

Don't stop believing

Jennifer Rose Carr

Department of Philosophy, University of Illinois at Urbana–Champaign, Champaign–Urbana, IL, USA.

ABSTRACT

It's been argued that there are no diachronic norms of epistemic rationality. These arguments come partly in response to certain kinds of counterexamples to Conditionalization, but are mainly motivated by a form of internalism that appears to be in tension with any sort of diachronic coherence requirements. I argue that there are, in fact, fundamentally diachronic norms of rationality. And this is to reject at least a strong version of internalism. But I suggest a replacement for Conditionalization that salvages internalist intuitions, and carves a middle ground between (probabilist versions of) conservatism and evidentialism.

ARTICLE HISTORY Received 14 November 2015; Accepted 18 November 2015

KEYWORDS Epistemology; conditionalization; diachronic norms; evidentialism; evidence

Epistemic rationality requires two kinds of coherence. Broadly speaking, an agent's beliefs must fit well together at a time, and also fit well together over time. At any particular time, we should avoid believing contradictions, believe the consequences of our beliefs, and so on. And over time, we should respect the evidence we've received and adapt our beliefs to new evidence.

The traditional Bayesian picture of epistemic rationality is simply the conjunction of a synchronic claim and a diachronic claim:

Synchronic coherence: Rational belief states form a probability function and are rationalized by one's evidence.

Diachronic coherence: Rational belief states evolve by retaining old certainties and conditioning on new evidence. Recently, a number of philosophers have pushed for the abandonment of diachronic coherence norms. Norms like conditionalization, that have traditionally been understood as constraints on beliefs at different times, have lately been reinterpreted as purely synchronic constraints. On this view, the norms of epistemic rationality apply only to time-slices of individuals.

CONTACT Jennifer Rose Carr  jrcarr@illinois.edu

© 2016 Canadian Journal of Philosophy

I want to resist this movement. I'll argue for the following claim:

Diachronic Rationality: There are diachronic norms of epistemic rationality. The problem that the opponent of diachronic rationality poses is this: diachronic norms of epistemic rationality are in tension with some form of *epistemic internalism*. Epistemic internalism, generically, is the view that whether or not an agent is epistemically rational supervenes on facts that are 'internal' to the agent. The relevant sense of 'internal' can be cashed out in a variety of ways, generating a variety of internalist theories. If there are diachronic norms of epistemic rationality, then whether an agent is epistemically rational *now* is partly determined in part by the agent's past. Facts about the past are not, in some epistemically important sense, internal.

I argue that the rejection of diachronic norms incurs a number of serious problems: most importantly, that it permits discarding evidence, and that it treats intuitively irrational agents as epistemically ideal.

In Section 1, I explain the framework in which much of my discussion takes place, i.e. the Bayesian view of rationality. I then flesh out the objection to diachronic epistemic norms, some of its common motivations, and how the debate is situated within epistemology. I introduce a few important distinctions.

In Section 2, I offer a series of objections to time-slice internalism. First, Sections 2.1 and 2.2 argue that time-slice rationality entails that discarding evidence and erratically changing one's beliefs without new evidence are rational. Then in Section 2.3 I argue against the view that epistemic *ought*-implies-*can*, and address the claim that cognitive limitations somehow limit our epistemic liability. I describe a notion of relative rationality, which allows us to accommodate many of the intuitions cited in favor of time-slice internalism. Section 2.4 argues that there are normative differences between agents who conform to diachronic norms and those who don't. The opponent of diachronic norms is committed to a strong claim: that no agent can ever be rationally worse than another in virtue of purely diachronic differences between them. There are intuitive counterexamples to this generalization. I discuss the ways in which the diachronic position can be weakened while still maintaining the existence of diachronic norms of epistemic rationality. Finally, in Section 2.5, I discuss the conception of rationality as subjective epistemic good, and the elaboration of this theory in epistemic decision theory.

Section 3 discusses an objection to diachronic norms prohibiting information loss. What if one can ensure a net gain in information only at the cost of losing some information? I discuss diachronic norms that can accommodate the idea that this sort of informational tradeoff can be rational. I conclude briefly in Section 4.

1. The conflict

1.1. Bayesianism

I will assume a partial belief framework. (Nothing hinges on this.) On this view, beliefs come in degrees ('credences'). Credences fall in the interval $[0, 1]$, where

credence 1 represents certain belief and credence 0 represents certain disbelief. A person's total belief state is represented by a credence function, i.e. a function from propositions to real numbers in the unit interval.

According to the classical Bayesian picture, rational agents conform to both synchronic and diachronic rationality constraints. The primary Bayesian synchronic constraint is probabilism:

Probabilism: a rational agent's credences form a probability function.¹The traditional diachronic Bayesian constraint is conditionalization:

Conditionalization: let E be the strongest proposition an agent learns between t and t' . Then the agent's credences should update such that²

$$Cr_{t'}(\cdot) = Cr_t(\cdot \mid E)$$

One of the consequences of conditionalization is that once an agent rationally learns a proposition, she can't rationally unlearn it. One can't rationally lose information. (The set of live possibilities only shrinks.)

There are analogs to conditionalization in the full belief framework. For example, Friedman (2000), defends the following norm of inquiry: when a question has been closed, don't reopen it. This is a close analog to conditionalization's controversial consequence: that possibilities with credence 0 cannot recover positive probability. Other diachronic norms are weaker: for example, some forms of epistemic conservatism say that if an agent rationally believes a proposition at an earlier time, then it remains rational for her to continue believing it at later times, as long as she doesn't receive any new, disconfirming evidence.

In this paper, I discuss a number of general diachronic norms that cross-cut whether we treat belief states with the full belief framework or the partial belief framework, and also cross-cut whether we treat the overriding diachronic norm as conditionalization, or whether we accept alternative diachronic norms on credences (e.g. Jeffrey conditionalization).³ Here is one candidate:

Diachronic evidentialism: An rational agent will only change her epistemic state by updating on new evidence. A consequence of diachronic evidentialism is that when we speak of 'an agent's evidence', we are speaking of all of the evidence the agent has received, not merely the evidence that is accessible to the agent, or internal to the time-slice. It would therefore be question-begging to motivate time-slice internalism by claiming that agents sometimes lose evidence and are only responsible for believing what it supported by 'their evidence at a time'. By the lights of the diachronic evidentialist, 'evidence at a time' is just all the evidence the agent has received by that time.

Note that this is, on its face, a fairly strong norm. One needn't endorse this strong a norm in order to accept that there are diachronic constraints on rationality. But we'll start with the strong claim and see what can be said in favor of it, before considering weakenings.

1.2. *The rejection of diachronic rationality*

Arguments against the existence of diachronic epistemic norms appear in Talbott (1991), McGrath (2007), Christensen (2000), Williamson (2000), Meacham (2010), Moss (2014), and Hedden (2015). There are a variety of motivations for a time-slice-first epistemology. Some, e.g. Williamson, simply find diachronic constraints like Diachronic Evidentialism implausible. For others, time-slice internalism follows from a more general principle – in particular, some form of epistemic internalism. Here, for example, is Meacham (2010):

In Bayesian contexts, many people have appealed to implicitly internalist intuitions in order to support judgments about certain kinds of cases. But diachronic constraints on belief like conditionalization are in tension with internalism. Such constraints use the subject's beliefs at other times to place restrictions on what her current beliefs can be. But it seems that a subject's beliefs at other times are external to her current state. (87)⁴

There are a number of different forms of epistemic internalism. The two varieties that are perhaps most familiar are *mentalist internalism* and *access internalism*.

Mentalist Internalism: the facts in virtue of which a subject is epistemically rational or irrational supervene the subject's mental states.

Access Internalism: the facts in virtue of which a subject is epistemically rational or irrational supervene on those of the subject's mental states that she has access to, in some epistemically relevant sense.

Neither of these immediately conflicts with diachronic constraints on rationality, at least as stated. After all, it might be that what's rational for an agent to believe at one time supervenes on her mental states at another time, or her mental states at many different times, or currently accessible mental states from other times, or mental states that were at one time or other accessible.⁵

Opponents of diachronic norms often appeal to a form of access-internalism: facts about our past mental states are irrelevant to our current rationality because they are, at least in some circumstances, inaccessible to us.⁶ However, not all opponents of diachronic norms accept access internalism, in the sense of access noted above. So the form of internalism endorsed by the opponent of diachronic norms should be characterized neutrally:

Time-slice Internalism: the facts in virtue of which a subject is epistemically rational or irrational at a particular time *t* supervene on the subject's mental states at *t*.

Here's an example statement of this sort of internalism:

Whether it is rational to retain or abandon a belief at a time is a matter of which of these makes sense in light of your current epistemic perspective, i.e. in light of what you currently have to work with in revising your beliefs. (McGrath (2007), 5)

Time-slice internalism is typically taken to entail that the norms governing epistemic rationality are purely synchronic, and hence Diachronic Rationality is false.

In some instances (e.g. Meacham (2010), Hedden (2015)), the motivations for time-slice internalism draw on an analogy between facts about the past and

facts about the external world. Our access to our past mental states is, at least in principle, limited in just the same way as our access to the external world.

In particular, this form of time-slice internalist emphasizes the analogy between an internally coherent agent who's deceived about the external world and a synchronically coherent agent whose memories are periodically systematically scrambled. Both agents are doing the best they can under strange, externally imposed circumstances. What more could rationality demand?

The proponent of diachronic norms responds that the scrambled agent should instead be understood on analogy to someone who is given a drug that makes him believe contradictions. They are both doing the best they can under strange, externally imposed circumstances – but nevertheless, they are not ideally rational. I'll argue for this claim in greater detail in Section 2.3. First, however, I'll discuss some general objections to any form of time-slice internalism, including those that are access externalist.

1.3. Diachronicity in normative theories

I will ultimately defend a comparatively weak claim: that there are diachronic norms of epistemic rationality. Advocating diachronic epistemic norms does not entail advocating conditionalization.

There are weaker diachronic requirements that could constrain rational belief: for example, that one shouldn't reduce or increase confidence in a proposition (in which her previous credence was rational) unless she receives new evidence or loses evidence. The time-slice internalist endorses a strong claim: that rationality is in no way sensitive to diachronic features of agents.

The framework for epistemic normativity that I use in this paper is epistemic utility theory. Following Berker (2013), I distinguish three components of this form of epistemic consequentialism: a theory of final value, a theory of overall value, and a deontic theory. In epistemic utility theory, the theory of final value is encoded in the epistemic utility function: a function from epistemic outcomes to cardinal utilities. Typically, in epistemic utility theory, epistemic outcomes are characterized in terms of the (gradational) accuracy of a doxastic state. The theory of overall value delivers a ranking of epistemic options, as a function of epistemic utility and other parameters, e.g. probabilities. For example, options may be ranked in terms of their expected epistemic utility. Finally, the deontic theory delivers verdicts about what's epistemically required, permissible, and impermissible. For example, expected utility maximization might require choosing an option with highest expected utility.

All three elements of the theory are normative, not simply the deontic theory. And so diachronic normativity could enter into the framework at any of the three points. There could, in principle, be diachronic norms that affect the theory of overall value (ranking of options) without affecting the deontic theory. (For example, diachronic consistency could be universally supererogatory.) For this

reason, it's useful to frame the discussion in terms of the facts in virtue of which an agent is *more or less* epistemically rational.

The opponent of diachronic norms under discussion opposes diachronic norms entering at any point within the normative theory. It's not obvious that any actual opponents of diachronic norms endorse this strong claim, or instead the weaker claim that diachronic norms don't affect the deontic theory. Absence of an explanation for why diachronic normativity in the theory of overall value shouldn't affect the deontic theory, it seems most likely that actual opponents of diachronic rationality defend the stronger claim: diachronic facts have no effect on an agent's rationality.

2. Problems for time-slice rationality

2.1. Problem #1: *permissibly discarding evidence*

One of the benefits that time-slice internalists claim for their view is that, by rejecting conditionalization, they are able to vindicate the idea that forgetting doesn't render a person irrational. If conditionalization applies, without qualification, over the whole of an agent's life, then any instance of forgetting would be sufficient for irrationality.⁷

The flip side is that time-slice internalism also makes any instance of discarding evidence epistemically permissible. Discarding evidence is, at least *prima facie*, a canonical example of a violation of epistemic norms. The reason that time-slice internalism has this effect is that discarding evidence is a fundamentally diachronic phenomenon. At some time, you receive evidence. At a later time, your attitudes fail to reflect the fact that you've received that evidence.

Discarding evidence

Suppose an agent has strong beliefs about whether capital punishment has a deterrent effect on crime. Then he learns of a study that provides evidence against his view. So he should reduce his confidence in his belief. But instead our agent (involuntarily) discards the evidence; he loses any beliefs about the study; it has no enduring effect on his attitudes regarding capital punishment. Now he can go on confidently endorsing his beliefs without worrying about the countervailing evidence.

This is a standard example of irrationality. One might object: an agent like this is epistemically irrational only if he voluntarily discards the evidence. But cognitive biases are not voluntary; so this objection would have the consequence that cognitive biases never result in irrational belief. I take this to be uncontroversially false.

Another possible objection: it's not the discarding of evidence as such that's irrational, but rather the disposition to discard evidence. This disposition is possessed at a time-slice. So the time-slice internalist has the means to explain the irrationality of discarding evidence.

Two replies: first, it seems to me that an agent who is disposed to discard evidence but never manifests this disposition is not epistemically irrational. It's the discarding itself, not the disposition to discard, that is irrational. Second, as Hedden (2015) points out, an agent might discard evidence and afterward lose her disposition to discard evidence. The time-slice internalist has no explanation for the epistemic non-ideality of such an agent.⁸

Discarding evidence is epistemically irrational. Therefore, there are diachronic norms of epistemic rationality. There's not much more to say about this. But to my mind it is a serious challenge to time-slice internalism; perhaps the most serious.

2.2. Problem #2: erratically changing beliefs without new evidence

Some kinds of belief change are plausibly described as deviating from some sort of epistemic ideal, even when no synchronic norms are violated. It might be controversial whether, by virtue of deviating from the ideal, the agent is irrational. Nevertheless, if there are purely diachronic epistemic ideals to deviate from, it follows that there are diachronic epistemic norms.

Consider again an agent whose total belief state is entirely overhauled at regular, and perhaps frequent, intervals (every minute? every second?). At every instant her credences are probabilistically coherent. And they uphold any other synchronic constraints on rational belief: for example, they are appropriately sensitive to chance information, they reflect whatever the epistemically appropriate response is to whatever phenomenological inputs the agent has at that instant, etc. However strong you make the norms of synchronic rationality, our agent obeys all of those norms at each instant.

But her total belief state at one moment is largely different from her total belief state at the next. If you asked her a minute ago where she was from, she'd say Orlando; if you asked her now, she'd say Paris; if you ask her a minute from now, she'll say Guelph. These changes are random.

The time-slice internalist is committed to the claim that such an agent can be *ideally rational*. I think this is false. Whether or not the agent rises to the level of rationality, it is clear that she is epistemically subideal: with respect to epistemic rationality, she is doing worse than someone whose credences are more stable over time.⁹

Objection: If her evidence changes with each belief overhaul, then perhaps it is rational for her to overhaul her beliefs so frequently.

Reply: In order to assess whether her evidence changes with each belief overhaul, we would need to say more about what 'her evidence' is. For example, if you believe her evidence is what she knows¹⁰ – i.e. $E = K$ – then we can stipulate that it too will overhaul, since her beliefs overhaul. It might just be that she irrationally stops believing various propositions that she previously knew.

Similarly for other views wherein what an agent's evidence is depends at least in part on her beliefs.

For a toy example, suppose an agent initially knows the outcomes of 50 independent tosses of a coin that is biased either 3/4 toward heads or 3/4 toward tails. Unfortunately, the coin happens to have fallen about half heads over the 50 tosses, so the information the agent has received happens not to be informative about the bias of the coin. If our agent stops having beliefs about the outcome of most of the tosses where the coin landed heads, she'll come to be confident that the coin is biased toward tails; vice versa if she stops having beliefs about most of the tosses that landed tails. If she alternates between these two states of mind, then she will change from confidence that the coin is biased toward heads to confidence that the coin is biased toward tails, back and forth, as quickly as you like. Because she loses belief, she loses knowledge; because she loses knowledge, on this conception of evidence, she loses evidence.

Imagine having a conversation with someone like this about the coin. I take it the person would seem more than merely confused. She would not seem like a paragon of rationality.

On any conception of evidence where some of an agent's beliefs count as evidence (because they count as knowledge, or for some other reason), it's unclear how the time-slice internalist can make such beliefs rationally mandatory. In losing those beliefs, the agent loses evidence; if she loses evidence, the time slicer says, then it's permissible to change beliefs. The proponent of $E = K$ owes us some story about how beliefs that are knowledge can ever be rationally mandatory even at a time-slice: in other words, a story about why one *should* have evidence.¹¹ Moreover, this toy case shows that endorsing *Uniqueness* – the thesis that anybody of total evidence determines a unique rationally permissible doxastic state – does not prevent wild fluctuations in belief when 'evidence' is interpreted synchronically.

On the other hand, if diachronic evidentialism is correct, then 'an agent's evidence' is *all the evidence the agent has received*. It is not just the evidence that is accessible to her in the moment, or internal to a time-slice. As for the agent whose beliefs overhaul every second, the proponent of diachronic evidentialism will say: her evidence does not dramatically change. So it's irrational for her beliefs to dramatically change.

2.3. *Epistemic ought implies can*

One might object: so much the worse for any of diachronic norms! 'Ought' implies 'can.' Consider one of the primary motivating cases for time-slice internalism: forgetting.

Forgetting is not irrational; it is just unfortunate. (Williamson (2000), 219).

But forgetting is not merely unfortunate. it's *epistemically* unfortunate. 'Epistemic misfortune' is simply a gentler name for epistemic subideality.

'Ought'-implies-'can' principles are questionable in epistemology. Our friend in his tinfoil hat can't make himself stop overtly believing contradictions. That doesn't mean he's doing what he epistemically ought to do. It is a commonplace in epistemology that a person can be irrational even when she is doing the best she can.¹²

Indeed, even if the epistemic *ought*-implies-*can* argument were successful against ideals like deductive closure, probabilism, or precise credences, it's not clear that it applies to forgetting. A norm against forgetting would prescribe maintaining one's current state (except in response to new evidence). Unlike deductive closure, probabilism, or precise credences, it's not physically or psychologically impossible to be in the recommended belief state. After all, the relevant agent has already been in the recommended belief state.

The relevant 'can' for 'ought'-implies-'can' principle is difficult to provide a semantics for. It's a challenge for the proponent of 'ought'-implies-'can' arguments against e.g. conditionalization to provide truth conditions for the relevant sense of 'can.'

Importantly, the relevant 'can' must not be the 'can' of practical rationality, whereby the relevant acts are under the agent's immediate voluntary control. Beliefs are not subject to immediate voluntary control.¹³ So an inability to immediately control forgetting is consistent with an epistemic obligation not to forget.

Finally: suppose the 'ought'-implies-'can' argument was successful against strong diachronic norms like conditionalization or the more general diachronic evidentialism. This would be no argument against diachronic epistemic requirements. Even if 'not can' implies 'not ought,' still: sometimes 'can' and 'ought.' We are sometimes capable of retaining information over time. When we can avoid forgetting, *ceteris paribus*, epistemic rationality favors doing so. To say that diachronic constraints have exceptions is not to say that there are no diachronic constraints.

One fear we might have about accepting epistemic principles that ordinary agents can't perfectly realize is that we would then have to accept that the norms of rationality are, in some sense, only for ideal agents; they don't apply to any actual agents.

But that's rather like saying that if you're not ideally law abiding – you've already gotten a speeding ticket; there's nothing you can do to change that fact – then traffic laws no longer apply to you. Suppose the traffic laws say:

- (1) Don't get speeding tickets;
- (2) If you get speeding tickets, pay the speeding tickets;
- (3) If you don't pay your speeding tickets, go to your court hearing;
- (4) ...Then this set of legal norms generates different 'levels' of law-abidingness. 'Ideal law-abidingness' amounts to obeying *all* of these (where everything after 1 you satisfy trivially by virtue of satisfying 1). Still, if you *can't* obey all of the laws, you're legally required to obey 2, 3, ...; and

if you *can't* obey 2, then you're legally required to obey 3, etc.. What the traffic laws require of you in particular circumstances is relativized to what you are capable of. Still, though, if you are not capable of satisfying all of the laws, then you are not *ideally* law-abiding.

We can represent the norms of rationality as having a similar structure:

- (1) Be diachronically *and* synchronically coherent.
- (2) If you can't be both, be synchronically coherent.
- (3) ...etc. So, like law-abidingness, we can think of rationality as *relative* – in particular, relative to our cognitive limitations. Ideal rationality is a special case of relative rationality: it is the case where there are no limitations.

2.4. Rationality and epistemic ideality

The aim of this paper was to defend the claim that there are diachronic epistemic norms. Here are two stronger claims:

Rationality = Ideal Rationality In order to be epistemically rational, one must satisfy all epistemic norms, synchronic or diachronic; or even stronger:

Rationality Requires Lifelong Information Retention Epistemic rationality requires never having lost any information through the course of one's life.¹⁴ One could resist these extensions of diachronic rationality. For example, it might be that being immune to information loss would make an agent epistemically better, but that it isn't necessary for rationality. Or it might be that information retention is rationally required over stretches of time, but not an agent's entire life. Perhaps it's required between instances of some psychological event of forgetting, where this might be psychologically distinguished from discarding evidence. Perhaps whenever it's psychologically possible for an agent to retain information, she should. These hypotheses are not remotely as demanding as the strong claims above.

The opponent of diachronic norms defends a universal generalization: she insists that no one is *ever* rationally subideal by virtue of diachronic facts. The proponent of diachronic norms defends an existential: there need be only one instance where, e.g. discarding evidence is rationally subideal.

2.4.1. Diachronic normativity within an epistemic deontic theory

First, there's a case to be made that diachronic norms should affect our epistemic deontic theory: that is, our assessments of whether an agent is rational or irrational.

A sociological observation: formal and informal epistemologists tend to talk about rationality in quite different ways. For many informal epistemologists, the majority of people are basically rational. It is common to think, e.g. that

one is not rationally required to believe all the consequences of one's beliefs (though perhaps doing so would be epistemically better). By contrast, among formal epistemologists, it is more common to use 'irrational' to mean *rationally imperfect*. To be epistemically irrational, in their sense, is to deviate from epistemic ideals.

Formal epistemologists often accept that all of us are irrational. Because of our cognitive limitations – for example, the fact that we can't believe all mathematical truths – actual agents' beliefs are never actually closed under deduction. Deductive closure remains an epistemic ideal. Similarly for probabilism, which entails that we must, e.g. have credence 1 in all necessary truths. Probabilism also entails that we must have infinitely precise credences: that there be a difference between having credence .2 and credence .2000000000000001. But because of our cognitive limitations, on the most plausible theories of mind, actual human agents never have infinitely precise credences. Moreover, canonical forms of epistemic irrationality are also a consequence of cognitive limitations. Irrationality is involuntary. Confirmation bias, hypoxia, and paranoid delusions are involuntary.

So, the proponent of diachronic evidentialism might conclude, because of our cognitive limitations, no actual agents are epistemically ideal. There's no obvious reason to treat forgetting any differently. Why should we classify the misfortune of being a forgetter differently from how we classify the misfortune of confirmation bias? The time-slicer should provide an explanation for why some epistemic misfortunes do not constitute irrationality, while others do.

2.4.2. Diachronic normativity within an epistemic theory of overall value

Second, there's a case to be made that diachronic norms should affect our epistemic theory of overall value: that is, our assessments of an agent's comparative epistemic ideality.

This form of diachronic epistemic norm is less demanding than those that infect our deontic theory. It won't immediately follow from this hypothesis that anyone is irrational by virtue of diachronic facts. One might agree with Williamson that forgetting is rationally permissible – but still hold it to be epistemically subideal.

One might, for example, have a *satisficing* deontic theory of epistemic rationality, according to which rational permissibility doesn't require ideal rationality. Traditionally, this form of deontic theory isn't common among formal epistemologists. Formal epistemologists tend to accept that rationality just is ideal rationality, and so accept that none of us is rational. But satisficing views are commonly presupposed informal epistemology. For example, informal epistemologists typically accept that it's not rationally required that we believe all the consequences of our beliefs, while also accepting that we would be rationally better if we did.

Still, there is common ground between these views. The satisficing and maximizing epistemologists can agree on comparative assessments of epistemic ideality. These comparative assessments form a (possibly partial) ranking of epistemic acts, states, or agents. The satisficer and maximizer endorse different functions from rankings to sets of permissible acts. The maximizer permits all and only maximally ranked acts, while the satisficer appeals to some permissibility threshold. The ranking is normative: it represents the extent to which different acts compare in terms of value. If the satisficer and maximizer agree on the ranking, they agree on some of the normative facts.

For example: perhaps the satisficer and maximizer can agree on the following ranking of epistemic properties in terms of ideality:

- (1) Omniscience
- (2) Conformity to some epistemically privileged diachronic and synchronic constraints¹⁵
- (3) Conformity to the synchronic constraints
- (4) Conformity to the synchronic constraints except where one is susceptible to involuntary biases and delusions

Each property entails the subsequent properties. The omniscient agent is epistemically ideal; the non-omniscient agent who nevertheless obeys both synchronic and diachronic norms (e.g. probabilism and conditionalization) does worse, epistemically, than the omniscient agent, but better than someone who only obeys synchronic norms, who in turn does better than someone who sometimes disobeys the synchronic norms because of biases and delusions.

Where in this ranking do we draw the line between what's necessary for *rationality* and what isn't? It is uncontroversial that omniscience is epistemically ideal but not necessary for rationality. It is also uncontroversial that some involuntary biases and delusions are sufficient for irrationality. So the question is whether to draw the rationality line between 1 and 2 (with the friend of diachronic norms) or between 2 and 3 (with the foe of diachronic norms).

The time-slicer worries that requiring 2 for rationality slippery-slopes into requiring 1. She perhaps allows that it's epistemically better to conform to diachronic norms, but maintains that it's not necessary for rationality. Otherwise, what's to stop us from requiring agents simply to be more knowledgeable? Meanwhile, the proponent of diachronic norms worries that permitting 3 slippery-slopes into permitting 4. If we treat information loss (forgetting or discarding evidence) as rationally permissible, what's to stop us from permitting other involuntary biases and delusions? Why not allow that it's epistemically rational to fully believe a contradiction, so long as you can't help it? Most – perhaps all – forms of irrationality are involuntary; should we thereby conclude that no one, or virtually no one, is ever irrational?

What's important to note, however, is that friends and foes of diachronic norms can still agree on this ranking – part of a theory of overall epistemic value¹⁶ – while disputing the *deontic* epistemic theory. It's at the deontic stage that we make judgments of rational permissibility or impermissibility. But if the diachronic facts make a difference to the theory of overall value, then there are diachronic epistemic norms that affect the comparative assessment of epistemic acts.

Suppose we all accept that the agent who loses information (by forgetting or discarding evidence) is doing worse, with respect to overall (not final) epistemic value, than the agent whose credences only change by rational update on new evidence. Then for my purposes, it doesn't matter whether we call the information loss irrational or rational. It might be that diachronic norms do not factor into the deontic theory: the determination of whether an agent is on–off rational. Still, though, diachronic facts are relevant to an agent's epistemic status. You might think that rationality is comparative or comes in degrees, or you might accept an on–off view of rationality. But at the level over overall epistemic value – where comparative or degreed rationality comes into play – the ranking above are plausible, and plausibly reflect diachronic norms.

2.5. Subjective and objective epistemic value

I have emphasized that there's a clear sense in which the subject who violates diachronic norms is doing worse, epistemically, than the subject who doesn't. But the time-slice internalist might object: the person who happens to *know* less is also doing worse, epistemically, than a person who knows more. But that doesn't mean that the person who knows less is *irrational*. So, the time-slice internalist might conclude, not all epistemic norms are norms of rationality.

There is a natural way of drawing a distinction between norms of epistemic rationality and other epistemic norms. In the practical realm we sometimes distinguish 'objective' and 'subjective' norms. In epistemology, it's often accepted that objective epistemic value is determined as a function of accuracy: the value of believing (or having high credence in) true propositions and disbelieving (or having low credence in) false proposition. The norms of rationality, by contrast, are subjective in some sense to be spelled out. At minimum, they do not require omniscience. Where do diachronic norms fall on this divide? Which of the epistemic norms are norms of epistemic rationality?

I'll suggest two hypotheses about how to address this question that are friendly to the proponent of diachronic norms.

The first is the less conservative. There is no binary subjective/objective divide. Motivations for this view have received more attention in the semantics of deontic modals (e.g. Kolodny and MacFarlane (2010)), but are generalizable beyond the linguistic and the practical. Diachronic norms are more 'subjective' than the norm of truth, but more 'objective' than some synchronic norms (like

probabilism), which are in turn more 'objective' than other synchronic norms. Defending this hypothesis is outside of the scope of this paper.

Second, and more conservatively:

Schwarz (2012) defended conditionalization with this analogy: suppose we want to build a robot to gather information for us in whatever environment it ends up in. We have the option of programming it to obey diachronic evidentialism. Should we? It seems fairly obvious that we should. Then the robot will not lose information, and so will end up with more information.

One of the ways of cashing this out: the epistemic norms are the constraints that characterize the epistemic states of the ideal information gatherer. The ideal information gatherer is non-omniscient; none of her beliefs is guaranteed to be true except on the basis of evidence.

Epistemic rationality involves having beliefs that approximate the truth as much as possible, given our non-omniscience. On this view, though, there's no reason to think of diachronic norms as somehow external to rationality. Retaining information will, by and large, help you keep your belief state more accurate.

Accuracy-centered epistemology (e.g. James (1896); more recently, in epistemic utility theory, Rosenkrantz (1981), Joyce 1998, 2009, Greaves and Wallace (2006), Leitgeb and Pettigrew (2010)) supports this hypothesis. Objective epistemic norms are encoded in scoring rules, which characterize objective epistemic utility in terms of gradational accuracy. Subjective norms are decision rules, e.g. expected inaccuracy minimization or accuracy dominance avoidance. These norms can constrain not only doxastic states at a time, but the relation between doxastic states at different times. Epistemic rationality is a matter of believing in accordance with epistemic decision rules that tend to promote accuracy. Accuracy-centered epistemic utility theory provides a formal precisification of the hypothesis that the norms of rationality are the norms of the ideal information gatherer.¹⁷

3. Rational information loss

3.1. *Losing information to gain information*

There are complaints against conditionalization that have nothing to do with information loss: for example, that it only allows update when evidence justifies credence 1 in some new proposition (unlike, e.g. Jeffrey conditionalization), and that it doesn't allow believers to lower their credence in any proposition from 1 even in circumstances where no forgetting takes place (e.g. in Arntzenius (2003) Shangri-La example). But neither of these objections extends to diachronic evidentialism; so these considerations simply invite us to find a suitable diachronic replacement for conditionalization.

There's another argument against conditionalization that extends to diachronic evidentialism. Like Arntzenius, Christensen (2000) argues that we're not merely rationally permitted to violate conditionalization, but in fact in some cases conditionalization violations are rationally mandatory. Unlike Arntzenius' argument, Christensen's argument involves rationally mandatory information loss.¹⁸

First, note that it can't be that losing information *necessarily* makes an agent's belief state less accurate. For example: suppose that, by chance, you happen to forget only misleading evidence. Losing information thereby makes your belief state more accurate. But retaining information makes it more likely that your credences will be more accurate, roughly speaking. For example, it increases the expected accuracy of your credences.

Now, conditionalizing on new information is an example of pure information gain. And forgetting and discarding evidence are examples of pure information loss. But what should we say about mixed cases?

We can define an *informational tradeoff* as a case where an agent gains some information at the cost of losing some other information. If taking an informational tradeoff can be rationally permissible, then some strong diachronic epistemic norms are false. For example, conditionalization is false: rational informational tradeoffs would require rational information loss.

Christensen (2000) uses an example with the following structure to argue against the view that there are diachronic epistemic norms:

Doxastic downloader

Suppose you know that someone knows more than you about some topic. You know a few things she doesn't know, but on the whole she's more informed on the topic. Unfortunately, it would be gauche to ask her about the topic. Fortunately, you have the option of using a doxastic downloader to replace your credences on the topic with hers. Is it permissible for you to do so?

Christensen invites us to judge that it is indeed permissible.

It should be clear that this is at best an argument against *some* diachronic norms, not against diachronic rationality in general. But one interesting fact about this case is that if you take the tradeoff, you violate conditionalization – but you also increase the expected accuracy of your credences. So, if epistemic rationality consists in maximizing expected accuracy, then conditionalization can't be a norm of epistemic rationality.

Note that there are two possible objections one could make against conditionalization on the basis of this example.

Objection #1. Taking the tradeoff maximizes expected accuracy. So you're rationally required to violate conditionalization. This shouldn't trouble the proponent of conditionalization. The norms of epistemic rationality govern only epistemic states, not actions like using a credence downloader. If we were rationally required to perform actions that maximize the expected accuracy of our credal states, then we would, for example, be rationally required to perform constant experiments, to read all of Wikipedia, etc.

Objection #2. If you do take the tradeoff, your resulting epistemic state is rational. So it must be permissible to violate conditionalization. This objection is more troubling for the proponent of conditionalization. If this objection is correct, then conditionalization is false. At most, conditionalization holds across periods of time where no worthy informational tradeoffs are available.

There are the two options, then, for the proponent of diachronic norms:

- (1) She can endorse conditionalization and reject the claim that there are epistemically rational informational tradeoffs.
- (2) Alternatively, she can adopt diachronic norms that are more liberal than conditionalization.

There's little more to be said about the first option. Let's explore the second option. But first, we should say a little bit more about what expected accuracy is.

3.2. Epistemic utility theory

Epistemic utility theory formalizes the idea that rational credences are governed by epistemic decision rules that tend to promote epistemic utility. The aim of epistemic utility theory was to use the tools of decision theory, combined with an epistemic version of value, in order to give a foundational justification for various epistemic norms.

The most widely discussed epistemic utility functions are *gradational accuracy* measures. The accuracy of a credence is its nearness to the truth (by some measure). A credence function with maximal accuracy would assign credence 1 in all truths and credence 0 in all falsehoods. In other words, it would be omniscient.

Decision rules are adapted from decision theory, e.g. expected utility maximization and dominance avoidance. Paired with accuracy as the relevant measure of utility, we end up with epistemic decision rules:

Accuracy Dominance Avoidances: adopt a credence function that is not accuracy dominated. A credence function Cr dominates a credence function Cr' iff, at all worlds w in a possibility space \mathcal{W} , $U(Cr) > U(Cr')$.

Maximize Expected Accuracy: adopt a credence function that has maximal expected accuracy, by your own lights.¹⁹ The expected accuracy of a credence function is standardly calculated as the sum of a credence function's accuracy in each world, weighted by the probability of that world. In symbols:

$$EU^{Cr}(Cr') = \sum_{w_i \in \mathcal{W}} Cr(w_i)U(Cr', w_i)$$

With each of these decision rules, various results can be proven. Greaves and Wallace (2006) and Leitgeb and Pettigrew (2010) proved that from an agent's

own perspective, given the choice of all possible update policies, conditionalization uniquely maximizes expected accuracy. Similarly, Briggs (2013) and Robert and Williams (2006) argue that conditionalization provides an update strategy that is uniquely accuracy non-dominated. So, one might conclude hastily, in order to be an ideal information gatherer, your credences should update by conditionalization. Hereafter, we'll focus on expected accuracy maximization.

But the doxastic downloader case is intuitively a case where an agent is in a position to expect that some other credences than her own will maximize expected accuracy from her point of view. The agent rationally expects an increase in accuracy only if she updates by accepting an informational tradeoff, thereby violating conditionalization. Does that example conflict with the results of epistemic utility theory?

3.3. Assessing informational tradeoffs

No. There's no conflict between the idea that there could be rational informational tradeoffs (violating conditionalization) and the epistemic utility theoretic result that conditionalization is the only update policy that maximizes expected utility.

The reason: the epistemic utility theoretic results apply in cases where the relevant space of epistemic acts includes only credence functions, specified *de re*. But it's a feature of informational tradeoffs that you do not know, in advance, what credences you will adopt as a result of taking the tradeoff. (If you did, then you could update on that information directly, which would amount to pure information gain.) Indeed, on common assumptions,²⁰ it cannot be the case for any particular credence function that you can rationally assign it higher expected accuracy than your own credence function. But if you have the option of adopting some member of a set of possible credence functions – adopting a credence function specified *de dicto* as whichever satisfies some constraint – then that option can maximize expected accuracy from your perspective.

Let's consider a particular case of an informational tradeoff, specifying some of the details from the doxastic downloader case.

Coin toss

Suppose a particular coin is either fair or biased (with a $\frac{3}{4}$ heads bias), and it will land either heads or tails. You are uncertain about both matters. Now, you and your colleague start with the same priors:

$$\begin{aligned} w_{FH}: \text{fair, heads} & \quad Cr_0(w_{FH}) = \frac{1}{4} \\ w_{FT}: \text{fair, tails} & \quad Cr_0(w_{FT}) = \frac{1}{4} \\ w_{BH}: \text{biased, heads} & \quad Cr_0(w_{BH}) = \frac{3}{8} \\ w_{BT}: \text{biased, tails} & \quad Cr_0(w_{BT}) = \frac{1}{8} \end{aligned}$$

Then you learn whether the coin lands heads or tails. Your colleague learns whether the coin is fair or biased. Both of you conditionalize on your respective evidence. You are not permitted to know the answers to both questions.

Suppose you learn that the coin lands heads. You have a credence downloader that will allow you to perform the informational tradeoff. Is it epistemically rational for you to give up your knowledge in order to gain your colleague's?

Applying expected utility maximization isn't straightforward. Since we don't know what your colleague has learned, we don't know which credence function to assess. So it's not obvious how we can even determine the expected accuracy of your colleague's credence function.

Here is my suggestion: we can introduce a new kind of epistemic action. Call it *learning the answer to a question*. Learning the answer to a question involves taking an epistemic option when you're not in a position to know what credence function it will result in your adopting.

This kind of epistemic tool isn't just for science fictional cases where you are offered informational tradeoffs. We can do other things with our new epistemic acts. For example, they can be useful in decisions over whether it would be more informative to perform one experiment or another, in circumstances where it is impossible, or at least costly, to perform both. However, these cases don't involve informational tradeoffs in the relevant sense: they don't involve partial information loss.

For a question Q (i.e. a partition over the set of epistemically possible worlds), let Cr_Q be Cr_0 conditionalized on whatever the true answer to Q is (that is, whichever proposition in Q is true at the world of assessment).

In our example, we can call whatever credence function your colleague has after learning whether the coin is biased or fair $Cr_{Q_{BF}}$. Note that ' $Cr_{Q_{BF}}$ ' is a description: it picks out different credence functions in different worlds. Ex hypothesi, your colleague updates on B in B -worlds and on F in F -worlds.

Now, with a concrete example in hand, and a new tool (the epistemic act of learning the answer to a question), we can ask: should you take the tradeoff? We need to explain how to calculate the expected accuracy of $Cr_{Q_{BF}}$ from your point of view:

- (1) Calculate the accuracy of Cr_B at B -worlds and Cr_F at F -worlds.
- (2) Sum the values, weighted by their probability according to Cr_H . In symbols:

$$EU^{Cr_H}(Cr_{Q_{BF}}) = \sum_{w_i \in \mathcal{W}} Cr_H(w_i)U(Cr_{Q_{BF}}, w_i)$$

In this case, with plausible assumptions about the accuracy function U , taking the tradeoff maximizes expected accuracy. Retaining your current credences does not.²¹

This isn't surprising. Knowing that the coin landed heads isn't particularly informative about whether the coin is fair or biased, since it would be unsurprising either way. On the other hand, if you had instead learned that the coin had landed tails, then it would maximize expected accuracy to reject the tradeoff. After all, knowing that the coin landed tails gives you fairly strong evidence in support of the coin's being fair.

So, we have a concrete case where taking an informational tradeoff maximizes expected accuracy, and a decision rule for assessing informational tradeoffs.

3.4. Discussion

Again, the defender of diachronic norms has two options for responding to an objection like this. If she endorses conditionalization, then she must reject the claim that it's rational to accept informational tradeoffs. This might involve rejecting the idea that we should perform those epistemic acts that maximize expected accuracy, or it might involve rejecting the idea that taking an informational tradeoff is appropriately understood as an epistemic act.

On the other hand, if we allow informational tradeoffs as epistemic options, then accepting tradeoffs can lead to maximizing expected accuracy. And if we accept that this is rational, then we should reject conditionalization. The defender of diachronic rational norms should replace conditionalization with a more liberal diachronic rule.

These two options provide us with different pictures of what an ideally rational agent's credences will look like over time. On the conditionalization picture, the ideal rational agent's stock of information will strictly increase. But if we allow for violations of conditionalization in informational tradeoffs, then the ideally rational agent will in some circumstances take epistemic risks. These risks have two salient features that distinguish them from obeying conditionalization. First, they involve sure loss of information; second, they may lead to decreases in the agent's expected accuracy (from the perspective of her post-tradeoff credences).

Here is a candidate liberal diachronic norm (which is a variant on diachronic evidentialism):

Liberal norm: An ideally rational agent's credences change only in order to maximize their expected accuracy. Note that for cases of pure information gain, conditionalization will still hold. Furthermore, rational tradeoffs arguably only occur in sci-fi cases.²² In ordinary cases, the verdicts of the liberal norm will coincide with the verdicts of the traditional, strict norm:

Strict norm (diachronic evidentialism): An ideally rational agent's credences only change in response to new evidence.

4. Conclusion

I've argued that there is a conflict between diachronic norms of epistemic rationality and a form of epistemic internalism. I've also argued that diachronically coherent agents are epistemically better. We should think of epistemic rationality as providing constraints that allow us to be more informed about our environment, whatever our environment happens to be like.

The diachronic norms I've advocated are at a middle ground between epistemic internalism and externalism: they are sensitive to facts that are external to the time-slice, but not necessarily external to the person. Contrast this sort of view with process reliabilism, which is concerned with whether some belief-forming process *actually* conduces toward the truth. Whether it does will depend on contingent facts about the agent's environment. A norm like expected accuracy maximization is concerned with whether an update method is *likely* to conduce toward the truth, by the believer's own lights.

If we take the option of maintaining conditionalization, we are also given at a middle ground between epistemic conservatism and evidentialism. Like conservatism, conditionalization permits us to continuing to believe a proposition if we already believe it (with certainty). In fact, conditionalization requires it. But unlike conservatism, conditionalization doesn't permit continuing to believe a proposition after the evidence for it has been forgotten. Conditionalization requires remembering the evidence as well. Conditionalization doesn't permit violations of diachronic evidentialism. Hence, what we're required to believe is always determined by what our evidence supports.

Notes

1. That is, an rational agent's credences conform to the following axioms: where \mathcal{W} is the set of all worlds under consideration (which I suppose throughout this paper to be finite):

(1) *Nonnegativity*: for all propositions $A \subseteq \mathcal{W}$, $Cr(A) \geq 0$

(2) *Normalization*: $Cr(\mathcal{W}) = 1$

(3) *Finite additivity*: if A and B are disjoint, then $Cr(A \vee B) = Cr(A) + Cr(B)$

2. $Cr(A | B)$ is usually defined as follows:

$$Cr(A | B) = \frac{Cr(A, B)}{Cr(B)}$$

3. Indeed, while I will defend conditionalization against the time-slice internalist's objections, it seems to me obvious that *de se* information of the sort discussed in Arntzenius (2003)'s Shangri-La case are successful objections to conditionalization. Note, however, that these cases are not counterexamples to the norm I call 'diachronic evidentialism.'
4. Note that while Meacham argues that there is a conflict between conditionalization and internalism, and provides a synchronic alternative to

conditionalization, he is (at least in his (2010) not committed to the denial of traditional diachronic conditionalization.

5. These interpretations depend on a non-tensed reading of the principles, which I take to be charitable (since otherwise their synchronic commitments would be undefended).
6. Christensen's objection to diachronic norms, which I discuss in Section 3, doesn't require appeal to either mentalist or access internalism. Williamson and Moss both explicitly reject access internalism. Note: if evidence need not be accessible, then it's no longer clear what motivates restricting the evidence an agent's belief states should respect to evidence that is internal to a time-slice. Moss (p.c.) suggests that this restriction is not motivated by any more general principle and is normatively primitive.
7. Meacham (forthcoming) distinguishes *sequential* and *interval* updating rules. A sequential updating rule tells an agent how to adjust her doxastic attitudes whenever she receives a piece of new information. An interval updating rule tells an agent how her credences should harmonize over arbitrary intervals, given the cumulative information that she receives during an interval. Conditionalization has both sequential and updating interpretations. Both are properly diachronic. Interval conditionalization is a stronger norm than sequential conditionalization, since the latter doesn't rule out changes in belief that aren't responses to new evidence: for example, forgetting.
8. Hedden (2015) accepts that such an agent is rational.
9. Of course, it's entirely appropriate that an agent's beliefs should continuously change *a little* all the time: she should update on new information about, e.g. the passage of time, new events that she encounters, etc. But in the example I'm concerned with, a much greater proportion of her beliefs change, and not simply because she's exposed to new evidence.
10. See Williamson (2000) for the canonical defense of this identity.
11. On conceptions of evidence where one's own beliefs aren't evidence, there need be no epistemic norms on evidence possession – only on doxastic responses to evidence.
12. In a 2012 AAP talk (no manuscript currently exists), Wolfgang Schwarz argued, similarly, that the motivation for rejecting diachronic norms derives from the idea that they cannot be action-guiding, and this turns on an illicit conflation of the practical with the epistemic.
13. Agents can take actions to induce beliefs, e.g. gathering evidence, or take actions to slowly indoctrinate themselves over time. But there is an important sense in which one cannot believe a proposition merely by trying.
14. In Bayesian terms, this would amount to obeying conditionalization with respect to every past time-slice. This is a simplification: again, *de se* information and its effects on *de dicto* information make clear that lifelong conditionalization is not epistemically ideal. I will (temporarily) speak as though conforming to lifelong conditionalization is epistemically better than not doing so for ease of exposition (and because it's not obvious what the best update rule for *de se* information is).
15. I don't specify which diachronic constraints in the interest of generality (but at the possible expense of clarity).
16. It's consistent with this hypothesis that we treat *final* epistemic value as, e.g. true or comparative gradational accuracy.
17. As Titelbaum (2006) pointed out, the Greaves and Wallace (2006) defense of conditionalization isn't properly understood as diachronic. Rather, it provides

a justification for *planning or intending* to update by conditionalization in light of future evidence. But imagine we are programming our robot to be an ideal information gatherer. We have the option of programming it to plan to update by conditionalization and the option of programming it actually to update by conditionalization. We will choose the latter. Programming our robot to update by conditionalization is the best means of gathering information, from the (third personal) perspective of us, the theorists.

18. Fuller discussion of Arntzenius' counterexample to conditionalization would require framework for *de se* belief update, which would take the present discussion too far afield.
19. Carr (1997) argues against the conception of expected accuracy used by epistemic utility theorists. For the purposes of addressing this objection to diachronic rationality, though, I will take the appeal to accuracy at face value.
20. Namely, that epistemic utility functions must be *proper* in the sense that they yield the result that any coherent credence function maximizes expected accuracy by its own lights.
21. See Joyce (2009) and Leitgeb and Pettigrew (2010) for plausible constraints on epistemic utility functions.
22. One might make the case that clutter avoidance is a more psychologically realistic version of an informational tradeoff; see Harman (1986).

Disclosure statement

No potential conflict of interest was reported by the authors.

References

- Arntzenius, Frank. 2003. "Some Problems for Conditionalization and Reflection." *The Journal of Philosophy* 100 (7): 356–370.
- Berker, Selim. 2013. "The Rejection of Epistemic Consequentialism." *Philosophical Issues* 23 (1): 363–387.
- Briggs, Rachael. (manuscript) "An Accuracy-dominance Argument for Conditionalization."
- Carr, Jennifer. (manuscript) "How to Expect When You're Expecting."
- Christensen, David. 2000. "Diachronic Coherence versus Epistemic Impartiality." *Philosophical Review* 109 (3): 349–371.
- Friedman, Jane. (manuscript) "Inquiry and Belief."
- Greaves, Hilary, and David Wallace. 2006. "Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility." *Mind* 115 (459): 607–632.
- Harman, Gilbert. 1986. *Change in View: Principles of Reasoning*, Cambridge, MA: MIT Press.
- Hedden, Brian. 2013. "Options and Diachronic Tragedy." *Philosophy and Phenomenological Research* 87 (1): 423–451.
- Hedden, Brian. 2015. "Time-slice Rationality." *Mind* 124 (494): 449–491.
- James, William. 1896. "The Sentiment of Rationality." In *The Will to Believe and Other Essays in Popular Philosophy*, 63–110. Norwood, MA: Synthese.
- Joyce, James. 1998. "A Nonpragmatic Vindication of Probabilism." *Philosophy of Science* 65 (4): 575–603.
- Joyce, James, Franz, Huber and Christoph, Schmidt-Petri, eds. 2009. "Accuracy and coherence: Prospects for an alethic epistemology of partial belief". Vol. 342, In *Degrees of Belief*, New York: Synthese, Springer, 263–297.

- Kolodny, Niko , and John MacFarlane, . 2010. "Ifs and Oughts." *The Journal of Philosophy* 107 (3): 115–143.
- Leitgeb, Hannes, and Richard Pettigrew. 2010. "An Objective Justification of Bayesianism: The Consequences of Minimizing Inaccuracy." *Philosophy of Science* 77 (2): 236–272.
- McGrath, Matthew. 2007. "Memory and Epistemic Conservatism." *Synthese* 157 (1): 1–24.
- Meacham, Chris. Forthcoming. "Understanding Conditionalization." Vol. I.
- Meacham, Christopher J.G., Tamar Szabo, Gendler, John, Hawthorne, eds. 2010. "Unravelling the tangled web: Continuity, internalism, non-uniqueness and self-locating beliefs". Vol. 3, In *Oxford Studies in Epistemology*, 86–125. Oxford: Oxford University Press.
- Moss, Sarah. 2014. "Credal Dilemmas." *Noûs* 48 (3): 665–683.
- Rosenkrantz, R. 1981. *Foundations and Applications of Inductive Probability*, Atascadero, CA: Ridgeview Press.
- Schwarz, Wolfgang. 2012. *Worms first! in defence of diachronic rationality*, Wollongong: Unpublished AAP talk
- Talbott, W. J. 1991. "Two Principles of Bayesian Epistemology." *Philosophical Studies* 62 (2): 135–150.
- Titelbaum, Michael. 2006. "Comments on greaves & wallace". In *Pacific APA*,
- Robert, J. and G. Williams, . (manuscript) "A Nonpragmatic Dominance Argument for Conditionalization."
- Williamson, Timothy. 2000. *Knowledge and its Limits*, Oxford: Oxford University Press.