CANADIAN JOURNAL OF PHILOSOPHY, 2016
VOL. 45, NO. 5–6, 670–691
http://dx.doi.org/10.1080/00455091.2015.1124000

Routledge
Taylor & Francis Group

Continuing on

Michael G. Titelbaum

Department of Philosophy, University of Wisconsin-Madison, Madison, WI, USA.

ABSTRACT
What goes wrong, from a rational point of view, when an agent's beliefs change while her evidence remains constant? I canvass a number of answers to this question suggested by recent literature, then identify some desiderata I would like any potential answer to meet. Finally, I suggest that the rational problem results from the undermining of reasoning processes (and possibly other epistemic processes) that are necessarily extended in time.

ARTICLE HISTORY Received 18 November 2015; Accepted 19 November 2015

KEYWORDS Reasoning; belief; consistency; rationality; diachronic

Consider this case from Titelbaum (2013, 154):

> *Baseball:* The A's are playing the Giants tonight, and in the course of a broader conversation A's announcers Ray and Ken turn to the question of who will win the game. They agree that it's a tough matchup to call: the Giants have better pitching, but the A's have a more potent offense; the A's have won most of the matchups in the past, but the A's are weaker this year than usual. All in all, it seems like a reasonable person could go either way. Nevertheless, Ken asks Ray what he thinks, and Ray says 'I'm not certain either way, but I think it'll be the A's.' Ray then goes on to discuss how an A's win might affect the American League penant race, etc.
>
> Five minutes later, Bill comes in and asks Ray who he thinks will win tonight's game. Ray says, 'I'm not certain either way, but I think it'll be the Giants.'

This series of responses seems puzzling, and in need of explanation. Perhaps Ray gained some new relevant information between answering the two questions – he glanced through his A's media guide and saw a crucial statistic he wasn't aware of before. Perhaps Ray remembered a relevant fact about the matchup that he hadn't thought about for a long time and wasn't taking into account when he provided his initial answer. Perhaps Ray's responses don't really reveal his beliefs, and there's some pragmatic reason why he would give a different response to Bill than to Ken. Or perhaps, Ray simply changed his mind about the game between responding to the two prompts.

CONTACT Michael G. Titelbaum titelbaum@gmail.com

© 2016 Canadian Journal of Philosophy

https://doi.org/10.1080/00455091.2015.1124000 Published online by Cambridge University Press

Presumably there are questions we could ask Ray to test these explanations. Suppose we ask him those questions and it turns out that none of them is an accurate description of his experience between giving his two answers. We concoct some more explanations and ask him about those, but they aren't correct either. In the end, Ray admits that he just believed one thing at one time and another thing at another; in fact, he wasn't even aware that his beliefs on the matter had changed until Bill asked the later question.[1]

In that case it seems to me that there is something wrong, from a point of view of epistemic rationality, with Ray's sequence of beliefs. I hope the reader shares that intuition with me. If not, here's a piece of initial evidence for it: Observing Ray's behavior, why are we inclined to ask him the kinds of questions I mentioned above? Why do we look for an explanation of his change in position? I submit that, in Davidsonian fashion, we are attempting to *rationalize* Ray's pattern of responses. If Ray just shifted from one view to another – without any new or remembered information, or even a conscious mind-change in between – then we would level a charge of irrationality. Imputations of irrationality are a last resort in the game of interpretation, so our questions seek other viable options.

Going forward, I will take it as a datum that Ray's pattern of opinions is rationally problematic. This essay explores accounts of *why* it's problematic. I will begin by describing some of the approaches currently available, and explaining my disappointment with each. I will then lay out a list of desiderata I would like a theory of rational doxastic diachronic consistency to meet. Finally, I will tentatively propose a theory that might fit the bill.

1. Available approaches

Let's focus a bit more specifically on the questions I'll be asking. Any argument for a particular rational diachronic constraint on doxastic attitudes can be broken down into two steps: first, one establishes the existence of rational pressure for an agent's attitudes at different times to line up in *some* way; second, one argues for a specific way in which those attitudes ought to line up. The second of these steps is often the easier one. For instance, in Bayesian epistemology there are a number of good arguments that *if* rationality requires an agent's credences to line up over time, *then* updating by conditionalization is the rational way for them to align (at least for cases in which the agent's later evidence is a superset of her earlier). These arguments include Teller (1973) (the "Diachronic Dutch Book"), Greaves and Wallace (2006), Brown (1976), and my own Titelbaum (2013). Yet none of these establishes a rational demand for an agent's attitudes to line up over time in the first place.[2]

I'm interested in that first step, so I will simplify the proceedings by mostly working with cases (like Baseball) in which an agent's evidence remains constant

https://doi.org/10.1080/00455091.2015.1124000 Published online by Cambridge University Press

over time. I take it that in many such cases the constancy of the agent's evidence makes it rational for her to keep her beliefs constant as well. If we can explain this rational pressure to keep one's attitudes fixed in cases of constant evidence, I think that will go a long way toward completing the first step, after which we can carry out the project of determining exactly how an agent's attitudes should change or remain constant in the face of changing evidence. So I will focus on cases in which diachronic *consistency* requires attitudinal *constancy*.

I will be asking two questions about such cases:

Architecture Question: Why is it good (from a point of view of rationality) that our cognitive mechanisms tend to keep our beliefs constant when we neither gain nor lose evidence?

Belief Question: Why is it bad (from a point of view of rationality) when our beliefs don't remain constant in the face of constant evidence?

I will say more about the interaction between these questions. But roughly speaking, my answer to the Architecture Question will be that it's rationally good that we have a particular cognitive architecture because that architecture prevents us from suffering the rational badness that would occur if our beliefs were inconstant in individual cases. So my answer to the Belief Question will provide most of the philosophical substance needed to answer the Architecture Question.

Here are some available accounts that address one or both of the two questions. I note at the outset that these accounts are not all mutually exclusive, nor are they all mutually exclusive with the proposal I will ultimately offer. Thus it's not strictly necessary to think of them all as *rivals*.

(1) *Primitivism*. There exist primitive rational requirements. Requirements of diachronic belief consistency are among those. Such requirements can be discovered by considering cases, but they admit of no further explanation.[3]

On the primitivist account, our judgment about Ray in the Baseball case is a datum, which may help us discover theoretical rationality's diachronic demands. But past that point, nothing more can be said about why it's rationally desirable for Ray to have the same belief at both times. Primitivism answers the Belief Question by saying, 'It just is.'

I find primitivism unsatisfying; it leaves something interesting and important about rationality unexplained. Perhaps we'll have to settle for this approach in the end, but it's worth trying to see if a more rewarding answer to the Belief Question can be found.

One way to explain what primitivism leaves unexplained would be to provide a theory of reasons – specifically, a theory on which each agent has reasons to assign the same beliefs across times. Such reasons are tidily provided by our second account.

https://doi.org/10.1080/00455091.2015.1124000 Published online by Cambridge University Press

(2) *Time-slicing from uniqueness.* An agent's evidence at a given time rationally mandates a specific attitude toward any proposition; constant evidence over time therefore requires constant attitudes. Apparently diachronic rational requirements are ultimately grounded in purely synchronic requirements.

According to the Uniqueness Thesis (Feldman 2007; White 2005), for any body of evidence and any proposition there is a particular attitude that any agent with that body of total evidence is rationally required to take toward that proposition. If the Uniqueness Thesis is true, then any agent who has the same total evidence at two distinct times yet takes different attitudes toward some proposition violates the requirements of rationality at least once. We get the result that a fully rational agent keeps her doxastic attitudes constant across times when her evidence remains fixed.

This account yields the result that something is rationally wrong with Ray's attitudes in Baseball. Assuming Ray initially judged his evidence correctly, it also provides him with a reason to adopt the same attitude toward an A's win at the later time as he did at the earlier. Yet that reason isn't a reason to adopt the same attitude *as such*. Instead, his evidence provides him reason to adopt a particular attitude concerning the outcome of the ballgame at the later time. If he judged that same evidence (and therefore the same reasons) correctly at the earlier time, he will now have reason at the later time to adopt the same attitude as he did earlier. But the fact that he adopted that particular attitude earlier makes no difference to his reasons later on. While there is apparently a rational diachronic constancy norm here, it ultimately boils down to synchronic rational pressures supplied by Ray's evidential reasons.[4]

Thus, the account of diachronic constancy developed from the Uniqueness Thesis is a *time-slice view*. Kelly (forthcoming) defines 'current time-slice' theories of rational requirements so that on any such theory, facts about what an agent is rationally permitted to believe at a given time supervene on the non-historical facts obtaining at that time.[5] Relevant non-historical facts might include an agent's current attitudes, whether her faculties are currently functioning properly, what entities she can currently perceive, etc. Historical facts include attitudes the agent assigned in the past, how she came to have particular beliefs, her faculties' historical track record of reliability, etc.[6]

One need not endorse the Uniqueness Thesis to maintain a time-slice view.[7] But if a time-slicer wants to explain the rational pressure on agents like Ray to keep his attitudes fixed when his evidence remains constant, a Uniqueness-based account nicely fits the bill.

We might criticize this account by suggesting that it doesn't explain all the diachronic rational pressures it needs to. So far I've discussed cases in which an agent correctly forms the attitude rationally mandated by her evidence at the earlier time – but what about agents who assign an earlier attitude at odds with their evidence? One might suggest that even if an earlier attitude

https://doi.org/10.1080/00455091.2015.1124000 Published online by Cambridge University Press

was a mistaken response to the evidence, there is still some rational pressure for the agent to maintain it going forward. If our only explanation of the rational pressure for diachronic constancy is that an agent who initially assigns the evidentially correct attitude should also line up her attitudes correctly with the evidence later on, then the Uniqueness time-slicer seems unable to explain the rational diachronic pressure on agents who make initial errors.

Of course, the time-slicer may respond that on her view no such pressure exists.[8] So I will pursue a different line of complaint: I think the Uniqueness Thesis is false. In fact, I think it's rather dramatically false: Uniqueness conjoins two claims, each of which strikes me as incorrect. First, Uniqueness asserts an evidentialism about rationality, on which the rational requirements on an agent supervene on her current total evidence. I think this evidentialism is false, for reasons supplied in Titelbaum (2010) and then further defended in Titelbaum and Kopec (forthcoming).[9] For present purposes, it suffices to say that assessing large bodies of evidence is a complex task that often involves balancing multiple considerations. For example, scientists may trade off competing epistemic features of hypotheses such as simplicity, predictive strength, and explanatory power. The conclusions a scientist draws from data depend not only on the content of her evidence but also on her weighting of these various features, and different weightings are rationally permissible.

Second, the Uniqueness Thesis asserts than in any given situation exactly one doxastic attitude is rationally permissible to adopt toward each proposition.[10] Yet I'm interested in the possibility that no matter what kinds of facts rational permissions supervene upon – present or past, evidential or otherwise – there are situations in which multiple attitudes toward a given proposition are rationally available to a particular agent. And if there are such situations, agents have the ability to perform a very special epistemic maneuver. Faced with an epistemic situation that could permissibly be resolved in multiple ways, confronted by competing considerations that must somehow be weighed, we have the power to *make up our minds*, taking some of those considerations to be compelling and forming the beliefs they support as a result.[11]

I have tried to present Baseball as one such case – a case in which Ray initially has to make up his mind between two potential conclusions, each of which is genuinely rationally open to him.[12] If the Uniqueness Thesis is false in the ways I have described, there must be at least some such cases, and I will suppose in what follows that Baseball is one of them. I'll admit that I find making up one's mind a fairly mysterious type of doxastic activity. But my concern here is: If an agent makes up her mind among genuinely rationally available options at an initial time, is there any rational pressure for her to maintain the attitude initially selected?[13]

https://doi.org/10.1080/00455091.2015.1124000 Published online by Cambridge University Press

The following accounts suggest that there is:

(3) *Identity/self-governance.* Self-governance and/or maintaining one's identity requires constancy of belief given constancy of evidence. This provides at least a *prima facie* reason to maintain rationally permissible beliefs.

This is really an umbrella category for multiple specific accounts that could be offered. I'm imagining those accounts running parallel to accounts that have been offered in the action theory literature of rational constancy of permitted intention. One could develop an account parallel to that of Korsgaard (2009) on which belief constancy would be a constitutive norm of rational agency, the violation of which would threaten the integrity of our epistemic identity.[14] Or one could develop an account like that of Bratman (2012), on which self-governance would require belief constancy and thus provide agents who aim to be self-governing with a reason to maintain their beliefs.

While these accounts in action theory still need many details filled in (even before we get to their correlates for belief!), I find them attractive and don't think the proposal I'll make necessarily rules them out. Nevertheless, I have a couple of concerns about any account in this category. First, these accounts are *global* rather than *local*; they indict individual episodes of inconstancy on the grounds that those episodes contribute to a broader problem. Too much inconstancy, and the agent either will no longer be self-governing or will simply cease to be. I think, however, that when doxastic inconsistency is rationally deleterious, this is because of features of that instance of inconsistency taken alone, and not because of its contribution to any broader effects. Here's an analogy: Suppose you criticize me for dropping the plastic holder from a six-pack of soda on the ground at a park. I ask you what's wrong with what I did. You might offer the global response that too much plastic left in the park will make that park unusable and will generally harm the environment. Or you might say that *this particular* six-ring of plastic could get picked up by a seagull and choke it to death. The latter response identifies a local problem with my action irrespective of how it might contribute to a larger pattern.

I prefer local answers to the Belief Question, local explanations of the rational problem with diachronic inconsistency. It seems to me that there is something wrong with Ray's set of responses *considered in its own right*, regardless of whether a pattern of such responses would have troubling long-term effects.[15] This kind of explanation also avoids a problem with global approaches: When a behavior pattern is desirable because it allows an agent to maintain a particular property over the long term (or maintain her status as *that particular agent*), one-off deviations from the pattern may not threaten maintenance of the property. This might happen because the particular deviation is somehow causally isolated from the larger effect. (I happen to know that if I litter right now, you will indignantly pick up whatever I have dropped.) Yet, I imagine it's open for

https://doi.org/10.1080/00455091.2015.1124000 Published online by Cambridge University Press

identity/self-governance theorists to say that in such cases diachronic inconstancy is not rationally problematic. While I happen to find that response unpersuasive, there's a trickier case in the offing: It might be that while no particular deviation is causally isolated from maintenance of the property, the property can withstand a few deviations here and there. (A park is still usable with only one or two pieces of litter.) In that case, the global theorist needs a way to say that each and every deviation is rationally problematic, even when the feared consequences for identity or self-governance are not going to appear. On the other hand, an explanation that accounts for the rational badness of inconstant episodes one at a time – on their own terms, as it were – ensures that each and every such episode will exhibit a problem.[16]

Second, if identity or self-governance accounts of rational belief are patterned after their practical brethren, they will suggest the existence of a kind of reason that I don't think exists. It's difficult to discriminate intuitively cases in which *no* reason for a course of action exists from cases in which such a reason exists but is very *small* (Schroeder 2007, Ch. 5). Nevertheless, consider Bratman's case in which he has to choose between two routes (Highway 280 and Highway 101) to reach San Francisco and each route is equally good. Suppose that far before reaching the relevant junction, Bratman makes up his mind to take 280. Now he is approaching the junction at which the decisive turn must be made. As many authors have emphasized, there are often reasons not to reconsider a settled intention.[17] But suppose Bratman re-opens the question anyway, and is considering once more which way to go. It would be awfully strange for Bratman to treat his past intention as generating any kind of consideration in favor of taking 280. Similarly, when Bill prompts Ray to think again about who will win the game, it would be odd for Ray to cite the belief he formed while talking to Ken as part of a *reason* to think the A's will win.[18] Yet the identity and self-governance accounts seem to make these earlier attitudes into reasons for later actions, or at least allow those earlier attitudes to generate reasons influencing later actions. These accounts thus posit a kind of later-time reason (even during reconsideration) of which I'm suspicious.

2. Desiderata

As has already begun to emerge, I am curious whether there is an account of the rational desirability of diachronic consistency that displays particular features. I will now lay out those features, and (where this hasn't already been addressed above) try to explain why I am interested in an account that displays them.

- *Extensional adequacy.* I am interested in an account that matches certain settled opinions about particular cases. For example, I want to maintain the result that there is a rational problem with Ray's changing beliefs in

https://doi.org/10.1080/00455091.2015.1124000 Published online by Cambridge University Press

the Baseball case. Must the account detail a rational problem in *every* case in which an agent's evidence remains constant but her doxastic attitudes don't? As I will discuss below, it's not clear to me that there actually is a rational problem in every such case. So to the extent there is a distinction between belief-inconstancy cases that are rationally problematic and those that aren't, we want the account to track the contours of this extensional difference. It's also possible that in some rationally problematic belief inconstancy cases there is more than one rational problem present. Or it may be that different types of rational problem are present in different belief-inconstancy cases – there may be no single, unified account covering them all. I would like if possible to find a unified account that indicates the same kind of problem in all cases in which belief inconstancy is rationally problematic, but I am not committed to that being the only kind of rational problem present.

- *Explanatoriness.* I am also interested in an *explanation* of why there's a rational problem in diachronic inconsistency cases when such cases are problematic. Further, if there's a distinction between rationally problematic and rationally unproblematic cases of diachronic inconsistency, it would be nice for the account to explain that distinction.

- *Concerns only theoretical rationality.* I am interested to find an account that works purely in terms of theoretical rationality, without invoking concerns of practical rationality. If you like, I want to show that diachronically inconsistent beliefs are rationally flawed *as beliefs* – as a particular kind of representational mental state. In addition to their purely representational role, beliefs may also play a role in rationalizing actions and a causal role in bringing about certain consequences for agents. But I am interested to see if a diachronic consistency account can be developed without adverting to these additional roles.

  This does *not* rule out the account's attending to what I'll call the 'pragmatics' of belief management. For instance, we'll look later at an account of belief consistency centered on the limited computational and storage capacities of finite human minds. This account concerns the demands of *theoretical* rationality on a believer with limited representational abilities. It concerns the pragmatics of belief management for such an agent, but not the practical rationality of possessing certain kinds of beliefs.[19]

- *Local not global.* As discussed above, I'm curious to see whether we can find an explanation for rational flaws in instances of diachronic inconsistency that appeals only to the features of each instance as opposed to the properties of a general pattern into which the instance fits (or fails to fit).[20]

https://doi.org/10.1080/00455091.2015.1124000 Published online by Cambridge University Press

- *Compatible with a Strong Permissivism.* Earlier I indicated that I'm a 'permissivist' (to use Roger White's coinage) – I deny the Uniqueness Thesis. Moreover, I'm what I'll call a strong permissivist, in the sense that I deny *both* conjuncts of the Uniqueness Thesis. So I want an account of diachronic consistency that applies even when the agent's initial doxastic attitude was one of a number that were rationally available to her at the earlier time.

- *Genuinely diachronic.* Given this strong permissivism, I will not be interested in time-slice accounts of diachronic rationality. I think there are cases in which all the purely synchronic features of a situation underdetermine what an agent is rationally required to believe, yet there are rationally significant relations between what the agent believes now and what she made up her mind to believe in the past. Put another way, I think there are rational evaluations of agents' series of beliefs that genuinely require looking at conditions across multiple times. So I will not be opting for a time-slice account.

- *Consistent with mind-changing.* As I just noted, the strong permissivism I'm exploring here sometimes allows an agent to make up her mind among multiple options that are genuinely rationally open to her. Intuitively, if my theory of rationality is going to allow an agent to make her own doxastic bed in this fashion, it ought to allow her to re-make it later. That is, once we allow an agent to *make up* her mind at a given time, it seems we ought to allow her to *change* her mind later on.[21] I find changing one's mind no less mysterious than making it up to begin with, but the general idea here is that if making up one's mind involves, say, determining how to weigh up various considerations in a manner not entirely driven by one's evidence, then it ought to be equally rationally permissible to *re*consider one's weighting later on in a manner not entirely driven by evidence (or a change therein). So I want an account that lends some rational permissibility to mind-changing.

I'm interested in a plausible account of diachronic rational consistency that has all of the features listed above. On the other hand, I will not demand that an account display certain features that many authors simply assume are required.

- *Need not be prescriptive.* While this use of terminology may be a bit idiosyncratic, I use 'normative' as an umbrella term for an entire category of notions contrasted with the descriptive. Thus for me the normative includes both evaluations (e.g. assessments of goodness) and prescriptions (e.g. 'ought' statements). I would be perfectly happy with an *evaluative* account that explained why there is a rational flaw in sequences of doxastic attitudes that are not consistent over time even if it did not entail that an agent *ought* at the later time keep her attitudes consistent with what she had assigned

https://doi.org/10.1080/00455091.2015.1124000 Published online by Cambridge University Press

earlier. I also will not demand that the account yield rational *requirements* of any kind.

- *Later reasons not provided.* In a similar vein, I do not think an account of diachronic rational consistency needs to provide the agent at a later time any *reason* to keep her attitudes in line with what she assigned earlier, especially not a reason that bears weight should she come to explicitly reconsider that earlier assignment. This was part of my point in discussing the Bratman highways example.

As I said, these last two desiderata contravene what many authors assume is the job of a story about theoretical rational diachronic consistency. For instance, Sarah Moss writes,

> The norms articulated in an epistemology classroom govern deliberating agents.... It may be true that people often chug along without deliberating, responding to any indeterminate claim as they did before, without reconsidering [the attitude] they are acting on. It may even be true that people cannot survive without acting in this way. But this does not challenge norms that tell agents what they should do when they deliberate. To compare: it may be true that people often fall asleep and hence fail to consider or assess any reasons at all, and it may even be true that people cannot survive without sleeping. But this fact about human nature does not challenge ordinary norms governing lucid agents. (2015, 184)

I don't know whether consideration of non-deliberators should influence the norms we apply to deliberating agents. I want to challenge the passage's first sentence, and suggest that there are norms worth articulating in the epistemology classroom that do something other than govern deliberating agents.[22] For instance, a theory of rationality might criticize an agent who fails to deliberate at all on important occasions (perhaps she takes a nap instead). This is a genuinely normative negative evaluation of the agent or of her cognitive processes, yet its content bars it from governing agents who are (already) deliberating.

A similar point could be made about the rational assessment of memory loss. I don't know whether it's ever rationally problematic for an agent to forget something. But if norms of rationality must provide the agent with reasons in deliberation, we can generate an impressively quick argument against the possibility of genuinely diachronic rational memory norms. Suppose an agent forms the belief that p at t_1 and then forgets that she did so by t_2. She cannot at t_2 have any reason to believe that p related to her t_1 belief, because she does not remember having assigned that earlier belief.[23] So there's the argument: Norms of rationality must provide reasons in deliberation, the forgetful agent can't have any reasons while deliberating at the later time to retain her earlier belief, so there are no norms of rationality related to memory loss.

This brief bit of reasoning does not seem to me to settle the subtle matter of whether there is ever anything rationally bad about forgetting. But that's in part because I reject the premise tying all rational norms to reasons in deliberation. On the other hand, Luca Ferrero accepts an argument fairly similar to this one.

https://doi.org/10.1080/00455091.2015.1124000 Published online by Cambridge University Press

Assessing the possibility that there might be an irreducibly diachronic constraint of structural rationality against forgetfulness, he writes

> A systematic failure to retain judgment-sensitive attitudes – like beliefs and intentions – is a failure in securing the necessary background for the proper functioning of the rational psychology of temporally extended agents like us. But it does not seem to me that, for each particular judgment-sensitive attitiude that one might have at any particular time t_1, one is under a rational constraint to preserve it at a later time t_2. (2012, 153)

The 'But' here reflects a crucial step in the argument. Ferrero grounds his denial that there exists *any* diachronic structural constraint related to memory loss in a denial that the agent is under a rational constraint to preserve the attitude in question at t_2. This move succeeds only given the assumption that if there is an extant diachronic constraint in place, it must issue a rational prescription applying directly to the agent at t_2.

One might grant that the assumptions I'm questioning here are indeed assumptions, yet still maintain that they flow naturally out of an intuitive understanding of the metanormative landscape. After all, how can we have a diachronic norm of rationality relating two times without its generating prescriptions for the agent at the later time, or at least *reasons* for the agent at that time? I'll try to demonstrate the possibility of a normative account that sheds these assumptions by providing one later on. But for the time being, let me try another analogy about what goes on at the park:

> Suppose that after picking up my litter, I lay out cones to mark the goals for a soccer game. As the game is being played, it's important that those cones stay where they are. (You'd be annoyed, for instance, if a young child picked one up and started moving it around.) But if we stop the game to have a discussion about moving the cones (perhaps I chose a hilly spot, or the sun is in one team's eyes), the fact that they're already in a particular location provides no reason to keep them there. If we decide in the end to move the cones, the child may be confused about why we were angry at him for doing so earlier. But that's because he doesn't understand the normative structure of the situation: Playing soccer is a downstream activity that depends on the cones' maintaining a constant location. In the course of that activity, the fixity is important for reasons that aren't dependent on the cones' particular coordinates. But if we halt that activity and turn to assess the coordinates themselves, the fixity no longer has any normative significance.

It seems to me that in this case it is genuinely bad if the cones move while we are playing soccer. Yet if, say, at half-time we deliberately reconsider their location, the fact that they have been in one place so far provides no reason not to move them. (Setting aside the slight effort it would take to pick them up and set them down.) So we have a genuinely normative evaluation dependent on a cross-temporal relation that nevertheless provides no reason to maintain that relation under deliberate reconsideration. I will exploit precisely this kind of pattern in my account of diachronic attitudinal consistency.[24]

https://doi.org/10.1080/00455091.2015.1124000 Published online by Cambridge University Press

3. A positive account

I will now try to develop an account that explains a rational flaw that occurs in many cases of diachronic inconsistency. I offer this account tentatively, and honestly don't know if I believe all the particulars myself. But even if it fails, I hope it will indicate that there is logical space for a kind of normative account that satisfies the desiderata laid out in the previous section.

The first thing to note about Ray's succession of beliefs in Baseball is that it's weird. Typical human cognitive architecture is such that once we adopt a particular attitude toward a proposition, we generally keep it – especially over a short period of time. In fact, I think it's an interesting question about the metaphysics of mind whether a mental state must have some nontrivial duration in order to even *count* as a belief. That will depend on whether beliefs are functional states, computational states, etc. and on whether such states could be realized only momentarily. But setting that question aside, humans are certainly built to have beliefs that persist. And it's possible that our negative intuitive reaction to Baseball is simply a reaction to Ray's behaving in a manner people typically don't.[25]

Yet I think there's more to it than that. As the Architecture Question indicates, I think that it's good – good from a point of view of theoretical rationality – that we are built such that our beliefs typically remain constant once we adopt them. Or at least, it would be bad if we didn't have that feature.

At this point, someone will inevitably bring up computational limitations. Because of our limited cognitive capacity, we can't recalculate all our beliefs from our evidence (or whatever else) every time we need one of them. So it's useful to be able to settle a matter and then retain that settlement for use at later times.

While I think that's true, one might object that this is only an argument for *having* beliefs about propositions one has considered before, not for keeping those beliefs *constant*. Imagine a creature who, once she adopts a doxastic attitude toward a proposition, always has an attitude toward that proposition at future times. But imagine that this creature's attitudes toward propositions slide around over time: if she has degrees of belief, their numerical values drift; if she has binary beliefs, some of them might flip from belief to suspension to disbelief or *vice versa* at random. Since this creature always has attitudes toward propositions at the ready, she need not calculate a new stance on a proposition should it rearise. But this creature certainly doesn't display the belief constancy we're interested in.

Let's be clear what exactly this objection comes to. I am trying to point out rational benefits of our tendency to keep beliefs constant. The objection proposes that the same benefits could be achieved by a mental architecture lacking such constancy. But that's fine – showing that some other architecture also achieves those benefits doesn't deny that our architecture does too, nor does it undermine the claim that it's a good thing for us that we have the architecture

https://doi.org/10.1080/00455091.2015.1124000 Published online by Cambridge University Press

we do. An analogy: For the sake of navigating our environment, it's good for us that we can see. The ability to see isn't any less beneficial in navigating our environment just because it would also be possible to navigate using echolocation.

Nevertheless, I do think there's another aspect of belief constancy that's crucial for theoretical rationality, and this aspect really depends on the *constancy*. Consider first the synchronic requirements of theoretical rationality. Rationality requires an agent's beliefs to be consistent, and requires her to believe (at least some of) the consequences of what she believes. A rational agent attempts to make her beliefs consistent, and attempts to believe (at least some of) the consequences of what she believes. The rational agent need not think of her efforts to meet these constraints as such. For instance, after Ray forms his belief that the A's will win tonight's game, he goes on to consider the consequences of an A's win for the playoff race. This is an exercise in forming beliefs that follow from what he believes, but Ray probably doesn't think of it that way. Instead, he simply thinks about what follows from the proposition that the A's will win (or really, what follows from *the A's winning*). Similarly, Ray might at some point consider how the A's manager will arrange his lineup to face the Giants' pitching. Once he sorts that matter out, Ray may wonder how his new conclusions mesh with his prediction of an A's win. This is an exercise in maintaining belief consistency.

While the norms here are synchronic, an agent puts herself in a position to satisfy them by engaging in various processes – in particular, *reasoning processes*. Since reasoning is a causal process, it extends over time.[26] And this creates a need for belief stability. Consider an agent who believes p and wants to determine some consequences of that belief. She begins reasoning from p to various conclusions that follow. As her train of reasoning continues, she considers the consequences of those conclusions, and the consequences of those consequences. She may no longer be mentally attending to p at all. But now suppose that as her reasoning worked through these further consequences, the agent's attitude toward p swung around so that she believed $\sim p$. (I realize this doesn't happen to normal human agents – the point is to see why it would be bad if it did.) An abrupt change in belief like that would vitiate the agent's reasoning from p; the reasoning process would no longer be one of drawing out consequences of her beliefs. A similar thing goes for checking belief consistency: imagine how difficult it would be to draw a large cluster of beliefs into harmony if they kept changing as you went along. (An analogy to herding cats feels appropriate here.)

There are other temporally-extended epistemic processes for which it's important that our beliefs remain constant. Sometimes a belief prompts us to launch an investigation out in the world. Sir Arthur Eddington was one of the earlier purveyors and defenders of general relativity. He famously organized a 1919 expedition to confirm the theory by measuring the sun's bending of starlight during a solar eclipse. Now imagine that once the expedition had been organized and launched, Eddington's allegiance had suddenly switched to a

https://doi.org/10.1080/00455091.2015.1124000 Published online by Cambridge University Press

theory compatible with any amount of bend. Again we have a situation in which inconstancy of belief would rob an ongoing process (in this case quite a costly one) of its significance for an agent. While I don't want to put too much emphasis on such external investigations – and will continue to make my argument largely in terms of internal reasoning – it's worth considering the significance of belief constancy in making them possible.

So there's the core of the account: reasoning is a crucial rational activity; being causal, it extends over time; instability of belief would vitiate reasoning's efficacy.[27] I will now expand upon this account in the course of explaining how it satisfies the desiderata I listed above (though not in the order in which I presented them).

Hopefully it's clear that this account, if correct, is Explanatory of at least one problem that occurs when beliefs do not remain constant over time. It also Concerns Only Theoretical Rationality – the explanation has to do with the pragmatics of coming to meet theoretical rationality's synchronic requirements. Further, the explanation is Local not Global: If an agent were to start reasoning from her belief that *p*, then cease to have that belief while the reasoning was still ongoing, this would generate a problem *in that specific case* – not because of any contribution it made to a larger phenomenon.

Assessing the account's Extensional Adequacy is a bit more complicated. It gets the Baseball case right: Being a baseball broadcaster, Ray engages in all sorts of reasoning downstream from his opinion about who will win the game, which is then undermined when that opinion changes out from under the reasoning. This explanation of the rational problem with a change in belief is Compatible with Strong Permissivism; the explanation still goes through if Ray's initial opinion about the game was not rationally required of him. (In fact, it goes through even if Ray's initial opinion was rationally *forbidden* given his evidence at the initial time.)

But what about other examples of diachronic inconsistency? One might object that all I've done is argue for the rational constancy of *some* of our beliefs *while* we're engaged in reasoning that depends on them. Thus, I haven't provided a full explanation of why diachronic inconsistency is rationally bad in general.

I have two responses to this objection. First, one of my goals is to answer the Architecture Question, to exhibit what's rationally good about a cognitive architecture that keeps beliefs intact over time. One way of reading the objection is that I haven't explained why that sort of cognitive architecture is better than an architecture that only keeps beliefs constant while we're engaged in reasoning dependent on them. Now it's a bit difficult to imagine a cognitive architecture that would achieve this trick without going *any* farther – in determining which beliefs to keep constant, such an architecture would have to track exactly what's involved in ongoing chains of reasoning, which investigations have been put into the field, and which beliefs are and are not in consistency relations with

https://doi.org/10.1080/00455091.2015.1124000 Published online by Cambridge University Press

other ongoing beliefs. But more to the point, this is the echolocation objection again. The fact that another kind of cognitive architecture could achieve the feature being considered here doesn't change the fact that that feature is a good feature of the architecture *we* have.

Second response: As an answer to the Belief Question, my account may be getting the extension of the phenomenon correct. Typical human mental faculties certainly don't keep all beliefs intact for all time. Intuitively, some of those losses seem problematic, some less so, and some perhaps not at all. I find quick, repeated belief flip-flops to be much more objectionable than changes over an extended period.[28] It also seems worse to flip on central issues than on trivial matters or matters on the cognitive periphery. The latter asymmetry might be explained by notions like centrality to epistemic identity. But both of these differences could also be explained by the potential of a switch to undermine ongoing reasoning and investigative processes. If I have some trivial belief – say, about the name of the guy who took my order at the coffee place this morning – and over the course of a week it switches from 'Stan' to 'Steve,' that's not a horrible problem. This tracks the fact that during that interval I probably don't have many ongoing reasoning processes or investigations involving this belief, nor is it heavily tied into relations of consistency with my other beliefs. On the other hand, when I imagine a case in which an agent's beliefs rapidly flip all over the place while she is engaged in actively deliberating about their consequences, that seems cognitively crippling. A spectrum of cases lies in-between.

The account may also explain another asymmetry in our intuitive assessments of diachronic inconstancy. In Baseball, we negatively evaluate Ray's belief switch in part because it seems to have taken place without his noticing. On the other hand, if he were to explicitly reconsider his opinion about the game's outcome, and come to change his mind about which way the evidence points, that wouldn't seem as rationally problematic. My account of diachronic inconsistency is Consistent with these intuitions about Mind-Changing: reconsidering whether to believe p is a process of explicitly reasoning *about* p, not a process during which p is taken as a premise while another proposition is considered. So changing one's opinion about p via explicit reconsideration does not have the same negative rational fallout as if one's opinion about p had drifted without one's attending to it.

The account therefore explains what I call the Eyes-On Asymmetry concerning belief consistency. *Very* roughly, it's rationally odd for an agent's beliefs to switch when she doesn't have her eye on them, but when a belief is the explicit focus of a reasoning process, change in that belief is much less rationally problematic. The latter category includes not only cases in which an agent consciously changes her mind, but also cases in which she gains new evidence and changes her belief as a result, and cases in which she discovers that her belief conflicts with other beliefs she possesses and modifies her opinions for that reason.[29] Though I'm not sure whether this gets my intuitions about the

https://doi.org/10.1080/00455091.2015.1124000 Published online by Cambridge University Press

rational status of different types of belief change exactly right, something does seem right to me about this Eyes-On Asymmetry.

The asymmetry also makes a great deal of sense in light of an analogy between this account of diachronic consistency and the soccer cones example from earlier. It's bad if a kid moves the cones while you're busy playing soccer, but it's perfectly fine for you to explicitly reconsider their location and move them during half-time. Reasoning from a premise is a downstream activity that relies on the fixity of that premise in your set of beliefs. But if we come to explicitly reconsider the premise itself, its fixity is no longer necessarily desirable.

This analogy can also help us understand the normative structure of the rational verdicts yielded by my account. The account doesn't generate anything like a general rational requirement that an agent who adopts a particular doxastic attitude *ought* to retain that attitude as long as her evidence remains constant. At best, we get that it's rationally harmful when an agent's attitudes shift around in certain kinds of cases. The rational verdict is evaluative, Not Prescriptive. Moreover, Later Reasons are Not Provided. When we reconsider the cones' location at half-time, their current position provides us with no reasons influencing our decision. Similarly, the fixity of an agent's doxastic attitudes is rationally significant when she is engaged in downstream reasoning. But when she explicitly reconsiders her attitude toward a particular proposition, my account of diachronic constancy does not provide her with any reason to keep her old opinions intact.[30]

One might have a different type of Extensional Adequacy concern: One might worry that my account hasn't captured the intuition that there is something problematic *with changing one's beliefs over time* when one's evidence remains constant. At best, I've shown that if an agent's beliefs shift around, and she nevertheless continues processes of reasoning based upon those beliefs, then this *combination* is rationally problematic. But perhaps there's nothing wrong with the belief-shifting; perhaps the agent's mistake lies entirely in carrying on with her reasoning after the shift occurred.[31]

The complaint here can't be that my account reduces to the nearly trivial 'You shouldn't continue a piece of reasoning once you've stopped believing its premises.' Because of the Eyes-On Asymmetry, the cases with which I'm concerned are precisely those in which the agent is unaware that the premises have been dropped. So the rational problem in the rationally problematic cases can't be that the agent has failed to properly enact this prescription. Yet one might still maintain that a negative *evaluation* of the reasoning is appropriate: while there's nothing wrong with the agent's having dropped the premises, we may negatively evaluate her continuing to reason based upon them.

This sounds to me a bit like the kid at the soccer game saying, 'There wasn't anything wrong with my moving the cones! *You guys* shouldn't have kept playing once I moved them!' Be that as it may, in answering the Belief Question I am primarily concerned to explain why something rationally bad happens in

https://doi.org/10.1080/00455091.2015.1124000 Published online by Cambridge University Press

cases of belief inconstancy. I don't feel a need to demonstrate that *the bad thing* was the belief inconstancy itself as opposed to the resultant vitiated reasoning or a combination of the two. Certainly such a demonstration isn't necessary to answer the Architecture Question: As long as something bad happens somewhere in cases of belief inconstancy, it can be a good thing that our architecture helps us avoid such cases. Moreover, I'm not sure my intuitions about cases like Ray's are fine-grained enough to demand that I locate the rational problem in one particular place rather than another.[32]

Final desideratum: Is my account Genuinely Diachronic? This is a bit subtle, for on one level my account is driven by synchronic concerns. An agent engages in reasoning to meet particular synchronic demands: to make her beliefs consistent, to believe what follows from her other beliefs, to square her beliefs with her evidence. The account explains the rational value of diachronic consistency in terms of what's required for such reasoning to be effective. Yet I don't think this kind of grounding in synchronic norms makes the account a time-slice view. The diachonic norms I've discussed certainly don't *reduce* to synchronic norms as they do on a Uniqueness account. And the rational evaluations in question supervene on genuinely diachronic relational facts. While the driving demands behind the account may be synchronic constraints on belief, much of the view's substance derives from the pragmatic problems of an agent who must develop and maintain such representational attitudes through temporally-extended causal processes. So I would characterize the account as genuinely diachronic, but amenable to those primarily concerned with synchronic norms.[33]

Perhaps an agent's ability to keep her beliefs intact is something like an 'executive virtue' – given the rest of her cognitive equipment, this ability helps the agent achieve synchronic features that are significant for theoretical rationality. And thinking along the lines of virtues may illuminate the relevant normative landscape as well. We are willing to recognize a character trait as virtuous or vicious even when we understand that an agent who possesses that trait may be incapable of changing it (or at least incapable any time soon). When an agent's beliefs go missing, or change without her noticing, there may be nothing she can do about that fact. (Or nothing she can do about it *after the fact*.) But that shouldn't bar us from evaluating the episode as rationally unfortunate.[34]

Notes

1. The previous two paragraphs are adapted from my discussion of the Baseball case in Titelbaum (2013).
2. For what it's worth, a similar point could be made about rational constraints on how an agent's various attitudes should relate at a *given* time. Arguments that rational credences satisfy Kolmogorov's probability axioms typically begin by *assuming* that there's some rational pressure for an agent's synchronic credences in different propositions to line up with each other. The assumption is rarely commented upon only because it seems so obviously true.

https://doi.org/10.1080/00455091.2015.1124000 Published online by Cambridge University Press

3. Thanks to Greg Novack for helping me better articulate this position.
4. There is a small logical step here from the synchronic to the diachronic: Just because one is required to have attitude A at time 1 and one is required to have attitude B at time 2, that doesn't necessarily mean that there's a requirement to have *A at 1 and B at 2*. The details will depend on one's deontic logic of rational requirements.
5. This is not Kelly's *definition* of time-slice views, for two reasons: First, Kelly works with facts about justification rather than facts about rational permission. Second, Kelly ultimately defines a current time-slice theory as one on which the normative facts are *grounded* in non-historical facts. Moving from supervenience to grounding helps Kelly deal with cases in which certain non-historical facts are themselves grounded in historical facts. Yet clearly the grounding definition implies the supervenience condition in the text above. Moss (2015), meanwhile, defines time-slice epistemology in terms of two claims: (1) 'what is rationally permissible or obligatory for you at some time is entirely determined by what mental states you are in at that time'; and (2) 'the fundamental facts about rationality are exhausted by these temporally local facts.' I slightly prefer Kelly's definition because it leaves open the possibility that non-historical facts other than facts about an agent's attitudes might affect what is rationally permissible for the agent. (See also Hedden (2015) for another approach to defining time-slice views.)
6. Kelly notes that on some epistemologies the constitution of an agent's current total evidence will count as a historical fact (or as grounded in historical facts). For instance, an epistemology might combine Williamson's (2000) position that evidence is knowledge with a theory of knowledge invoking the etiology of beliefs. Yet most time-slicers assume that evidential facts are non-historical, so I will go along with that assumption here.
7. Hedden (2015) does, while Moss (2015) doesn't.
8. Instead, she might respond that an agent who forms a rationally-incorrect initial belief will *take* it to be correct, and so will think there is some rational pressure to maintain that belief later on because she will continue to take it to be what's (synchronically) required. In other words, the agent who initially errs is actually under no rational requirement to remain diachronically consistent, but our intuitions about her can be explained by the fact that in some sense it's reasonable for her to *think* she's under such a requirement. This line of response strikes me as unpromising, because we can always ask about cases in which the agent does not maintain her initial belief, nor does she remember its content. If there is a rational problem with the agent's adopting a different belief in response to the same evidence in at least some such cases – as it will emerge I think there is, even when the initial belief was an incorrect response to the evidence – the time-slicer's response will fail to account for that problem.
9. For additional arguments against the Uniqueness Thesis, see Kopec (forthcoming), Kopec (2015), Meacham (2014), Schoenfield (2014), and Kelly (2014).
10. Some formulations of Uniqueness replace 'exactly' with 'at most' to allow for doxastic rational dilemmas.
11. Sometimes the expression 'make up my mind' is used as follows: 'I wanted to go to graduate school for a long time, but in the fall of 2002 I finally made up my mind to do so.' I'm not sure whether this use of the expression (to indicate resolution, as it were) is different from the use just indicated in the main text, but if so let me stipulate that this isn't the sort of mind-making I'll be discussing in this essay.

https://doi.org/10.1080/00455091.2015.1124000 Published online by Cambridge University Press

12. Moreover, I've tried to make Baseball *acknowledgedly permissive*: not only are there multiple conflicting beliefs rationally available to Ray when he makes his initial judgment; Ray is also aware of this permissiveness in his own situation. For the significance of acknowledgedly permissive cases, see Ballantyne and Coffman (2012) and Titelbaum and Kopec (forthcoming).
13. The 'selection' talk may be a bit misleading here; I don't want to take a stand on whether making up one's mind need be or can be a volitional activity. Moreover, my anti-Uniqueness stance doesn't require the outcome of making up one's mind to be underdetermined by *causal* factors – there may be a perfectly deterministic story by which I could look at your brain right now and figure out how you're going to make up your mind about any given issue. I merely want to maintain that the *epistemically* relevant factors bearing on a particular instance of belief-formation may rationally underdetermine which directon a given agent goes.
14. The type of epistemic identity maintained in part by belief constancy need not be one of the types of identity considered in the personal identity literature, nor in the metaphysics of identity more generally. To my mind, Hedden's (2015) attacks on the identity approach assimilate it too quickly to these potentially independent discussions.
15. Earlier, I suggested that my answer to the Belief Question has a certain sort of priority over my answer to the Architecture Question. How does this interact with my preference for local vs. global answers to the Belief Question? Answer: Not at all, as far as I can see. The local/global issue is: Given a rationally problematic instance of belief inconstancy, is it rationally problematic only because (and when) it fits into a larger, undesirable pattern? This is different from the Architecture Question about the rational properties of cognitive architectures that causally *generate* constant or inconstant beliefs. One could give either a local or a global answer to the Belief Question and still consider that answer prior to one's answer to the Architecture Question.
16. Compare the traditional problem with rule-utilitarian theories that agents may on occasion be able to break a rule without causing the general harm that's supposed to motivate that rule. We may think that even in such cases breaking the rule is wrong, but the rule-utilitarian struggles to explain why that's the case.
17. See especially Holton (2008) on the reconsideration of intention, and Paul (2015) for investigation of related questions in the belief case.
18. If Ray is uncertain at the later time what would be a rational response to his initial evidence, yet is fairly confident he was thinking rationally at the earlier time, he might take the fact that he settled on the A's at the earlier time as evidence that that belief was supported by his evidence at that time. But then Ray's later body of total relevant evidence isn't the same as his initial body of evidence, because it includes an evidentially significant fact about his earlier attitudes.
19. Here's one way to think about the distinction: Imagine a purely receptive entity (agent?) whom I'll call the 'passive believer.' The passive believer takes in information about the world, is concerned to develop the most accurate representations of that world that it can, yet has no ability to act in the world as a result. If finite, the passive believer will still have pragmatic theoretical rationality concerns about how to best manage its representational resources, but presumably there are no demands of practical rationality on a passive believer.
20. Compare Ferrero's use of 'local' terminology in his (2012) and the fourth desideratum he lists in Section 1.3 (which he in turn attributes to Bratman (2012, Section 1.5) and Bratman (2010, 10–11 and 20–21)). Elsewhere in the action theory literature, a 'local'/'global' distinction is sometimes used to frame

https://doi.org/10.1080/00455091.2015.1124000 Published online by Cambridge University Press

the question of whether a rational requirement that an agent displays certain general properties over time can generate requirements on particular attitudes considered singly. Notice that even if this question is answered in the affirmative, the resulting account of rational requirements on individual attitudes would not be 'local' in my sense.

21. To quote Jeff van Gundy on a May 2nd, 2014 broadcast of the NBA playoffs: 'Man's greatest right is to change his mind.'
22. Compare Nomy Arpaly's distinction between theorizing about rationality that creates a 'rational agent's manual' and theorizing that creates an 'account of rationality.' As she puts it, 'Not everything which is good advice translates into a good descriptive account of the rationality of an action, or vice versa' (2000, 489).
23. Here, I'm availing myself of the distinction between reasons that apply to an agent (so to speak) and reasons that the agent *has*. One might be willing to grant that the agent's t_1 belief makes it the case that there *is* some reason for the agent to believe *p* at t_2, but because the agent is unaware of any such reason it's not a reason that she *has*.
24. While I don't, some people *define* the normative in terms of the presence of reasons. If we want to satisfy such a definition, we will have to find a reason *somewhere* in the vicinity of the soccer example, or else the statement that it would be bad for the cones to move during the game cannot be genuinely normative. Here's a suggestion: We have a reason to tell the young child not to move the cones while we're playing soccer. While this reason couldn't appropriately figure in our half-time deliberations, it nevertheless is a reason somewhere in the normative mix of the situation, and seems to pair nicely with the claim that it'd be bad if the cones were moved. Later on I'll point out where we could make a similar move (if we felt the need to do so) in my account of diachronic consistency.
25. Sarah Moss suggested this to me in conversation.
26. John Broome writes, 'When you acquire a new attitude – for instance you learn something or you make a decision – many of your other attitudes may need to adjust correspondingly, to bring you into conformity with various synchronic requirements of rationality.... That some of our attitudes take time to catch up is a limitation of our human psychology.... Ideally rational beings would instantly update their attitudes when things change.' (2013, 153) Broome therefore sets aside the time lag in reasoning because the rational requirements he is considering are based on ideal agency.

    Given that human minds are realized in physical brains, human reasoning must be a causal process. (One might argue that as the *generation* of one attitude from another, reasoning would have to be a causal process even if minds weren't physically realized. But I don't know how to argue about causality among the non-physical.) And I believe it's a metaphysical truth that causal processes take time. My conception of rational ideality does not involve idealizing away from the physical realization of minds. So even if we grant that considerations of rationality should attend to the conditions of ideal agency, I do not agree with Broome that ideally rational beings could update instantly.
27. While I came up with the basics of this account independently, it bears strong affinities to some of what John Broome says in his (2013, especially Ch. 10). Kelly (forthcoming) also discusses the relevance of reasoning processes to time-slicing. Meanwhile, Abelard Podgorski has in a number of works developed the idea of reasoning as a temporally-extended process into a much broader account of diachronic rationality.

https://doi.org/10.1080/00455091.2015.1124000 Published online by Cambridge University Press

28. Though this is a general tendency, not an ironclad rule. And that may be tied to the fact that some reasoning processes are extremely extended over time. Philosophers are certainly familiar with mulling over a particular argument or piece of reasoning over the course of *years*.
29. One might worry that even these cases are rationally dangerous: Given the requirement to keep one's beliefs consistent, there's always the risk that a given belief change will generate unnoticed inconsistencies with some of the agent's other beliefs. Yet that seems to me just a cognitive fact of life. The key point here is that my account lumps the threat level of mind-changing in with that of the other two reasoning processes listed, and does *not* put it in the same boat as background flips-flops of opinion.
30. So (going back to the potential concerns of note 24), are there reasons involved at *any* level in my account of rational diachronic consistency? If one wanted, one could say that my answer to the Architecture Question provides reasons at the level of Bratmanian 'creature construction': I have explained some reasons why a designer of cognitive creatures like us might want to provide them with a cognitive architecture that keeps beliefs intact.
31. I'm grateful to Sarah Moss, Michael Bratman, and Sergio Tenenbaum for discussion of this concern.
32. Forget about the *target* of the rational problem; one might complain that I still haven't explained what the *nature* of that problem is. What exactly *is* this 'vitiation' that occurs when a reasoning process loses its premises? I haven't answered that question because I think a number of possible answers are available, and I haven't settled on one that I like best. But just because it's interesting, here's one option: We might see reasoning as an attempt not just to generate new beliefs that follow from one's other beliefs, but instead to generate new beliefs *grounded* or *based* in the appropriate way on one's other beliefs. Viewed this way, reasoning is a process of constructing a cognitive state with a particular justificatory structure. If that's right, then the disappearance of the beliefs from which a process of reasoning began makes it impossible for that reasoning to be successful.
33. Going back to the last point of the Extensional Adequacy discussion, one might worry that if the rational flaw in cases of belief inconstancy lies entirely with the continuation of reasoning, this will open the door once more for a time-slicing account. The thought would be that it's a purely synchronic matter whether one is engaged at a given time in processes of reasoning whose premises one believes. Yet shifting the locus of negative evaluation in cases of diachronic inconstancy to reasoning processes seems to me like a bad move for the time-slicer. A reasoning *process* is a temporally extended affair; evaluations of such processes seem intrinsically diachronic to me. Podgorski's work is once more illuminating on this point.
34. I am grateful to Sarah Paul for copious discussion of this essay and for suggestions concerning many of the references. I am also grateful to audiences at the spring 2014 Informal Formal Epistemology Meeting and the fall 2015 conference on Belief, Rationality, and Action over Time, both held at UW-Madison.

References

Arpaly, N. 2000. "On Acting Rationally Against One's Best Judgment." *Ethics* 110: 488–513.
Ballantyne, N., and E. Coffman. 2012. "Conciliationism and Uniqueness." *Australasian Journal of Philosophy* 90: 657–670.

https://doi.org/10.1080/00455091.2015.1124000 Published online by Cambridge University Press

Bratman, M. E. 2010. "Agency, Time, and Sociality." *Proceedings of the American Philosophical Association* 84: 7–26.
Bratman, M. E. 2012. "Time, Rationality, and Self-governance." *Philosophical Issues* 22: 73–88.
Broome, J. 2013. *Rationality through Reasoning*. Oxford: Wiley Blackwell.
Brown, P. M. 1976. "Conditionalization and Expected Utility." *Philosophy of Science* 43: 415–419.
Feldman, R. 2007. "Reasonable Religious Disagreements." In *Philosophers without Gods: Meditations on Atheism and the Secular Life*, edited by L. M. Antony, 194–214. Oxford: Oxford University Press.
Ferrero, L. 2012. "Diachronic Constraints of Practical Rationality." *Philosophical Issues* 22: 144-164.
Greaves, H., and D. Wallace. 2006. "Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility." *Mind* 115: 607–632.
Hedden, B. 2015. "Time-slice Rationality." *Mind* 124: 449–491.
Holton, R. 2008. *Willing, Wanting, Waiting*. Oxford: Oxford University Press.
Kelly, T. 2014. "Evidence may be Permissive". In *Contemporary Debates in Epistemology*, edited by M. Steup, J. Turri, and E. Sosa, 298–311. West Sussex: Wiley-Blackwell.
Kelly, T. (Forthcoming). "Historical versus Current Time Slice Theories of Epistemic Justification". In *Goldman and His Critics*, edited by H. Kornblith and B. McLaughlin. Blackwell.
Kopec, M. 2015. "A Counterexample to the Uniqueness Thesis." *Philosophia* 43: 403–409.
Kopec, M. (Forthcoming). "A Pluralistic account of Epistemic Rationality." Unpublished manuscript.
Korsgaard, C. M. 2009. *Self-constitution: Agency, Identity, and Integrity*. Oxford: Oxford University Press.
Meacham, C. J. G. 2014. "Impermissive Bayesianism." *Erkenntnis* 79: 1185–1217.
Moss, S. 2015. "Time-slice Epistemology and Action Under Indeterminacy". In *Oxford Studies in Epistemology*, edited by T. S. Gendler and J. Hawthorne, Vol. 5, 172–194. Oxford: Oxford University Press.
Paul, S. K. 2015. "Doxastic Self-control." *American Philosophical Quarterly* 52: 145–158.
Schoenfield, M. 2014. "Permission to Believe: Why Permissivism is True and What it Tells us About Irrelevant Influences on Belief." *Noûs* 48: 193–218.
Schroeder, M. 2007. *Slaves of the Passions*. Oxford: Oxford University Press.
Teller, P. 1973. "Conditionalization and Observation." *Synthese* 26: 218–258.
Titelbaum, M. G. 2010. "Not Enough There There: Evidence, Reasons, and Language Independence." *Philosophical Perspectives* 24: 477–528.
Titelbaum, M. G. 2013. *Quitting Certainties: A Bayesian Framework Modeling Degrees of Belief*. Oxford: Oxford University Press.
Titelbaum, M. G., and M. Kopec. (Forthcoming). "Plausible Permissivism." Unpublished manuscript.
White, R. 2005. "Epistemic Permissiveness." *Philosophical Perspectives* 19: 445–459.
Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.

https://doi.org/10.1080/00455091.2015.1124000 Published online by Cambridge University Press