

# A Reply to the Synchronist

(Forthcoming in *Mind*)

Abelard Podgorski

## 1. Introduction

On the face of it, in ordinary practices of rational assessment, we criticize agents both for the combinations of attitudes, like belief, desire, and intention, that they possess *at* particular times, and for the ways that they behave cognitively *over* time, by forming, reconsidering, and updating those attitudes. Accordingly, philosophers have proposed norms of rationality that are *synchronic* - concerned fundamentally with our individual time-slices, and *diachronic* - concerned with our temporally extended behaviour. In the former camp, we find familiar norms of consistency in belief and intention, of adherence to the probability calculus in credences, and of the constraint our current evidence places on our attitudes. In the latter camp, not quite so well-explored, we find Bayesian demands that we update our credences according to Conditionalization<sup>1</sup>, and requirements that our attitudes exhibit various sorts of stability across time.

The impulse to unsheathe Ockham's razor and trim the excess in our theory is strong, however. It is perhaps not surprising, then, that a handful of epistemologists have recently begun to question whether an account of rationality requires both kinds of norms. It is synchronic norms which have the deeper history, and with the notable exception of Barry Lam (2007)<sup>2</sup>, these would-be-barbers<sup>3</sup> have attempted to cast doubt on the plausibility of diachronic norms altogether. My aim in this paper is to address what I take

---

1 According to Conditionalization, one's credence in P after learning some evidence E should equal one's prior conditional credence  $\Pr(P/E)$ . See Teller 1972 for a canonical defense.

2 Lam argues for a thesis he calls *dynamicism*, the view that rational norms apply fundamentally to *changes* of belief. Although it is beyond the scope of this paper to argue, I join Lam in endorsing a purely diachronic approach to rationality, though the process-oriented picture of diachronic norms I sketch in this paper differs substantively from his view that such norms apply to mere attitude change.

3 See particularly Sarah Moss (unpublished, forthcoming), who explicitly sympathizes with this project, and David Christensen (2000), who does not, but who argues against a large class of diachronic norms.

to be the most direct and general recent attack on diachronic epistemic rationality, the arguments for so-called ‘time-slice epistemology’ by Brian Hedden (2015). I argue that Hedden's attempt to motivate the rejection of diachronic rational norms ultimately fails, and in particular that an independently attractive view about the nature of such norms, namely one on which such norms govern *processes*, escapes his assault unscathed.

## 2. Time-Slice Rationality

As Hedden understands it, time-slice epistemology involves the conjunction of two claims:

*Synchronicity*: What attitudes you ought to have at a time does not directly depend on what attitudes you have at other times.

*Impartiality*: In determining what attitudes you ought to have, your beliefs about what attitudes you have at other times play the same role as your beliefs about what attitudes other people have. (p. 4)

Although Hedden's paper largely consists in the application of arguments for Synchronicity and Impartiality to refute two specific norms, Conditionalization and Reflection, it is clear that Hedden takes the larger project of time-slice epistemology to be inconsistent with the existence of diachronic norms of rationality altogether<sup>4</sup>. Because my interest is in defending the possibility of diachronic norms, I will not discuss Impartiality, which functions mainly as a constraint on *synchronic* norms and which I take to be well-motivated<sup>5</sup>. I will focus attention instead on Synchronicity and the arguments meant to establish it, showing in the next section that even if the claim were true, it would not be sufficient to motivate the rejection of

---

4 Hedden informally characterizes time-slice epistemology as the thesis that ‘the relationship between time-slices of the same person are not importantly different, for purposes of rational evaluation, from the relationship between time-slices of different persons’ and that ‘the locus of rationality, so to speak, is the time-slice rather than the temporally extended agent’ (p. 1) and elsewhere as committed to the claim that ‘All requirements of rationality are synchronic.’ (p. 3)

5 Motivation for Impartiality can be found in Christensen 1991 and Arntzenius 2003.

diachronic norms (even, plausibly reformulated, those Hedden explicitly addresses), and sketching a view of such norms that I will defend in the face of the considerations Hedden marshals against them.

### 3. The Argument from Internalism

Hedden gives two arguments for Synchronicity. The first I will consider is an argument from *Internalism*. The core internalist intuition, as Hedden presents it, is that ‘being rational is a matter of believing and behaving sensibly, given your perspective on the world’ (p. 4). What is rational for an agent, according to the internalist, supervenes on her perspective. But, he argues, one's perspective on the world at a time is surely constituted by what one's mental life is like *at that time*. So, it looks like it follows that what is rational for you to believe at a time is fully determined by what your mental states are at that very time.

The internalist intuition is not uncontroversial, and Hedden does not himself go so far as to endorse it outright, but he notes that this formulation of internalism is weak enough to be compatible even with Timothy Williamson's view that what is rational to believe is determined by what one *knows*, given that knowledge is a mental state. (p. 5) I myself take the internalist intuition to have a compelling basis and will not reject it here<sup>6</sup>. Indeed, we may even safely grant that the argument just presented is sound. The problem, I suggest, is that Hedden's formulation of Synchronicity is simply too weak to motivate the rejection of diachronic norms. Recall: Synchronicity states that what attitudes an agent ought to have at a time is wholly determined by their mental states at that time. For this to fully vindicate a purely synchronic model of rationality, however, an additional claim must be added: *facts about what attitudes agents ought to have at particular times exhaust the demands of epistemic rationality*. And this, I argue, the diachronist should reject.

The mistake Hedden makes is much like the one underlying Zeno's infamous paradox of the arrow. Zeno notices that at each moment in time, an arrow does not change its position - it merely occupies its own space. So at each instant, the arrow does not move. Since there is no instant at which the arrow

---

<sup>6</sup> For a typical challenge, see Goldman 1999. It is defended in Conee and Feldman 2001.

moves, he concludes, it never moves.

Analogously, Hedden argues, according to internalism, what is rational for an agent at a time is determined by what they are like at that time. So at each instant, what is rational for an agent is determined purely synchronically. Since what is rational for an agent at each instant is determined purely synchronically, he implicitly concludes, rationality is purely synchronic.

But the correct diachronist response here is likewise analogous to the proper response to the paradox of the arrow. The final step in either argument is invalid - just because there is no instant at which a phenomenon occurs does not mean that it never occurs. For there may be phenomena which are *essentially diachronic*, properties that temporal slices cannot possess but which temporally extended objects can. Motion is such a feature. And diachronists should say that rationality and irrationality, in at least some of its forms, is also such a feature. Some rational requirements, they can claim, are like the policeman's command 'Don't move!' There is no individual time-slice at which one violates this command, only intervals during which one does so. Such norms tell us what is rational for agents *over intervals* in a way not reducible to what is rational for them *at times*.

This does not mean abandoning or trivializing the internalist insight. We may insist that the internalist claim, that what is rational for an agent supervenes on their perspective, is perfectly true. And it is perfectly true that what your perspective is at a moment is determined by what your mental states are like at that moment. But this is just a special case of a more general truth: what your perspective is like over *any* interval of time is determined by what your mental states are like during that interval. It follows that your rationality during an interval supervenes on your mental states during that interval. So the constraint internalism places on diachronic requirements is this: whatever fundamentally diachronic cognitive phenomena are assessable for rationality, their rationality will supervene on the agent's mental life during the interval in which they occur.

Hedden misses this because the diachronic norms he considers are formulated as requirements on

the *attitudes* one must have at a time, in virtue of facts about one's mental *history*, rather than requirements on essentially diachronic phenomena, in virtue of one's concurrent mental life. That is, they are *historical* norms on *states* rather than *ahistorical* norms on temporally extended objects.

It could not plausibly be suggested that norms of the latter sort are outside the scope of his paper, for even his explicit target, Conditionalization, can be formulated so as to claim nothing about what attitudes are rational at individual times, instead describing only what *changes in attitude* are rational or irrational, a change being a paradigmatic example of an essentially diachronic unit.

Though the view that rationality applies to mere changes is immune to the argument from internalism, it is not the one I would like to defend. I suggest a different approach to diachronic norms. Just as we expect synchronic norms to be requirements on the most natural synchronic cognitive units, *attitudes* such as belief, we should look for diachronic norms to govern the most natural diachronic cognitive units - *processes* such as reasoning. The notion of a process is a thicker one than that of mere change; it includes, significantly, a causal element. Norms on mere attitude change will not, in contrast to norms on processes, be able to distinguish shifts in opinion brought upon by pristine reasoning from those brought upon by repeated lightning strikes to the head. The view I propose, then, is that diachronic norms govern *processes*, temporally extended, causally continuous patterns of mental states. In light of internalism, those requirements will supervene on internal features of the agent during those processes.

This sort of view is one that Hedden and other recent time-slice theorists do not explicitly address - their paradigm candidates for diachronic requirements are either norms that apply to mere changes, like Conditionalization, or norms that require broad coherence between our cognitive behaviour now and our cognitive behaviour in the (possibly remote) past (Hedden 2013, 2015, Moss forthcoming). We have seen that by taking the norms to apply to diachronic phenomena rather than states, our picture avoids the argument from Internalism. In what follows, we will see how the causal patterns approach in particular has the resources to deflect Hedden's second, and in his eyes, more central, objection.

## 4. The argument From Personal Identity

Hedden's second argument goes something as follows: to determine whether a diachronic principle is satisfied, one needs to know facts about the agent's personal identity over time. But one can know all there is to know about what an agent rationally ought to believe without settling the facts about personal identity. So one can know all there is to know about an agent's rationality without invoking diachronic principles.

To defend this argument, Hedden invites us to consider a case of fission such as those described by Derek Parfit (1971, 1984). An agent (Pre) steps into a teletransporter, which vaporizes her body and creates two duplicates (Lefty and Righty) in separate cities. It is not obvious what happens to Pre - whether she survives as Lefty, or Righty, or both, or neither. But to determine what Lefty and Righty rationally ought to believe, he claims, we do not need to know whether they are identical with Pre - we just need to know their current evidence. A diachronic principle like Conditionalization, which constrains future credences by past ones, would require us to settle the question of identity before settling what Lefty ought to do.

## 5. R-Relatedness

Hedden anticipates a response that arises naturally from Parfit's own discussion of the fission cases. The objector rejects the significance of identity in favor of the significance of some psychological relation, call it 'R', which both duplicates may bear to Pre. Since whether R holds is settled in the case described, the case is no counterexample, provided the diachronic norms govern how we must be related to our R-ancestors rather than our past selves. But Hedden thinks that this response fails for two reasons.

First, as Parfit notes (1984, p. 298) the R-relation comes in degrees. So, Hedden suggests (p. 7), it is natural to expect the degree to which a person's time-slices exhibit the R-relation to have some upshot for the way they are rationally assessed. But it is hard to see how rational requirements like Conditionalization

can plausibly be made sensitive to these matters of degree.

Second, Hedden thinks there is an explanatory challenge - the defender of the R-relation account should explain why the R-relation has its unique significance for rational assessment - why collections of time slices united by the R-relation are importantly different, from the point of view of rationality, from other collections of time slices. He is skeptical that the challenge could be answered.

These worries may have force against norms like Conditionalization, but if we understand the R-relation in the usual way, as some kind of causal psychological connectedness, the view about diachronic norms I suggested in the previous section, on which they are norms governing *processes*, can answer both challenges. We can both provide a natural account on which rational assessment is sensitive to differences in degree of R-relatedness, *and* explain why the R-relation in particular distinguishes those collections of time-slices that are subject to rational norms from those that are not.

Notice that R-relatedness, on this picture, is a matter of how ones' mental states are causally related. So, differences in degrees of R-relatedness go hand in hand with differences in which causal patterns of mental states are exhibited. Processes, like reasoning, deliberation, and belief formation, in turn, just *are* causal patterns of mental states. Since these are exactly the things which are, on our view, the fundamental objects of diachronic rational assessment, facts about rationality will closely depend on the strength of the R-relation among time-slices.<sup>7</sup>

This same feature promises an explanation for why R-relatedness between time-slices, among the countless relations time slices might bear to each other, has special rational significance. Hedden's skepticism makes sense if one assumes that all diachronic norms resemble Conditionalization. Whether an agent conforms to Conditionalization depends only on what the temporal series of their credal mental

---

<sup>7</sup> It is worth adding: it is not clear that R-relatedness being a matter of degree generates a problem unique to the diachronist. *Any* view on which norms of any sort apply to subject matter vulnerable to Sorites or vagueness concerns will run into a similar problem. And even those features that distinguish agential *time-slices* from time-slices that are not rationally assessable will exhibit vagueness and continuity at the margins. So more needs to be said about why the diachronist is at a special disadvantage here.

states looks like, and not on any deeper relations between them. An agent can satisfy Conditionalization by having their credences rearranged by an appropriate sequence of lightning strikes. So it is indeed a mystery on such a view why rationality would only govern those collections of time-slices connected by the R-relation. But this stems from a feature of Conditionalization inessential to diachronic norms. On the process account, the explanation is simple. R-relatedness is a matter of causal psychological connectedness; this connectedness is a matter of the causal relations between states; and causal patterns of states is precisely what diachronic rationality is all about. So it's no mystery at all why only collections bound by the R-relation would be the proper subjects of diachronic rational requirements.

## 6. Synchronic Sufficiency

The R-relation response allows us to explain why we do not need to settle facts about personal identity to settle facts about rationality in the fission case. But Hedden's argument has a second layer. In addition to the claim that facts about identity are not *necessary* to account for the rationality of belief, Hedden suggests (p. 4, p. 7) that facts about the synchronic relation between each time slice and its evidence are *sufficient* to settle facts about rationality. If he is right about this latter claim, then appeal to the R-relation will not be enough to defend diachronic norms, for while the facts about R-relatedness may not, in the test cases, be in dispute in the way facts about identity are, neither are they doing any work.

The appeal Hedden makes to justify this thought is something like this: consider each time-slice and the evidence it possesses. Ask yourself “what ought this time-slice rationally believe?” Intuitively, Hedden expects, we will say “it ought to believe exactly whatever its evidence favors.” The matter is thus settled without need for more information. So purely synchronic considerations are sufficient to account for the rationality of belief, and the time-slice picture is vindicated.

But here we recall the main lesson gleaned from evaluating the Internalism argument: showing that what attitudes it is rational to have at particular times is determined synchronically is not sufficient to show that rationality is synchronic. For some rational requirements may apply not to the rationality of attitudes,



but to the rationality of diachronic phenomena like belief *change* or *reasoning*. Synchronicity, as formulated by Hedden, is too weak to refute the existence of diachronic norms. This response applies with equal strength to the argument here. Even granting that the evidence of each time slice suffices for what attitudes are rational for that time-slice, there may be *other* questions we can ask about rationality that are not so easily dispensed with. We may ask whether a certain instance of belief formation was rational, whether someone is reasoning rationally, whether a certain *pattern* of attitudes they exhibit is rational, and so on. And these are not straightforwardly answered merely by looking at isolated time-slices.

Moreover, there are positive reasons to think that important facts about rationality *cannot* be settled synchronically. I will give several examples concerning the rationality of *belief formation*.

## 7. Belief Formation

On the face of it, the synchronist has an explanation handy for the rationality of belief formation: belief formation is a way we bring ourselves into compliance with synchronic norms when our evidence changes. We gain some new evidence, our total evidence now supports a new belief, and we bring ourselves to comply with the synchronic demand that our beliefs match our evidence by forming a new belief. This picture, I suggest, is mistaken. Rational belief formation cannot be explained by appealing to synchronic relations of evidential support, falsifying the sufficiency claim.

I will be using a very weak assumption: that in the normal course of things, for finite human agents like us, responding to evidence takes time. For our purposes, this time may be vanishingly small. It is sufficient that the time at which we initiate forming a belief in response to our evidence and the time at which it is fully formed are not identical.

Imagine that my friend Minnie promises that she will come to my birthday party. Like most of my friends, Minnie is an odd duck; she is a pathological liar who delights in making promises she intends never to keep and rejoices in the disappointment of others. So the fact that Minnie tells me that she is coming to my birthday party is normally excellent reason to believe she will not. However, Minnie is also extremely

superstitious, and believes that breaking a promise made on the thirteenth of each month will curse her. So she always keeps promises made on the thirteenth of the month, which, it so happens, is today's date. All this is known to me. Having incorporated this evidence at  $t_0$ , suppose that the earliest time at which I can deliberately fully form a belief regarding Minnie's presence or absence (about which I am, at  $t_0$ , agnostic) is  $t_1$ .

Now imagine that, unbeknownst to me, exactly at  $t_1$ , I will suddenly forget that it is the thirteenth (so that I do not have this knowledge at  $t_1$ ). What belief is it rational for me to form at  $t_1$ ?

According to the synchronist, the belief that it is rational for me to form at  $t_1$  is the belief that Minnie will not be at the party, since this is the belief that is supported by the evidence I will have at that time. But, I claim, this is the wrong result. Because what belief I form at  $t_1$  is determined by the process of belief formation that operates before that time, forming the belief that Minnie will not be at the party would require me to, before  $t_1$ , ignore the perfectly compelling evidence I have that she will. It is true that at  $t_1$ , my epistemic state has changed in a way that may require me, going forward, to cease believing that Minnie will come to my party. But *at*  $t_1$ , I have not yet had any chance to respond to this sudden epistemic impoverishment, and so my failure to take it into account cannot be rationally impugned.

We may imagine a variation of this case which differs only in that I know, beforehand, that I am about to forget the decisive evidence concerning Minnie's reliability. I am deciding now what to believe, knowing that my current evidence supports Minnie's presence, but that by the time I form a belief, my impoverished evidence will support her absence. Which belief should I decide to form? Again, I maintain, forming the belief that is synchronically supported by my evidence at  $t_1$  is bizarre. I am deciding *now* what belief to form, am now better informed than I will be at  $t_1$ , and have every reason to think Minnie will be at the party. Whatever belief I form will be formed as a response to my *current* mental state, so my choice is between using all the information now available to me or effectively handicapping myself by treating perfectly good evidence as though it had no weight. To reject a belief we have every reason to think is true

in favor of a belief that coheres with evidence we know to be misleading is, it seems to me, to disrespect the fundamental epistemic concern with truth.

In both of these cases, then, it looks rational to form the belief that Minnie will make it to my party at  $t_1$ . But either case would be enough. By  $t_1$ , it is too late for that belief to bring me to satisfy a synchronic relation of evidential support. So if the judgment I suggest in either of these cases is correct, the synchronist is not in a position to explain the rationality of belief formation. The lesson here is that belief formation is something that happens *going forward*. But the attitudes that rationalize it, the evidence one is *responding* to, when one is being properly evidence-responsive, is evidence one has during the process of deliberation, *before* the belief is formed. The rationality of belief formation cannot be captured merely by looking at what attitudes are justified synchronically before, during, and after deliberation.

Another kind of case underscores the importance of causal relations between temporal parts in the rationality of belief formation. Consider the following two worlds. In the first, an agent **A** performs a typical act of good reasoning, forming a new belief in response to some newly acquired evidence **E**. In the second, there are two agents, **B<sub>1</sub>** and **B<sub>2</sub>**, similar to **A** except for the following bit of history: **B<sub>1</sub>** gains the evidence **E** and begins the very same act of reasoning completed by **A**, but halfway through, after the evidence has made its causal contribution but before the new belief is formed, is hit by lightning, causing them to forget **E** and cease the reasoning. Elsewhere in the world, **B<sub>2</sub>** is struck by lightning, causing them to gain or remember **E** and scrambling their brain as though it were halfway through the aforementioned act of reasoning, so the evidence makes no causal contribution but the belief is formed to completion. On the time-slice model, there should be nothing relevant to rationality that happens in the first world that does not also happen in the second world - cutting a sequence of time slices from **B<sub>1</sub>** and setting them against a sequence from **B<sub>2</sub>** allows us to replicate the pattern of consecutive mental states experienced by **A**. But this seems wrong. **A** formed a belief *in response to her evidence*. And this is an event of normative significance. This event does not occur in the second world - **B<sub>1</sub>** does not form a belief because she is

interrupted, and  $B_2$  forms a belief, but not in response to her evidence. If this is right, there is a certain kind of event - the formation of a belief as a causal response to evidence, whose rational properties are not reducible to the rational properties of its instantaneous parts. This kind of event is the sort of essentially diachronic object of rational assessment that only a correspondingly diachronic set of norms can adequately describe.

## 8. Uniqueness

I have argued that Hedden's positive arguments for the time-slice picture and, consequently, the rejection of diachronic rational norms, do not succeed. Hedden has one final explicit aim in the paper relevant to our concerns here: to show that a synchronic view has the resources to capture intuitions about rational agency that on their face seem best explained diachronically. In particular, synchronic norms look ill-poised to explain why agents should exhibit stability in their beliefs and credences over time. Individual time slices may differ drastically in their attitudes and still be, in isolation, rational. On the time slice picture, any sequence of such time-slices will exhibit no rational failure. Intuitively, though, agents who fluctuate wildly in their attitudes are not rational.

Hedden argues (pp. 14 -17) that we can accommodate this intuition on a synchronic picture, provided we accept *Uniqueness*, the claim that there is only one rational set of attitudes to have given a set of evidence. If Uniqueness is correct, then provided we do not gain or lose significant amounts of evidence, a rational agent's beliefs will remain relatively stable.

The problem with this is not just that Uniqueness is at best highly controversial<sup>8</sup>. Even granting it, the principle will not rule out the rationality of intuitively objectionable instability without a complementary notion of “evidence”. On one view, one's evidence is determined by states like belief. This will be no help to the synchronist, however, since an agent who wildly fluctuates in their beliefs is *thereby*

---

<sup>8</sup> Proponents of the thesis include Roger White (2007) and Ballantyne and Coffman (2011), but a rebuttal can be found in Kelly (forthcoming), and Schoenfield (2012) argues for a qualified rejection of uniqueness in favor of a moderate permissivism. Chapter Three of this dissertation discusses the issue of uniqueness in more depth.

fluctuating in their evidence. It is not much better if evidence is *knowledge*, since plausibly, wild fluctuations in beliefs can, in various ways, undermine knowledge as well. If evidence is something like perceptual experiences, then on a model that limits such evidence to the present time-slice, one's evidence will simply be too sparse to justify much at all. Hedden needs an account of evidence where it is both plausibly the sufficient grounds for our attitudes *and* resists being gained or lost through an agent's bizarre cognitive behaviour, and there does not seem to be one in easy reach. Meeting this challenge matters for Hedden's project because as long as a synchronic understanding of the irrationality of radical cognitive instability looks out of our grasp, we have an additional reason to be skeptical of the time-slice view.

## 9. Conclusion

I conclude, then, that there is no reason to think that diachronic norms have unacceptable implications concerning cases where personal identity is a matter of dispute, no reason to think that they violate plausible constraints of internalism, and no reason to think that they can easily be done away with without sacrificing an explanation of apparently rational cognitive inertia. Hedden's arguments rest on the assumption that it is an agent's attitudes at individual times that are the sole fundamental target of rational assessment - an assumption the diachronist should reject. Furthermore, we have sketched a picture of diachronic norms, one on which such norms govern cognitive processes like belief formation and reasoning, and shown that it is independently plausible and especially well placed to answer the criticisms leveled by time-slice epistemologists. It may be time, then, for enthusiasts of parsimony in the realm of rational norms to accept that the rational man is more than the sum of his parts.

## References

- Arntzenius, Frank 2003: 'Some Problems for Conditionalization and Reflection'. *Journal of Philosophy*, 100, pp. 356-370.
- Ballantyne, Nathan and Coffman, E.J. 2011: 'Uniqueness, Evidence, and Rationality'. *Philosophers' Imprint*, 11, pp. 1-13.
- Christensen, David 1991: 'Clever Bookies and Coherent Beliefs'. *Philosophical Review*, 100, pp. 229-247
- 2000: 'Diachronic Coherence versus Epistemic Impartiality'. *Philosophical Review*, 109, pp. 349-371.
- Conee, Earl, and Richard Feldman 2001: 'Internalism Defended'. *American Philosophical Quarterly*, 38, pp. 1-18.
- Goldman, Alvin 1999: 'Internalism Exposed'. *Journal of Philosophy*, 96, pp. 271-293.
- Hedden, Brian forthcoming: 'Time-Slice Rationality'. Forthcoming in *Mind*.
- 2013: 'Options and Diachronic Tragedy'. *Philosophy and Phenomenological Research*, 88.
- Kelly, Tom forthcoming: 'How to be an Epistemic Permissivist'. Forthcoming in Matthias Steup and John Turri (eds): *Contemporary Debates in Epistemology*. Malden, MA: Blackwell.
- Lam, Barry. (2007) 'The Dynamic Foundations of Epistemic Rationality'. Dissertation, Princeton University.
- Moss, Sarah. (Unpublished) 'Credal Dilemmas'
- forthcoming: 'Time-slice Epistemology and Action Under Indeterminacy'. Forthcoming in *Oxford Studies in Epistemology*.
- Parfit, Derek A. (1971). 'Personal Identity'. *The Philosophical Review*, 80, pp. 3-27.
- 1984: *Reasons and Persons*. Oxford: Oxford University Press.
- Schoenfield, Miriam. 2012: 'Chilling Out on Epistemic Rationality'. *Philosophical Studies*, 158, pp. 197-219.

Teller, Paul 1973: 'Conditionalization and Observation'. *Synthese*, 26, pp. 218–258.

White, Roger 2007: 'Epistemic Permissiveness'. *Philosophical Perspectives*, 19, pp. 445-459.