fact *be* differently good, say from the point of view of rationality. My aim is just to offer the anti-conservationist a psychological mechanism whereby the agent *can* actively and rationally get from one doxastic attitude-type to the other.

Even given all that, it's not immediately obvious *how*, exactly, this responds to the P_I–C_I-type argument. On one variant, picking this way allows us to reason to the doxastic attitude we pick. In that case, we would reject P₂. On

another variant, we might deny that this is reasoning to the new attitude, but by analogy with Buridan's ass cases, argue that we needn't be in a position to do that. That's because, in the Buridan's ass case, the person is rationally required to pick either A or B, even though (we are now assuming) they cannot reason to an intention to choose A or B. It's not obvious to me which variant is better, but either one constitutes a *prima facie* compelling response to the P_I-C_I-type argument.

Well, does this response actually work? I'm not very optimistic about it, though I have no knockdown argument against it. In Buridan's ass cases, a person can simply pick what to do because the action itself involves no *commitment*; the whole point is that of deciding to act nonintentionally is that you are not committed to the action you ultimately perform. But doxastic attitudes are essentially committal states. To believe that *p*, or to have a credence (precise or imprecise) of *x* (or [*x*, *y*]) that *p*, etc., involves a commitment to viewing the world a certain way. Minimally, we feel pressure of some kind not to violate our commitments; otherwise they would not be commitments. It's difficult to specify the exact nature of this pressure, or exactly how it works. But what you have simply isn't really a commitment unless you feel some pressure against violating it.

Now, suppose you know you seem to yourself simply to have *picked* a given credence, precise or imprecise. Would you feel any pressure to stick to it? Not only do I think you won't, but I think you *shouldn't*. Of course, sometimes we have reasons not to revisit pickings in the practical case: that would be very much a waste of our time. But suppose having picked doing one thing, you "accidentally" perform the other action, grab the other cereal indiscernible cereal box as it may be. You shouldn't feel any guilt or other reactive attitude toward yourself or your behavior—at most, you should be amused by the vicissitudes of your psychology in such situations. If you pick A, begin reaching for it, lose focus a bit and then find yourself going to B, you have no reason to divert yourself back to A. No real commitment was involved.

I think picking can at best generate guesses of various kinds, mental acts partly characterized by the fact that they *don't* involve commitment. ("If I had to guess..." elicits a guess because it cancels the suggestion of a commitment.) Though matters here are tricky, I think guessing that *p* rarely if ever leads to the

belief that p .

I say matters here are tricky because there is *some* reason to think that guessing can reliably lead to believing. Imagine someone asking:

(18) What kind of dog do you think she'll get?

One and the same person can answer with any of the following:

- (19)
- a. A corgi.
 - b. I think she'll get a corgi.
 - c. A corgi or a dachshund.
 - d. I don't know.

(19a) can be a guess, the person's best guess, which might still correspond to a very low likelihood. But as with (b), we can make our guess with apparently belief-reporting language. ('I believe' sounds worse here, I think, but I don't want to rest anything on that.) Insofar as 'think' and 'believe' are synonyms, we have reason to think that guesses can amount to beliefs.³⁹ Personally, I think the 'think' is being used *only* to make a guess; I think that because there wouldn't be the relevant commitment we expect of belief. But, as I said, matters here are tricky. So perhaps this response can be made to work. That is, perhaps picking a doxastic attitude *is* possible, because doxastic attitudes don't involve the commitment I think they do involve. As I said, I'm not optimistic, but it is one option for the anti-conservationist.

That is the extent to which I'll explore the argument for CONSERVATION OF ATTITUDES from the difficulty of reasoning from one type of doxastic attitude to another. There are, as I've tried to illustrate, a number of places it might be resisted, most easily simply by developing a clear and convincing counterexample involving good reasoning from one type of doxastic attitude to another. But perhaps it might also be resisted by developing the picking suggestion, or in some other way still.

³⁹ For an interesting theory that incorporates data like (19) to reach that conclusion, see Holguín (n.d.).

5 *Conclusion: Some Future Directions*

My aim, as I said, was exploratory: it was to fill in some of the contours of a debate that I think people should have, namely whether conservatism is true. I have only looked at one way to capture conservatism, CONSERVATION OF ATTITUDES, and only one argument for that thesis. So there is an enormous amount of work to do before it becomes especially clear whether conservatism is at the end of the day an attractive position. So, while I have no idea whether conservatism is right—my sympathies do incline in that direction, I should say—it strikes me as an important for epistemologists to work out whether it is.

In the spirit of exploration, I'll end with some pointers toward what I suspect will be rewarding next places to look. First, I think there may be a different route to CONSERVATION OF ATTITUDES, one that is more instrumentalist in character. To see how it might go, consider

NO POSITIVE EPISTEMIC DUTIES. There are no propositions p such that anyone at any time is epistemically required to believe p .

To get a general feel for (one) argument for NO POSITIVE EPISTEMIC DUTIES, consider that, in any evidential situation we may be in, there is an open-ended number of propositions that are sufficiently strongly supported by our evidence, at any level of stringency that 'sufficiently' may pick out. But we cannot be epistemically required to believe such a large number of propositions. In any given case, our interests can pick out some as worth special consideration; but that would go beyond strictly epistemic requirement.⁴⁰

There's a lot to say about whether such an argument is right, about the role that *interests* play in getting us to requirements for belief in specific propositions, and much more here.⁴¹ I want to bracket all of that right now. Just notice that there's no clear route from NO POSITIVE EPISTEMIC DUTIES to CONSERVATION OF ATTITUDES, or to any of CONSERVATION OF ATTITUDES's most important entailments (e.g., that we cannot be epistemically required to have credences). After all, an epistemic requirement to have beliefs *in general* or credences *in general* seems not to entail having any particular belief or credence. So any argument we give can't just be immediate from NO POSITIVE EPISTEMIC DUTIES.

⁴⁰ This is a rendering of Nelson (2010)'s argument, oversimplified in some important ways.

⁴¹ For some of the best recent discussion, in my opinion, see Friedman (2018).

But I do think there's a path, though filling it out at all persuasively will take a lot of work. The idea is roughly that forming different doxastic attitudes is, at least in the relevant cases, at least when belief involves rational deliberation about what to believe, a matter of answering a question that you ask yourself. Plato has Socrates voice something like this view of belief in the *Theaetetus*.⁴² Other philosophers have defended views like it.⁴³ Thus, to come to believe that p is for p to be your answer to some question you ask yourself. For example, to believe that your friend will get a corgi is to answer that she'll get a corgi when you ask yourself what kind of dog she'll get.

I think this view does a lot of explanatory work, especially in solving a difficult puzzle, but I'll say no more for it than that. The next thing to note is that if that's right, then it's very natural to think that when we form credences that ϕ , we do so in answer not to, e.g., $\lceil \phi? \rceil$, but rather, e.g., $\lceil \text{how likely is it that } \phi? \rceil$. Viewed from that perspective, forming a belief as opposed to forming a credence is just answering a specific sort of question rather than a different sort of question. And the last step in the argument is to say that we are *not* required to answer specific kinds of questions rather than others; epistemology does not tell us what to be curious about, wonder about, etc. This is a kind of instrumentalism about epistemic rationality: epistemology tells us how to achieve or aims in inquiring, but does not tell us *what* we ought to inquire into.⁴⁴ So it won't tell us to inquire into whether p or how likely p is, specifically. So it won't tell us to have beliefs or credences.

This is, of course, only the barest sketch of an argument. But it strikes me as very much worth looking into, as a kind of generalization of the argument against positive duties to have specific beliefs.

Finally, it would also be worthwhile to investigate whether conservatism will apply to non-doxastic attitude-types. In specific forms this debate has been

⁴² "The soul when it thinks is doing nothing other than dialoguing, asking itself questions and answering them itself, and affirming things and denying. And whenever it has determined something, either gradually or by leaping quickly, and affirms the same thing and does not disagree, we put that down as its doxa. So I call forming a doxa 'saying', and I call doxa a logos spoken not to another nor with voice, but silently to oneself" (*Theaetetus* 189e6–90a6). (I take this translation from Moss and Schwab (2019), page 10). See Moss (2014) for further discussion, and, e.g., *Philebus* 38b12–39a7 for another, similar passage.

⁴³ See, e.g., Hieronymi (2008) and Drucker (2020).

⁴⁴ For someone who seems to me sympathetic to this view of epistemology, see Friedman (2020).

happening for a very long time. Can we be rationally required to have a given desire solely because we have certain kinds of beliefs? But the issues are much more general than that, and it might be illuminating to introduce more abstract, structural considerations than people have in the past. And if conservatism is true for doxastic attitudes but not for non-doxastic attitudes, it would be very interesting to see exactly why, in order to help us to understand better the differences between doxastic and non-doxastic attitudes. So far, though, most of this work has yet to be done, because people have not paid very much attention to the specific structural features of attitudes other than belief, desire, and intention. I strongly suspect it would be illuminating to do so.

References

- Anscombe, G. E. M. 2000. *Intention*. Cambridge, MA: Harvard. Originally published 1957. 2
- Berker, Selim. 2018. “The Unity of Grounding.” *Mind* 127:729–777. 2
- Blackburn, Simon. 2010. “Dilemmas: Dithering, Plumping, and Grief.” In *Practical Tortoise Raising*, 47–63. Oxford, UK: Oxford. 20
- Boghossian, Paul. 2014. “What Is Inference?” *Philosophical Studies* 169:1–18. 6
- . 2018. “Delimiting the Boundaries of Inference.” *Philosophical Issues* 28:55–69. 5
- Briggs, R. A. and Pettigrew, Richard. 2020. “An Accuracy-Dominance Argument for Conditionalization.” *Noûs* 54:162–181. 1
- Briggs, R.A. 2009. “Distorted Reflection.” *Philosophical Review* 118:59–85. 9
- Broome, John. 2013. *Rationality Through Reasoning*. Malden, MA: Wiley-Blackwell. 1
- Carr, Jennifer Rose. 2020. “Imprecise Evidence without Imprecise Credences.” *Philosophical Studies* 177:2735–2758. 13, 16
- Chislenko, Eugene. 2016. “A Solution for Buridan’s Ass.” *Ethics* 126:283–310. 20

- Christensen, David. 2004. *Putting Logic in Its Place*. Oxford, UK: Oxford. 16
- . 2010. “Rational Reflection.” *Philosophical Perspectives* 24:121–140. 9
- Clarke, Roger. 2013. “Belief Is Credence One (in Context).” *Philosophers’ Imprint* 13:1–18. 8
- Dancy, Jonathan. 2018. *Practical Shape*. Oxford, UK: Oxford. 6
- Dogramaci, Sinan. 2018. “Rational Credence Through Reasoning.” *Philosophers’ Imprint* 18:1–25. 6, 7
- Drucker, Daniel. 2020. “The Attitudes We Can Have.” *Philosophical Review* 129:591–642. 25
- Easwaran, Kenny. 2016. “Dr. Truthlove or: How I Learned to Stop Worrying and Love Bayesian Probabilities.” *Noûs* 50:816–853. 3
- Elga, Adam. 2010. “Subjective Probabilities Should Be Sharp.” *Philosophers’ Imprint* 10:1–11. 16
- Friedman, Jane. 2013. “Rational Agnosticism and Degrees of Belief.” In Tamar Szabó Gendler and John Hawthorne (eds.), *Oxford Studies in Epistemology*, volume 4, 57–81. Oxford, UK: Oxford. 14
- . 2018. “Junk Beliefs and Interest-Driven Epistemology.” *Philosophy and Phenomenological Research* 97:568–583. 24
- . 2020. “The Epistemic and the Zetetic.” *Philosophical Review* 129:501–536. 25
- Gärdenfors, Peter and Sahlin, Nils-Eric. 1982. “Unreliable Probabilities, Risk Taking and Decision Making.” *Synthese* 53:361–386. 13
- Graham, Peter A. 2011. “‘Ought’ and Ability.” *Philosophical Review* 120:337–382. 14
- Hacking, Ian. 1967. “Slightly More Realistic Personal Probability.” *Philosophy of Science* 34:311–325. 7

- Hall, Ned. 1994. "Correcting the Guide to Objective Chance." *Mind* 103:505–518. 9
- Hawthorne, John, Rothschild, Daniel, and Spectre, Levi. 2016. "Belief Is Weak." *Philosophical Studies* 173:1393–1404. 8
- Hieronymi, Pamela. 2008. "Responsibility for Believing." *Synthese* 161:357–373. 25
- Holguín, Ben. n.d. "Thinking, Guessing, and Believing." Manuscript. 23
- Ismael, Jenann. 2008. "Raid! Dissolving the Big, Bad Bug." *Noûs* 42:292–307. 9
- Joyce, James M. 2005. "How Probabilities Reflect Evidence." *Philosophical Perspectives* 19:153–178. 16
- . 2009. "Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief." In Franz Huber and Christoph Schmidt-Petri (eds.), *Degrees of Belief*, 263–297. Springer. 3
- Kaplan, Mark. 1981. "Rational Acceptance." *Philosophical Studies* 40:129–145. 8
- Kelly, Thomas. 2003. "Epistemic Rationality as Instrumental Rationality: a Critique." *Philosophy and Phenomenological Research* 66:612–640. 4
- Kolodny, Niko. 2005. "Why Be Rational?" *Mind* 114:509–563. 17
- Konek, Jason. forthcoming. "Epistemic Conservativity and Imprecise Credence." *Philosophy and Phenomenological Research* . 3
- Kornblith, Hilary. 1983. "Justified Belief and Epistemically Responsible Action." *Philosophical Review* 92:33–48. 4
- Lasonen-Aarnio, Maria. 2020. "Enkrasia or Evidentialism? Learning to Love Mismatch." *Philosophical Studies* 177:597–632. 17
- Levi, Isaac. 1985. "Imprecision and Indeterminacy in Probability Judgment." *Philosophy of Science* 52:390–409. 16

- Lewis, David. 1980. "A Subjectivist's Guide to Objective Chance." In R. C. Jeffrey (ed.), *Studies in Inductive Logic and Probability*, volume 2, 263–93. Berkeley: University of California Press. 9
- . 1988. "Desire as Belief." *Mind* 97:323–332. 2
- Lyon, Aidan. 2017. "Vague Credences." *Synthese* 194:3931–3954. 13
- Moss, Jessica. 2014. "Plato's Appearance–Assent Account of Belief." *Proceedings of the Aristotelian Society* 114:213–238. 25
- Moss, Jessica and Schwab, Whitney. 2019. "The Birth of Belief." *Journal of the History of Philosophy* 57:1–32. 25
- Moss, Sarah. 2019. "Full Belief and Loose Speech." *Philosophical Review* 128:255–291. 8
- Nelson, Mark T. 2010. "We Have No Positive Epistemic Duties." *Mind* 119:83–102. 24
- Pettigrew, Richard. 2016. *Accuracy and the Laws of Credence*. Oxford, UK: Oxford. 3
- Rescher, Nicholas. 1959. "Choice without Preference." *Kant-Studien* 51:142–175. 20
- Schoenfield, Miriam. 2012. "Chilling Out on Epistemic Rationality: A Defense of Imprecise Credences (and Other Imprecise Doxastic Attitudes)." *Philosophical Studies* 158:197–219. 16
- . 2017. "The Accuracy and Rationality of Imprecise Credences." *Noûs* 51:667–685. 3
- Schwitzgebel, Eric. 2002. "A Phenomenal, Dispositional Account of Belief." *Noûs* 36:249–275. 11
- Searle, John R. 1970. *Speech Acts*. Cambridge: Cambridge. 2
- Seidenfeld, Teddy, Schervish, Mark J., and Kadane, Joseph B. 2012. "Forecasting with Imprecise Probabilities." *International Journal of Approximate Reasoning* 53:1248–1261. 3

- Sinnott-Armstrong, Walter. 1984. "‘Ought’ Conversationally Implies ‘Can’." *Philosophical Review* 93:249–261. 14
- Staffel, Julia. 2013. "Can There Be Reasoning with Degrees of Belief?" *Synthese* 190:3535–3551. 6
- Sturgeon, Scott. 2008. "Reason and the Grain of Belief." *Nous* 42:139–165. 16
- Swanson, Eric. 2016. "The Application of Constraint Semantics to the Language of Subjective Uncertainty." *Journal of Philosophical Logic* 45:121–146. 9
- Ullmann-Margalit, Edna and Morgenbesser, Sidney. 1977. "Picking and Choosing." *Social Research* 44:757–788. 20
- Valaris, Markos. 2017. "What Reasoning Might Be." *Synthese* 194:2007–2024. 6
- van Fraassen, Bas C. 1984. "Belief and the Will." *Journal of Philosophy* 81:235–256. 9
- Weatherson, Brian. 2019. *Normative Externalism*. Oxford, UK: Oxford. 17
- Wedgwood, Ralph. 2013. "Rational ‘Ought’ Implies ‘Can’." *Philosophical Issues* 23:70–92. 14
- Williamson, Timothy. 2007. "How Probable Is an Infinite Sequence of Heads?" *Analysis* 67:173–180. 11
- Yalcin, Seth. 2007. "Epistemic Modals." *Mind* 116:983–1026. 6
- . 2011. "Nonfactualism about Epistemic Modality." In Brian Weatherson and Andy Egan (eds.), *Epistemic Modality*. Oxford: Oxford. 6