Reasons without Persons: Reply to Döring and Eker, Snedegar, and Lenman

In their contributions to this symposium, Döring and Eker, Snedegar, and Lenman advance a variety of objections to the time-slice-centric theory of rationality I defend in *Reasons without Persons*. I am grateful for their thoughtful and incisive comments. I consider their arguments in turn.

1 Döring and Eker

Döring and Eker focus on one element of time-slice rationality, namely Synchronicity, the claim that all requirements of rationality are synchronic. In this reply, I will address their skepticism about the motivations I use to support Synchronicity, leaving aside their more specific arguments against it. I motivated time-slice rationality by appeal to a mentalistic version of internalism, on the one hand, and puzzle cases for personal identity over time, on the other. Döring and Eker argue that these considerations do not motivate Synchronicity. Let us take these in turn.

Internalism is, in my view, best motivated by the thought that rationality is a matter of believing and behaving in ways that are sensible, given your perspective on the world. Access internalists, on this way of thinking, hold that your perspective on the world is constituted by facts to which you have some kind of special access. Assuming that you lack the relevant special access to facts about the past (including facts about your past mental states), this

would straightforwardly motivate Synchronicity. But access internalism is increasingly out of favor, and for good reason. Arguably, there are no facts to which we have the relevant kind of special access (Williamson 2000), not even facts about one's present experiences. Mentalist internalism can be seen as conceiving of your perspective on the world as constituted by your mental states, regardless of whether you have special access to those states.

I think that on such a view, your perspective on the world is best seen as constituted by your *present* mental states, and in this way mentalist internalism would support Synchronicity. But Döring and Eker are skeptical and make several points in opposition. First, they seem to suggest that past experiences can sometimes partly constitute your present perspective. This strikes me as mistaken. Surely experiences that you had long ago and no longer remember are not part of your present perspective on the world. Insofar as more vivid or more recent experiences can seem to be part of your present perspective, that must be because you remember them. Better, then, to say that your present memories can be part of your present perspective, but that past experiences themselves cannot be.

They also note that your perspective evolves over time and that these changes might be governed by rational requirements. They present this thought as one that is consistent with the underlying motivation for internalism, which they dub PERSPECTIVE. But in my view they are in fact arguing against this motivation. After all, if there are rational constraints on how your perspective should change over time, and these constraints wind up making a difference to what you ought to believe at a time, then what you ought to believe does *not* just depend on your perspective on the world, but also on whether a given set of beliefs would be part of a permissible sequence of belief states. To see the tension with internalism, consider a case where you no longer remember what beliefs you had in the past. I think that internalists should say that those past beliefs do not at all constrain what you ought to believe now, whereas Döring and Eker's view seems to suggest that they do. After all, regardless of your having no memory of your past beliefs, some belief states you could adopt now would constitute permissible evolutions of your perspective, whereas others wouldn't.

Döring and Eker are correct though that I also think the best version of mentalist internalism is an evidentialist one. On this view, your evidence supervenes on your mental states and is the part of your perspective on the world that determines what you ought to believe. Pursuing their thought about rational constraints on the evolution of your perspective, might there be rational constraints not just on what to believe *given* your evidence, but also what your evidence should itself be? For instance, Williamson (2000) argues that your evidence consists of all and only the propositions that you know (E=K). But then, arbitrarily dropping beliefs can constitute a change in your evidence, and one might think this is sometimes irrational, even if you wind up with beliefs supported by your then-attenuated evidence. Similarly, Siegel (2012) argues that implicit bias can sometimes affect what evidence you acquire from a given experience, and that this is a failure of rationality, even if your resulting beliefs are those best supported by your skewed evidence. These phenomena might motivate diachronic norms governing evidence, though it may also be possible to deal with them using only sychronic norms. For instance, instead of a diachronic norm saying that you ought not arbitrarily drop beliefs and thereby lose evidence, we might instead adopt a synchronic norm saying that at a time, you ought to have all the beliefs such that, were you to have them, they would constitute knowledge, and hence evidence (or, more generally, that at each time you ought to possess all the evidence that you are in a position to have). So while there may be grounds for adopting rational requirements governing the possession of evidence, further argument would be needed to show that these requirements must be diachronic.

I turn now to the second motivation, namely puzzle cases for personal identity over time. In these cases, the facts about personal identity over time are sometimes murky and controversial, but I argued that even in those cases we can settle what attitudes a given agent ought to have without settling the facts about the metaphysics of personal identity over time. This makes it plausible that what you ought to believe at a time doesn't depend on facts about personal identity over time. And this, in turn, supports the claim that there are no diachronic requirements of rationality. Döring and Eker are unconvinced. They make

two claims: First, diachronic requirements of rationality can easily be reformulated to make reference to psychological continuity or some other surrogate notion. Second, there are also puzzle cases for personal identity at a time, and so a parallel argument to the one I make against diachronic norms could be made against even synchronic norms. Since this is absurd, something must have gone wrong in my argument.

Start with the first claim. Döring and Eker are correct, I think, that one of my arguments against reformulating diachronic norms in this way rests on an evidentialist presupposition that might not be granted by the defender of diachronic norms. I argued that it is not clear why facts about whether you are or are not psychologically continuous with some past time-slice should affect what you ought to believe now, given that such facts do not in general constitute evidence relevant to the question at hand (RwP 34). But one might object to this emphasis on evidence, as they do, and hold that facts about psychological continuity matter because rationality is partly about whether your perspective on the world evolves in the right sort of way. (Note, though, that even this way of talking presupposes personal identity, since it is the evolution of 'your' perspective on the world that is taken to be important.)

Nevertheless, I think Döring and Eker are too optimistic about the feasibility of reformulating diachronic norms so that they refer to psychological continuity, say. Psychological continuity comes in degrees. Thus we could set some threshold level of continuity at which the relevant diachronic norm kicks in, though the arbitrariness of the threshold threatens to reintroduce the explanatory problem. Why should facts about psychological continuity above this arbitrary threshold make such a big difference to what you ought to believe now? Alternatively, we might try to make the diachronic norms sensitive to degrees of psychological continuity. But it is not at all clear how to do so. Intuitively, the thought would be that what you ought to believe now depends on the beliefs of some past time-slice to the extent to which your current time-slice is psychologically continuous with that past time-slice. But making this intuitive thought precise is no easy task. Consider Conditionalization, which (put in terms of personal identity) says that if your past time-slice has credence function P and

you gained evidence E (and nothing stronger) in the meantime, your new credences should be $P(- \mid E)$. I see no attractive way to reformulate this norm so that it is sensitive to the degree to which you are psychologically continuous with this past time-slice.

Now turn to Döring and Eker on personal identity *at a time*. They argue that a parallel argument to the one I make against diachronic norms can be run against synchronic norms as well. They say that in cases where the facts about personal identity at a time are murky and controversial, we can still settle the facts about rationality without settling the facts about identity at a time, and so by parallel reasoning, rational requirements can't make reference to personal identity at a time, thereby ruling out synchronic norms as well. But this is absurd, and so my own argument must be unsound.

I do in fact discuss personal identity at a time, albeit briefly (RwP ch. 11). But I come to the opposite conclusion. I think that in cases where the facts about personal identity at a time are murky and controversial (for instance, split-brain patients and organisms with more distributed nervous systems, like cephalopods), it is in fact extremely difficult to know what to say as far as rationality goes. Indeed, it is hard to even know how to describe the case. Who is it that ought or ought not have a given belief? Which experiences are part of a given entity's evidence? And so on.

Consider Döring and Eker's case (from Olson 2003, 330) of Even and Odd:

> Imagine [...] a human being with two personalities, one of which is "out" or active on even days, and the other of which is active on odd days. Each day at midnight, like clockwork, one personality goes dormant and the other takes over. Call the being who thinks and acts on even days Even and the one who thinks and acts on odd days Odd [...]

They suppose that Even intends to ϕ but has no beliefs about necessary means to that end, while Odd believes that ψ-ing is a necessary means to ϕ-ing, but doesn't have ϕ as an end. Moreover, neither Even nor Odd intends to ψ. They then ask whether Even or Odd violate a synchronic norm of means-end coherence, and they say that 'The answer seems "crystal clear": they do not commit any rational mistakes by violating [means-end coherence]' (6).

I confess I do not share their verdict about the case. If it really is to be a puzzle case for personal identity at a time, then it seems question-begging to talk of Even and Odd as separate agents. If Even and Odd really are two separate agents, then I agree that there is no violation of means-end coherence by anyone. But if 'Even' and 'Odd' are names for modules or fragments of a single agent, then in order to determine whether there is any violation of means-end coherence, we need to settle what the whole agent believes and intends. And those facts are then quite murky. More generally, in puzzle cases for identity at a time, I feel the need to settle how many agents there are and which attitudes and evidence are had by which agent in order to settle facts about rationality. For this reason, I think that rationality doesn't depend on the metaphysics of identity over time, but it depends heavily on the metaphysics of identity at a time.

2 Snedegar

Snedegar argues that there are diachronic norms of practical rationality. While I address and rebut various motivations for diachronic norms of practical rationality in Chapters 3, 6, and 7, Snedegar's argument is novel. He defends a diachronic norm that requires you to fill in partial, or coarse-grained, plans with additional detail over time as the need arises. He also anticipates my response to his argument, namely that it is possible to account for the phenomena he highlights using only synchronic norms, but also raises worries for this alternative account. I begin by briefly reviewing Snedegar's argument, then turn to elaborating and defending the time-slice-centric alternative that he sketches.

Snedegar says that often you ought to make coarse-grained or partial plans. These are plans that specify an act but do not settle in minute detail the exact way in which to perform that act. According to Snedegar, these partial plans are useful for at least two reasons.

First, it can be cognitively demanding to make maximally specific plans. As limited agents, we are often unable to make maximally specific plans due to the immense amount of detail we would have to entertain in order to do so. And even when we are able to make

relatively specific (albeit perhaps not maximally specific) plans, trying to do so can consume so much cognitive resources that it is not worth our while.

Second, partial plans are more flexible, and hence can be more stable over time, than highly specific ones. And this stability in turn is important for plans to fulfill their roles in coordinating action. As Snedegar writes, 'intentions and plans let us coordinate both our own actions across time and our actions with those of others with whom we need to cooperate' (4). Suppose you need to go to the store in order to buy groceries to cook for dinner with your partner. If you make the highly specific plan to drive to the store at 8:35am, taking the scenic route, this plan may have to be abandoned in the face of information about roadworks on the scenic route, thereby throwing your dinner plans into doubt. But if you had just made the partial plan to go to the store sometime before dinner, you'll work around potential obstacles to particular ways of carrying out that plan and ensure that one way or another, you have the groceries in hand when needed.

I am skeptical of this second motivation for partial plans. For one thing, we might conceive of maximally specific plans as Gibbardian (2013) hyperplans which specify what to do in any possible contingency. They're like huge conjunctions of conditionals of the form *if situation X arises, do Y*. These plans needn't be abandoned in the face of recalcitrant information, for they already say how to go on when confronted with any possible new information. For another, even setting aside this conception of maximally specific plans, whether abandoning some particular plan would frustrate interpersonal coordination depends on what happens following the abandonment of that plan. If upon learning of the closing of the scenic route you would straightaway adopt some other plan for getting to the store, and if your partner knows this about you, then coordination is easily achieved despite your tendency to adopt highly specific plans and abandon them if necessary.

But I think that this consideration leads to a different motivation for sometimes adopting partial plans. Suppose that for some highly specific plans (even Gibbardian hyperplans), you are uncertain whether you'd actually carry them out. If you think that your abandonment

of such a plan would lead not to your adoption of a different plan which effectively serves much the same goal, but would instead lead to apathy or some other bad outcome, you may have good reason to instead adopt a partial plan that you think you'll actually see through to fruition. (I think this motivation for partial plans applies to ideally rational agents as well as non-ideal ones. For I doubt that ideal rationality requires having full control over one's future selves, and even if it does, ideally rational agents certainly needn't be certain that they have such control.[1]) For example, suppose that you have the goal of adopting a healthy lifestyle. There are many maximally specific plans for how to live your life, for instance different plans involving specific programs of diet and exercise. Suppose also that you think, for each such maximally specific plan, that there is a non-negligible chance that you will abandon it, and that having abandoned it, you will revert to your unhealthy ways. Better then to adopt some partial plan that you think you would actually follow through on.

I grant, then, the importance of partial plans, both for non-ideal agents and for ideally rational ones. But as Snedegar notes, if you start off with a partial plan, you'll need to eventually fill it in with enough detail to actually carry it out. At some point you'll have to settle on how and when to get to the store, or what workout program to start. He claims that planning is therefore governing by a diachronic norm he calls **Fill In**, which say that 'If at t_1 you have a partial plan to A at t_2, then by t_2 you ought to have sufficiently filled in that plan' (5).

I think that neither the need to make and fill in partial plans, nor the more general role of plans in coordinating actions across time and with other people, requires diachronic norms. These phenomena can be accounted for by a synchronic norm that says that at a given time, you ought to have the plan that is best supported by the reasons you have at that time (Snedegar's **Planning**). This, in turn, falls out of expected utility theory, coupled with my view that a decision-maker's options at a time are all and only the decisions (or plans) that she can make at that time.

[1] As Christensen (2007) argues, ideally rational agents needn't be certain of their own ideality.

For reasons noted above, this synchronic picture accounts for the rationality of making partial plans. Sometimes a partial plan is the one that is best supported by one's reasons (or, in decision-theoretic terms, the one that maximizes expected utility). Moreover, it accounts for the rationality of filling in plans over time, on the plausible assumption that if one's reasons do not come to support abandoning that partial plan, they will eventually come to support adopting some more filled-in version of that plan.

I also think that this synchronic picture can account for the role of plans in coordinating action across time and with other people. Snedegar, following Bratman, is correct that plans must exhibit a certain stability in order to play this coordinating role. But this stability does not need to be directly underpinned by diachronic norms on planning. First, we can think of the stability of plans as a fact about the nature of plans. Plans by their nature tend to persist over time, and that is part of what makes them a useful part of our mental toolkit. We do not need to go further and say that it is a requirement of rationality that our plans persist over time in some way. (See RwP, pp. 123-8 for further discussion.) Second, provided that one's reasons remain relatively stable over time, and provided that one's reasons at a time significantly narrow down the plans that it is rational to have in light of those reasons, it follows that if at each successive time you have plans are that permissible given your reasons, your plans will exhibit the requisite stability over time as a byproduct. Compare the epistemic case. If your evidence severely restricts the set of beliefs states that are rational for you to have (in the limit making it a singleton), then your beliefs cannot fluctuate wildly unless your evidence fluctuates wildly, or unless you often fail to have beliefs that are permissible in light of your evidence. Moreover, even if new reasons support a change in your plans, in most cases they will not support big enough changes to undermine your ability to coordinate with others. Return to the case of going to the store. If you learn that the scenic route is closed, this supports abandoning the plan of going to the store via the scenic route. But provided you retain your goal of making dinner with your partner, your new reasons will still support adopting some plan or other that gets you to the store in time.

Toward the end of his article, Snedegar raises two worries for the synchronic account of planning that he sketched and that I have elaborated upon. First, he writes that:

> One of the most important advantages for agents like us is that forming partial plans lets us put off thinking about and evaluating more detailed ways of carrying them out until we have more time. If the coarse-grained decision and the many fine-grained ways of carrying it out are all options, then in order to determine that what we have most reason to do is make the coarse-grained decision, we will have to consider all of these more fine-grained decisions, as well. (9)

Second, he writes that on my account, deliberation 'is in an important way a kind of higher-order deliberation *with* or *about* decisions or plans' and that the 'reasons we weigh up are reasons that support decisions themselves, rather than actions' (9).[2] This, he says, does not mesh well with how we ordinarily think about deliberation. Typically, deliberation seems to involve considering reasons bearing on the actions themselves, and not only on plans or decisions to perform actions.

In response to both of these points, I claim that we should think of this view of rational planning and rational decision-making as making claims about what plans or decisions one rationally ought to adopt, rather than about how one does, or ought, to reason one's way to a conclusion about what to do. It is a criterion of rightness, not a decision procedure, to use a distinction from normative ethics. And while deliberating by weighing up reasons bearing on actions themselves, rather than on plans or decisions, won't always lead one to adopt plan or decision best supported by one's reasons (after all, one would thereby be ignoring a whole class of reasons), this may be a rough and ready way of deliberating that is easier and more natural for cognitively limited agents and that at least won't typically lead us very far astray. I doubt Snedegar would have much quarrel with the distinction between criteria of rightness and decision procedures. But he worries at the end that a theory which involves such a disconnect between the considerations that actual determine what we ought to do, and those we tend to consideration in deliberation, 'may give a picture of rational agency that is

[2]He is not objecting that my account allows for state-given reasons for plans rather than content-given (or object-given) reasons; he agrees with Schroeder (2012) that the case of plans shows that state-given reasons need not always be the 'wrong kinds of reasons.'

not very useful or familiar for agents like us' (10). This is a significant issue and is related to some of the worries pressed forcefully by Lenman as well, which I attempt to deal with in the next section.

3 Lenman

Lenman objects that the picture of rationality I defend is insufficiently sensitive to the cognitive needs and limitations of real-life humans and instead presents a highly idealized view that thus fails to provide adequate guidance to creatures like us.

Lenman is correct that many elements of my theory are very idealized, or as I prefer to say, demanding. But only some are essential to time-slice rationality. For instance, some of the demandingness comes from the background Bayesian framework within which I am working. This framework assumes that rational agents have perfectly precise credences and complete preferences, whereas we humans almost certainly do not. But time-slice rationality could also be combined with a theory which allows rational agents to have more coarse-grained doxastic states as well as incomplete preferences. Indeed, time-slice rationality is in at least one way particularly congenial to such an approach. There are money pump-style arguments (or as I call them, 'diachronic tragedy arguments') purporting to show that imprecise credences and incomplete preferences are irrational, but I show in Chapter 6 that a time-slice-centric perspective yields a principled rebuttal of these arguments.

Moreover, there are some respects in which time-slice rationality is less idealized, or less demanding, than its alternatives. For instance, my view allows that forgetting, while unfortunate, need not be irrational (RwP, pp. 42-44). Similarly, the fact that we humans lack total control over our future selves motivates my time-slice view of a decision-maker's options. I argue that an agent's options consist of all the *decisions* she can make at the time and that which option she ought to make depends (*inter alia*) on how confident she is that she will carry out that decision.

One element of time-slice rationality that is indeed very demanding, and that Lenman

objects to, is uniqueness theses. In the epistemic case, uniqueness is the claim that for any body of evidence, there is a unique most rational doxastic state to have in response to that evidence. I endorse uniqueness for independent reasons, but here is the role it plays in my defense of time-slice rationality (briefly touched upon above). Suppose we take as a *datum* that rational creatures have belief states that are relatively stable and evolve smoothly over time rather than jumping around erratically. On a non-time-slice view, this *datum* can be accounted for using diachronic norms, such as Bayesian conditionalization. In effect, such a norm might say that while different people can permissibly differ with respect to their evidential standards (or ways of weighing up different theoretical virtues of simpliclty, explanatory power, and so on), each person must use the same evidential standards at each time in her life. You can choose your standards, but once you've chosen, you have to stick with them. By contrast, a time-slicer can account for the *datum*, more plausibly in my view, by endorsing uniqueness. In effect, this is to say: sure, you should use the same evidential standards at each time in your life, but that's not because there's some special reason in favor of stability, but rather because there's a unique best set of evidential standards, and you should always use that one.

A similar move can be made in the case of preferences. If we suppose that rational agents necessarily have preferences that exhibit a certain sort of stability over time (what I call the 'Alleged Datum,' RwP 149), this can be accounted for either through a special diachronic norm on preferences, or through a principle of uniqueness for preferences. And holding fixed the Alleged Datum, I prefer to account for it in the latter way. Lenman is right to think that the uniqueness principle for preferences is rather implausible. For that reason, I explicitly said that I do not fully endorse the principle. I do, however, think it is less implausible than might be supposed, and so I sketched how I think it could best be defended. Lenman is entirely correct, though, that that argument sketch relies at various points on very strong and contentious assumptions. I think that if uniqueness for preference is false, as it probably is, then the Alleged Datum is likewise false. It is no requirement of rationality that one's

preferences be stable over time in any specific way, though there are often good pragmatic reasons to try not to undergo too frequent and too dramatic shifts in one's preferences.

Could one endorse time-slice rationality without endorsing even uniqueness for beliefs? I think so. This would involve denying that rational agents necessarily exhibit cross-temporal stability of beliefs of the sort required by diachronic conditionalization. On such a permissivist but time-slice-first view, while it is irrational to *plan* to abandon your evidential standards in favor of some particular alternative standards, it need not irrational to in fact do so (cf. Easwaran 2012, Pettigrew 2016; see also Moss 2014). Then, at any given time, your beliefs ought to be those that are best supported by your evidence, according to the evidential standards you endorse at that time. But if your standards change, so should your beliefs. While I find uniqueness independently compelling and think that the best version of time-slice rationality is one that espouses uniqueness, it is possible for a permissivist to be a time-slicer too and thereby wind up with a less demanding, or less idealized, theory of rationality.

Nevertheless, the theory as I presented it is indeed very demanding. Is this really problematic? There are at least two reasons why one might think so. The first is that such a highly demanding theory violates the principle that *ought* implies *can*. For instance, we simply cannot have precise credences, let alone those that are uniquely rational given our evidence. But we must be careful about what sort of modality is in play. It is clearly not metaphysically impossible to have the precise credences best supported by one's evidence. Nor, I think, is it nomologically impossible. And the same goes for the other requirements of rationality that I espouse. So the thought must be that it is *psychologically* impossible to satisfy these requirements. But I find the notion of psychological possibility rather unclear, and I know of no one who has managed to make it more precise. More importantly, even granting some intuitive notion of psychological possibility, it is doubtful whether requirements of rationality should be such that it is psychologically possible to satisfy them. Greco (2012, 350) argues that 'a hopeless paranoid schizophrenic may be unable to give up the belief that he is the target of a Martian conspiracy' but '[t]hat he is incapable of giving up this belief does not mean that

he's not irrational for holding it.' It may also be psychologically impossible to avoid falling prey to the sorts of mistakes highlighted by research on biases and heuristics, but that does not mean that those mistakes are not failures of rationality. For these reasons, I think that the principle that *ought* implies *can* is true when the modality is metaphysical and perhaps even nomological, but false when the modality is psychological.

Second, it might be that with such a demanding theory, it is sometimes impossible to know what the theory demands of you, and in this sense fails to give usable guidance to real-world agents. But I am convinced by recent arguments, especially from Williamson (2000) and Srinivasan (2015), that no theory can be fully operationalizable, in the sense that you are always in a position to know what the theory demands of you. Regardless of whether our theory is a highly demanding one, like mine, or a more permissive one, you will sometimes be uncertain what you ought to believe and do by the lights of that theory.

This fact, if it is a fact, raises a deep question. If no theory can be fully operationalizable, why not just go for the theory that says that you rationally ought to believe all truths and disbelieve all falsehoods? Sure, you often aren't in a position to know how to conform to this norm, but as we've learned, that's everyone's problem. Here, I think we should rely on the slogan I used to motivate mentalist internalism, namely that rationality is a matter of believing and behaving sensibly given your perspective on the world. This means that what you ought to believe and do is a function of your perspective on the world, and not a function of how the world in fact is (thus ruling out the norm of believing all and only the truths).

In my view, then, the sense in which a theory should provide guidance is that it should be sensitive to your evidence and uncertainty about the world. This, I think, is compatible with its often being very hard, and sometimes even impossible, to figure out what is the rationally best response to that information. My opponent might claim that this is such an attentuated guidance-providing role as to be of no use whatsoever. There is some justice in this claim. But I my view, we should respond not by adopting some less demanding theory of the rational ideal, but rather by distinguishing between two quite different questions. First,

what ought one rationally believe, desire, and do at any given time? Second, given that no one will always do as they ought, what sorts of things can one do to try to come closer to approximating the rational ideal? My theory primarily addresses the first question. The second question is no doubt important, but it is important to note that the it is a largely empirical question. Figuring out what to do to try to become less prone to implicit bias and increase one's willpower, for example, requires attention to research in psychology. It cannot be answered by a pure theory of rationality. Lenman complains that this 'outsources' a lot of work that should be done by a theory of rationality. But it strikes me instead as a principled division of labor into largely *a priori* investigation of the rational ideal, on the one hand, and largely empirical investigation of effective means to better approximating that ideal, on the other. Both projects are important, but they are different projects.

So I think any theory of the rational ideal, even a less demanding one than mine, should be supplemented with a theory of what may be effective means of trying to better approximate the ideal. Then, even if the theory of the ideal doesn't provide much guidance that is implementable by ordinary agents, the theory of the means to better approximating the ideal will step in to fill this role.

We have seen, then, that time-slice rationality needn't be a particular demanding or idealized theory. Its essence can be divorced from the Bayesian framework I adopted, and it could perhaps even be divorced from the uniqueness thesis I defend in the epistemic case. But even the extreme demandingness of my particular version of the theory is not, I think, problematic, provided we recognize the importance of supplementing a theory of the rational ideal with a theory of how to better approximate that ideal given our contingent limitations.

References

Christensen, D. 2007. 'Does Murphy's Law Apply in Epistemology? Self-Doubt and Rational Ideals.' *Oxford Studies in Epistemology*, vol. 2, 3-31.

Easwaran, K. 2013. 'Expected Accuracy Supports Conditionalization—and Conglomerability and Reflection.' *Philosophy of Science* 80 (1):119-142.

Gibbard, A. 2003. *Thinking How to Live*. Cambridge, MA: Harvard University Press.

Greco, D. 2012. 'The Impossibility of Skepticism.' *Philosophical Review* 121 (3):317-358.

Moss, S. 2014. 'Credal Dilemmas.' *Noûs* 48 (3):665-683.

Olson, E. 2003. 'Was Jekyll Hyde?' *Philosophy and Phenomenological Research* 66 (2):328 - 348.

Pettigrew, R. 2016. *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.

Schroeder, M. 2012. 'The Ubiquity of State-Given Reasons.' *Ethics* 122 (3):457-488.

Siegel, S. 2012. 'Cognitive Penetrability and Perceptual Justification.' *Noûs* 46 (2):201-222.

Srinivasan, A. 2015. 'Normativity Without Cartesian Privilege.' *Philosophical Issues* 25 (1): 273-299.

Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.