Mental Processes and Synchronicity

Brian Hedden

Abstract

I have advocated a time-slice-centric model of rationality, according to which there are no diachronic requirements of rationality. Podgorski (2015) challenges this picture on the grounds that temporally extended mental processes are epistemically important, rationally evaluable, and governed by diachronic requirements. I argue that the particular cases that Podgorski marshals to make his case and unconvincing, but that his general challenge might motivate countenancing rational requirements on processes like reasoning. However, so long as such diachronic requirements are merely derivative on more fundamental synchronic requirements, so that a pattern of reasoning counts rational or irrational only insofar as it tends to lead one to better satisfy these fundamental synchronic requirements, we can meet Podgorski's challenge without significantly deviating from a time-slice-centric approach to epistemology.

I have advocated Time-Slice Rationality, according to which the relationship between time-slices of a single agent is treated the same, as far as fundamental requirements of rationality are concerned, as the relationship between time-slices of distinct agents (Hedden 2015a, b). On this view, rationality is concerned with how you are at particular times, rather than with how you are over extended periods of time. Moss (2014, 2015) also gives important defenses and applications of this approach. Podgorski (2015) challenges this picture, arguing that temporally extended mental processes, such as belief formation and reasoning, are rationally evaluable and subject to diachronic norms (I shall use the terms 'requirement' and 'norm' interchangeably). He argues that there are distinctive facts about rationality that cannot be accounted for by appeal only to synchronic norms governing an agent's attitudes at a single time. I will argue that Podgorski's main example is not compelling. However, his more general claim, that essentially diachronic epistemic phenomena are subject to diachronic norms, poses a serious challenge to a synchronic, time-slice-centric theory of rationality. I will aim to show how this challenge can be met by allowing for requirements of rationality governing reasoning which are derivative on more fundamental synchronic requirements governing what attitudes you have at a time. Allowing for such derivative requirements of rationality involves deviating from the letter of Time-Slice Rationality, but not its spirit.

1 Mental processes and synchronic norms

The above gloss on Time-Slice Rationality is picturesque. In Hedden 2015a, I gave a more specific statement of the view as the conjunction of two theses, Synchronicity and Impartiality. Synchronicity is the claim that what attitudes you ought to have at a time does not directly depend on what attitudes you have at other times. Your attitudes at other times may indirectly affect what attitudes you ought to have now, for instance by affecting what actions you performed in the past and thereby affecting what evidence you now possess, but the impact of past attitudes on what you ought to believe now is screened off by facts about your present mental states. Impartiality is the claim that your beliefs about what attitudes you have at other times play the same role in determining what attitudes you ought now to have as your beliefs about what attitudes other people have. Podgorski accepts Impartiality, and so it will play no role in this paper.

Podgorski argues that Synchronicity, as stated above, is insufficient to vindicate a purely synchronic, time-slice-centric conception of rationality. He points out that Synchronicity is only a thesis about the attitudes you ought to have at a particular time. It does not entail that the only objects of rational evaluation are attitudes held by an agent at some particular time. As he puts it, vindicating a purely synchronic model of rationality also requires the claim that 'facts about what attitudes agents ought to have at particular times exhaust the demands of epistemic rationality' (p. 5).[1]

The defender of diachronic norms can accept Synchronicity while holding that there are other objects of rational evaluation that fall outside its purview, and that norms governing these other objects are diachronic. In particular, Podgorski holds that there are what he calls 'essentially diachronic' epistemic phenomena. These phenomena chiefly include mental processes of reasoning and belief formation, which are essentially diachronic simply because all processes take time. Podgorski further holds that there are diachronic norms that govern these 'continuous, temporally extended causal patterns of mental states' (p. 7).

Podgorski rightly notes that my main arguments in favor of a synchronic account of what you ought to believe (i.e. for Synchronicity) do not straightforwardly carry over to support the rejection of norms governing reasoning and other processes. First, I argued that internalism about rationality supports Synchronicity. Internalism about rationality is a view about the factors that determine what one ought to believe at a time, namely that those factors are in some sense internal to the agent (in the version of internalism I favor, they are internal in the sense of supervening on the agent's mental states). One can be an internalist in this sense while still accepting diachronic norms

[1]Kagan (1992) introduces the useful terminology of *evaluative focal points*, which are the things which are evaluated by a normative theory. In the case of ethics, we can evaluate acts, or dispositions, or people, for instance, and moral theories differ with respect to which focal point(s) they deal with and regard as central. In the present, epistemological case, Synchronicity is a thesis about just one evaluative focal point, namely what beliefs you ought to have at a time. Podgorski's point is that there are other focal points (for instance, patterns of reasoning) about which Synchronicity is silent.

such as norms governing reasoning, provided that these norms do not help determine what you ought to believe at a time. Second, I argued that facts about personal identity over time or surrogate notions like R-relatedness are evidentially irrelevant, and hence do not affect what an agent ought to believe at a time, as diachronic norms for what you ought to believe would have it. But facts about identity or R-relatedness are not clearly irrelevant to the question of whether an agent has reasoned correctly, for instance, since they matter for whether the tokened premises and conclusion were tokened by time-slices of a single agent (or by time-slices R-related to each other).

So I concede Podgorski's point that Synchronicity does not entail that there are *no* diachronic norms, but only that there are no diachronic norms governing what attitudes an agent ought to have at a particular time, norms which say that the attitudes an agent ought to have at one time are in part a function of what attitudes the agent has at other times. But I think that the defender of a synchronic, time-slice-centric model of rationality can meet his more general challenge and resist the need for any fundamental diachronic norms governing mental processes like reasoning or belief formation. I begin by looking at his main example, which aims to show that some 'important facts about rationality cannot be settled synchronically' (p. 12), before turning to how a defender of Time-Slice Rationality should treat temporally extended mental processes in general.

2 Time lags in responding to evidence

Here is Podgorski's main example. At t_0 I have the following evidence: (i) that my friend Minnie promised today to come to my birthday party, (ii) that Minnie is a pathological liar who delights in making and breaking promises, so usually her promising to come to a birthday party means that she will not attend, (iii) that Minnie is superstitious and so avoids breaking promises made on the thirteenth of each month, and (iv) that today is the thirteenth. Now suppose that, owing to the fact that responding to evidence takes time, 'the earliest time at which I can deliberately fully form a belief regarding Minnie's presence or absence (about which I am, at t_0, agnostic) is t_1' (p. 13). Finally, suppose that, unbeknownst to me, I will at t_1 suddenly forget (iv) that today is the thirteenth.

Podgorski asks, 'What belief [about Minnie's presence or absence] is it rational for me to form at t_1?' According to those like me who reject diachronic requirements, I rationally ought at t_1 to believe that Minnie will not be at the party, for at t_1 my total evidence will include only (i)-(iii) above, and this evidence supports the proposition that Minnie will not attend. Having beliefs at t_1 which are proportioned to my evidence at t_1 requires my believing that Minnie will not attend. But Podgorski argues that this is the wrong result:

> Because what belief I form at t_1 is determined by the process of belief formation that operates before that time, forming the belief that Minnie will not be at the party would require me to, before t_1, ignore the perfectly compelling evidence I have that she will. It is true that at t_1, my epistemic

> state has changed in a way that may require me, going forward, to cease believing that Minnie will come to my party. But *at* t_1, I have not yet had any chance to respond to this sudden epistemic impoverishment, and so my failure to take it into account cannot be rationally impugned. (p. 13-14)

Now, a first thing to note about Podgorski's case is that, while he talks of belief formation, his conclusion has implications for what beliefs I ought to have at a particular time.[2] If the belief I ought to form at t_1 is that Minnie will not attend, then at t_1 I ought to believe that Minnie will not attend. This entails that sometimes, the beliefs that an agent ought to have at a time are not those which are supported by the evidence she has at that time. Instead, sometimes the belief that an agent ought to have at a time t_1 is one which is supported by the evidence she possesses at a slightly earlier time t_0, rather than one which is supported by the evidence she possesses at t_1. So Podgorski's judgment about the Minnie case conflicts even with Synchronicity as I stated it, as a thesis about what attitudes an agent ought to have at a particular time.

Podgorski's verdict on the case is the wrong one, in my view, for it entails that sometimes one ought to be in a belief state which is incoherent. Once I forget what today's date is, presumably I ought to have high credence that today is not the thirteenth (after all, only one day per month is the thirteenth). Given that I remain certain that Minnie promised today to attend the party, I should therefore have high credence that Minnie will break her promise to attend the party. So if Podgorski is correct that I ought to be confident that Minnie will attend, since that is what my t_0 evidence suggested, then at t_1 I will wind up confident that Minnie will break her promise to attend and also confident that Minnie will attend. But this is incoherent.[3]

Here is a case which parellels Podgorski's but which involves evidence gain rather than evidence loss. Suppose that at t_0 my evidence includes the proposition that P and the proposition that if P then Q, but the earliest time at which I can form a belief regarding Q (about which I am, at t_0, agnostic) is t_1. Suppose also that at t_1 I will suddenly gain $\neg$Q as evidence. Applying Podgorski's reasoning to this case, I ought at

[2]It is also worth noting that Podgorsk's worry has an analogue in the practical case. The practical analogue of the epistemic requirement of having beliefs which are proportioned to your evidence is the requirement of making the decision which has highest expected utility relative to your credences and utilities. Just as there is a worry about needing time to respond to changes in evidence in the epistemic case, there is a worry about needing time to respond to changes in your credences and utilities in the practical case. When your credences and utilities change, you might be unable to take this into account and make the decision which has highest expected utility relative to those credences and utilities until after a short time has passed. My response to this worry in the practical case is identical to my response to Podgorski's objection in the epistemic case.

[3]Perhaps Podgorski would respond that once I forgot today's date, I will have no doxastic attitude whatsoever regarding what today's date is. If that is correct, then coming to believe that Minnie will attend will not result in an incoherent belief state. But it seems at least permissible for me to have high credence that today is not the thirteenth once I forget that today is in fact the thirteenth. And this high credence that today is not the thirteenth (and hence that Minnie will break any promises made today) yields an incoherent belief state if I also form the belief that Minnie will attend, which Podgorski think I ought to do.

t_1 to believe that Q, since this is what my t_0 evidence supported. But I also ought at t_1 to believe that $\neg$Q, for by hypothesis my t_1 evidence includes the proposition that $\neg$Q, and a proposition must be believed in order to count as part of my evidence. So in this case as well, if it takes me time to fully respond to evidence, I will wind up with inconsistent beliefs in the interim.[4]

I conclude that Podgorski's verdict on cases of time lags in responding to evidence is incorrect. Your beliefs at a time ought to be proportioned to the evidence you possess at that same time, rather than the evidence you possessed a short time before, on pain of having an incoherent belief state. (Note that it may be, given my cognitive limitations, that were I to actually have the beliefs supported by the evidence I possess at that same time, that these beliefs would nonetheless not be *based* on that evidence. It would be a stroke of luck that I happened to wind up with beliefs proportioned to your evidence. If so, then it may be that believing as Podgorski thinks I ought (i.e. believing at t_1 that Minnie will attend) is better, epistemically speaking, than getting lucky in believing that she will not attend. My claim is only that the ideal case - the one in which I am perfectly rational - is one in which I believe at t_1 that Minnie will not attend, and moreover believe this not as a matter of luck, but rather as e.g. a result of a general disposition to have beliefs which are proportioned to my evidence. To see this, imagine a god with infinite processing speed who believes (i)-(iv) at t_0 but forgets (iv) at t_1 and consider what this god ought to believe about Minnie's presence or absence at t_1. The answer, it seems clear, is that Minnie will not attend.[5])

This, of course, assumes that there are synchronic requirements of rationality which proscribe being in an incoherent belief state. Podgorski could reject this assumption and hold that his case motivates moving to a framework in which there are *only* diachronic requirements which state how your beliefs should evolve over time, without saying anything about what beliefs you should (or should not) have at any particulary time. Such a framework is defended by Lam (2007), but I will not pursue this option here.[6] Alternatively, Podgorski could hold that cases in which you cannot immediately respond to changes in your evidence are cases in which it is impossible for you to satisfy all the requirements of rationality at once. Requirements of rationality conflict in this case, with one (synchronic) requirement saying that one's beliefs at a time ought to be coherent, and another (diachronic) requirement saying that I ought to form at t_1

[4]Podgorski also considers a slight modification of his original case in which I know in advance that at t_1 I will forget (iv) that today is the thirteenth. He asks what belief I should decide at t_0 to form at t_1 and suggests that the defender of a purely synchronic model will have to give the intuitively incorrect answer that I ought to decide to form the belief that Minnie will attend, for that is the belief that is supported by my t_1, post-forgetting evidence. But I can recognize that that is the belief that would be rational for my t_1 self to have without desiring or deciding now to adopt that belief when the time comes, for I know that it will be rational only given an attenuated evidence base. I can recognize that some belief would be rational at some later time while also thinking that it is less likely to be accurate than some other belief, and in such a case I should attempt to cause myself to form the accurate belief rather than the rational one.

[5]Thanks to Caspar Hare for helpful discussion of this point.

[6]Thanks to Daniel Greco for raising this possibility.

the belief which seemed to me at t_0 to be supported by my evidence. In the Minnie case, the former requirement tells me to believe at t_1 that Minnie will not attend, while the latter requirement tells me to believe at t_1 that she won't. The view that sometimes requirements of rationality conflict has been defended by David Christensen (2007), although his alleged cases of conflict center on the role of higher-order evidence - respecting higher-order evidence about how competent you are at evaluating first-order evidence sometimes requires you to have inconsistent beliefs (or incoherent credences). But judging a case to be one in which there is a genuine conflict between requirements of rationality is a last resort. In my view, we should simply reject Podgorski's judgment about the Minnie case (and other cases of time lags) rather than treat it as a case of conflicting requirements of rationality, though I admit to not having any knock-down argument against the latter option.

3 Ought implies can

My view is that at each time, you ought to have the beliefs which are supported by the evidence you possess at that time. This means that your beliefs ought to change instantaneously in response to changes in your evidence. Does this not violate the principle that *ought* implies *can*?

Arguably not. Even if we finite, cognitively limited humans cannot respond instantaneously to changes in our evidence, so that it is psychologically impossible for us to respond instantaneously, this does not mean that it is nomologically or metaphysically impossible to do so. And it may be that the *can* in an epistemic *ought* implies *can* principle should be read as expressing something like nomological or metaphysical possibility, rather than something more narrow like psychological possibility.

It might be thought, however, that responding instantaneously to changes in your evidence is at least nomologically impossible, for causation takes time. But this thought rests on a mistaken conception of what it is to respond to evidence. It is not as though your evidence is something completely separate from your belief state, such that you notice changes in your evidence and then react by changing around your beliefs. Rather, what evidence you possess is linked to what your belief state is. For plausibly, in order for some proposition to count as part of your evidence, you must believe that proposition.[7] This is not to say that everything you believe counts as part of your evidence, but only that everything that is part of your evidence must be something you believe. If you do not believe P, then P is not part of your evidence, even if it is something that could become part of your evidence in the future.

If this is right, then it is trivially true that at least some of your beliefs do respond instantaneously to changes in your evidence, namely your beliefs about propositions that are part of your evidence. If you didn't previously believe P, and P becomes part of your evidence, then at that very instant you come to believe P.

[7]This is true, for example, on Williamson's (1997) E=K view, on which your evidence consists of all and only the propositions that you know.

I hold that in such a case, there should be no time lag between when you come to believe P and when you make the requisite further changes in your beliefs. (More accurately, I hold that at each time, your beliefs ought to be the ones best proportioned to the evidence you possess at that same time, and this entails that, if at all times you have the beliefs you ought to have at that time, then there will be no time lag between a change in your evidence and your coming to have beliefs proportioned to your new evidence.) In the ideal case, your beliefs in certain evidence propositions do not change first, with these changes then causing the rest of your beliefs to eventually change as well. Instead, your belief state should change all at once, without some parts lagging behind others. After all, if some parts of your belief state lag behind changes in other parts of your belief state, then you will wind up with incoherent beliefs in the interim. So insofar as you have a time lag in coming to fully respond to changes in your evidence, you fall short of perfect rationality in the meantime.[8]

Now, it might be that, as finite, cognitively limited agents, we cannot have our belief states change all at once. Perhaps as a result of our limitations, some parts of our belief states must inevitably lag behind changes in other parts. If this is right, then it really is psychologically impossible to always have beliefs which are proportioned to the evidence one possesses at that time. But I would draw the conclusion, not that we sometimes ought not have beliefs which are proportioned to our evidence, but rather that sometimes we cannot help but fall short of doing what rationality demands of us. This should come as no surprise, for we already have a great deal of independent evidence that we fall short of that ideal in systematic, predictable ways.

Moreover, it is appropriate to have a model of rationality which is not sensitive to our contingent cognitive limitations. First, contingent cognitive limitations vary not only between species (e.g., between humans and far smarter aliens whose belief states change all at once rather than piecemeal), but between individuals within a single species. A model of rationality which is sensitive to all these particular limitations would be objectionably disunified. Second, adopting the sort of highly demanding model of rationality that I favor does not commit us to holding that ordinary people are epistemically blameworthy for failing to always have the beliefs which are supported by their evidence. Being psychologically unable to satisfy the requirements of rationality may excuse one from blame without making one exempt from their demands.

The sort of model of rationality that I support would typically be regarded as a heavily idealized one, but I think that it is better described as demanding rather than idealized. The strict requirements of rationality that I endorse are not to be thought of as descriptions of how some mythical ideally rational creatures would conduct themselves, descriptions which have only indirect relevance for how we earthly beings should be. Rather, these requirements of rationality apply directly to all of us, though we can be expected to fall short of them on a regular basis.[9]

[8]Compare Broome 2013, Ch 9.

[9]Compare Broome 2013, 155-6: 'most requirements of rationality are necessary within...the domain of rationality. They apply to you at all worlds where you are a rational being. This means that, if a requirement would apply to you were you a superior sort of rational being such as an angel, it applies

Compare Utilitarianism in ethics. If Utilitarianism is true, then most of us violate the requirements of morality much of the time, and do so predictably. Moreover, just as we may find it 'psychologically impossible' to fully respond immediately to changes in evidence, so we may find it psychologically impossible to reliably meet the demands of Utilitarianism, not only because of our epistemic limitations, but also because we may be hard-wired not to be completely selfless. But this does not mean that Utilitarianism is to be thought of as a cleaned up, idealized model of morality which is just more tractable than the messy true model of morality. Rather, it just means that Utilitarianism is very demanding.[10] By the same token, I take the correct theory of rationality to be a highly demanding one. We regularly fail to satisfy its requirements, but this does not mean that those requirements do not apply to us.

4 Reasoning

So far I have addressed Podgorski's main example involving time lags in responding to evidence and argued that it does not show that you sometimes ought not have the beliefs which are supported by your present total evidence. But what about Podgorski's more general point, that there are essentially diachronic epistemic phenomena such as reasoning, and that these will be governed by diachronic requirements of rationality?

In my view, the defender of a synchronic, time-slice-centric model of rationality should deny the existence of rational requirements governing the reasoning process, except if those requirements are derivative on more fundamental synchronic requirements (more on this below).[11] The reasons are foreshadowed in my discussion of Podgorski's example of time lags. If we were perfectly rational, we wouldn't need to engage in reasoning in order to satisfy the requirements of rationality and have beliefs which are proportioned to our evidence. We reason precisely because we fall short of perfect rationality. Reasoning is a tool we can use to get ourselves to come closer to satisfying the requirements of rationality. In this way, its value is contingent and instrumental - contingent because it stems from our contingent cognitive limitations, and instrumental because reasoning serves as a means to the end of having beliefs proportioned to one's evidence (compare Broome 2013, p. 207, though he does think we need to formulate norms for reasoning).

Importantly, reasoning is not alone in being a tool we have for cognitive self-improvement. There are physical actions such as taking naps and consuming caffeine which we can perform with an eye toward getting our beliefs more in line with our evidence, and there are other mental actions such as brainstorming and imagining which can serve the same purpose.

---

to you as a human being.'

[10] See Colyvan 2013 for discussion of idealisation in normative models and how such idealisation compares to the sorts of idealisations we find in scientific models.

[11] This section reiterates the stance I defend in Hedden 2015, Ch. 10, on which norms for reasoning can be accommodated provided that they are derivative on more fundamental synchronic norms.

Because reasoning is a cognitive self-help tool, I regard its deployment as being governed by the same requirements of instrumental rationality that apply to other sorts of actions. Whether, when, how, and for how long you ought to reason will be determined by the same expected utility considerations that apply to other actions, including those that serve to help you better satisfy the demands of epistemic rationality such as drinking coffee and brainstorming.

One sympathetic to diachronic requirements might not be happy with this stance on reasoning. Even if reasoning in some fallacious way happens to maximize expected utility, that pattern of reasoning is still irrational. I am not wholly convinced by this thought, but I think that the defender of a purely synchronic model of rationality can go some way toward accommodating it. The key is to distinguish between fundamental and derivative requirements of rationality. I claim that the fundamental requirements of epistemic rationality are synchronic requirements governing beliefs held at particular times, in particular the requirement that you have the beliefs which are supported by your evidence at that time. But given this fundamental requirement of rationality, we can also evaluate dispositions, processes, and the like in terms of how conducive they are helping you satisfy this fundamental requirement. Some belief-forming disposition, or some pattern of reasoning, can be evaluated positively to the extent that it can be expected to reliably get you closer to having beliefs which are proportioned to your evidence. There would be little point in putting one's foot down and refusing to use ordinary normative terms like 'ought,' 'rational,' 'requirement,' and 'norm' when making these evaluations. Better to be permissive about what can count as, say, a requirement or norm of rationality, provided we keep track of which requirements or norms are fundamental and which are merely derivative.[12]

In this way, we can allow for a rational requirement that you not reason fallaciously, provided we are clear that this requirement derives whatever normative force it has from a more fundamental requirement that you proportion your beliefs to your evidence.[13]

[12]In distinguishing between fundamental and derivative norms, I have been influenced by Williamson (forthcoming), who uses the distinction to defend a knowledge-first approach to epistemology from the New Evil Demon Problem (Cohen 1984). A brain in a vat who falsely believes she has hands may violate the fundamental norm to have only beliefs which constitute knowledge. But she nonetheless merits some positive evaluation in virtue of satisfying a derivative norm to believe what someone generally disposed to have beliefs which constitute knowledge would believe in that situation. This allows Williamson to vindicate our judgment that there is *something* that the brain in the vat is doing well, but without abandoning a knowledge-centric approach to epistemology. Similarly, an agent might have a belief which constitutes knowledge but nonetheless merit some criticism, if she nonetheless is generally disposed to have lots of beliefs that fall short of knowledge. In the same way, I can appeal to derivative norms to capture the sense in which someone who reasons well but winds up with beliefs which aren't proportioned to her evidence (for instance, due to forgetting one of her premises in the course of the reasoning) has done something well, and also the sense in which someone who reasons fallaciously but winds up through sheer luck with beliefs which are proportioned to her evidence has done something badly. The first agent violates the fundamental norm of having beliefs which are proportioned to her evidence but satisfies a derivative norm of reasoning in ways which tend to lead to beliefs which are proportioned to one's evidence, and the second satisfies the fundamental norm but violates the derivative norm.

[13]Note that because reasoning is an action, the voluntary aspects of reasoning, such as when to start

Insofar as it is irrational to reason fallaciously, this is because it is irrational to have beliefs which are not proportioned to your evidence, and fallacious reasoning can be expected to lead to your having such irrational beliefs.[14]

Admittedly, this is more a promisory note than a rigorous defense of particular norms governing reasoning and mental processing more generally. There are a number of different ways in which one might define a derivative norm for reasoning. One might evaluate a pattern of reasoning in terms of whether it in fact reliably leads to beliefs better proportioned to one's evidence, or whether it reliably leads to better proportioned beliefs in a certain (perhaps contextually determined) range of nearby possible worlds, or whether it can reasonably be expected to be reliable in leading to better proportioned beliefs. And there are many other options besides. I will not take a stand on which of these ways of evaluating reasoning to adopt in devising a derivative norm for reasoning. Indeed, in keeping with an ecumenical approach to norms, it may be best to allow for each of them to give rise to a derivative norm, with different derivative norms being relevant for different evaluative purposes. A given pattern of reasoning can be evaluated along many different dimensions, since it bears a variety of relevant relations to fundamental goal of having beliefs which are proportioned to one's evidence.

The strategy of adopting derivative norms needn't be confined to norms for reasoning, either. Reasoning isn't the only thing other than beliefs held at particular times that we might wish to evaluate. We also want to evaluate people, for instance. Taking the norm of having beliefs which are proportioned to one's evidence as our fundamental epistemic norm, we can define a variety of norms and evaluations applying to people. A person can be said to be rational to the extent that she usually (i.e. at most times) has beliefs which are proportioned to her evidence (yielding a norm that one ought usually to have beliefs proportioned to one's evidence), or to the extent that she usually

reasoning, what to reason about, and when to stop reasoning, will still be governed by ordinary norms of instrumental rationality, such as expected utility theory, and perhaps also governed by derivative norms of epistemic rationality. The particular pattern of reasoning one engages in (for instance, whether the reasoning is fallacious) may not be under one's voluntary control, and if so it may be subject only to derivative norms of epistemic rationality relating the pattern of reasoning to its tendency to lead to beliefs better proportioned to one's evidence.

[14]As Daniel Greco pointed out to me, one might worry that even the requirement of having beliefs which are proportioned to your evidence may be derivative on a more fundamental norm of believing the truth. Insofar as you ought to proportion your beliefs to your evidence, this is only because you ought to believe truths and refrain from believing falsehoods (though as James (1986) notes, the values of believing truths and abstaining from believing falsehoods will sometimes conflict). I am not sure this is right (see Kelly 2003 and Berker 2013 for objections to certain ways of making this guiding thought precise). But if it is, then norms for reasoning may be doubly derivative; they will be derivative on the norm of proportioning your beliefs to your evidence, which in turn is derivative on the norm of believing the truth. However, I concede that there is an alternative picture, one which is less congenial to a time-slice-centric approach, on which the norm of proportioning your beliefs to your evidence and norms for reasoning are both derivative on the norm of believing the truth - reasoning well is epistemically good not because it tends to lead to beliefs which are proportioned to one's evidence, but rather because it tends to lead to true beliefs. On this approach, the norm of having beliefs proportioned to your evidence and the norm of reasoning well would be on a par, with neither being derivative on the other.

has a disposition to have beliefs proportioned to her evidence (yielding a norm that one ought usually to be disposed to proportion one's beliefs to one's evidence), etc.[15] We needn't choose between these different ways of evaluating persons. They are all legitimate appraisals, different ones of which may be of particular concern in different contexts. The important point for present purposes is simply that the defender of Time-Slice Rationality has the option of adopting an ecumenical approach on which there is a proliferation of requirements of rationality, provided they are all grounded in synchronic requirements governing what attitudes an agent has at particular times.

5 Conclusion

For us mortals, it takes time to respond to evidence, and we often have to resort of reasoning in order to do so. Because of this, we frequently fall short of satisfying the requirements of rationality, which make no allowances for our contingent cognitive limitations. Before we have had time to reason and to get our beliefs to catch up with changes in our evidence, we will have incoherent belief states, which are paradigmatically irrational. Insofar as we cannot help being in such incoherent belief states, we may be blameless for it, but this does not mean that we ought to have incoherent beliefs, as Podgorski's verdict on his case would entail.

He is right, however, that diachronic processes such as reasoning play an important role in our epistemic lives. However, we should resist positing fundamental diachronic

[15] Adopting derivative norms may also help address another objection raised by Podgorski. He notes that Uniqueness, the thesis that each body of total evidence uniquely fixes which doxastic state it is rational to be in, will not rule out intuitively irrational fluctuations in belief if what your evidence is depends on what you believe. For in that case, irresponsible changes in belief could yield changes in your evidence such that your new beliefs are nonetheless proportioned to your new evidence. For instance, if your evidence consists of all and only the propositions that you know, as Williamson (1997) argues, then by dropping a belief willy-nilly, you will thereby lose a piece of evidence, if that belief was one that constituted knowledge. If your new beliefs are the ones supported by your newly impoverished body of knowledge, then this change in your beliefs is not proscribed by Uniqueness. In my view, we should say in such a case that at no time are your beliefs irrational. However, you, as a temporally extendend agent, are irrational, but only in virtue of manifesting a disposition not to have beliefs which are proportioned to your evidence. After all, on an ordinary spelling-out of the case, it was a miraculous stroke of luck that you happened to wind up with the right set of beliefs upon frivolously letting one belief drop out of your doxastic state. Your beliefs may at all times be rational, but you the person are irrational in virtue of violating a derivative norm, that of being disposed to satisfy the synchronic norm of having beliefs proportioned to your evidence.

I hasten to add, however, that it may also be possible to respond to this challenge about improper manipulation of evidence without appealing to derivative norms. For instance, we might adopt a slight modification of Williamson's conception of evidence, holding not that your evidence consists of all and only the propositions you in fact know, but rather all and only the propositions you are in a position to know. Dropping beliefs willy-nilly may affect what you in fact know, but it does not obviously affect what you are in a position to know. Provided we can understand what it is to be in a position to know something in a synchronic way, this will yield an account of evidence which is immune to Podgorski's worry about irresponsible management of evidence while remaining within a fully synchronic framework.

requirements of rationality governing the reasoning process. First, reasoning is a cognitive self-help tool, and so its deployment and use is governed by garden variety norms of instrumental rationality, just like other epistemically important actions like brainstorming, imagining, consuming caffeine, debating with colleagues, and the like. Second, insofar as we want norms which specifically govern the reasoning process, these norms can be derivative on the more fundamental norm of having beliefs which are proportioned to one's evidence. A given pattern of reasoning is then rational or irrational insofar as it can be expected to reliably lead one closer or further from satisfying this fundamental norm. Podgorski is thus correct that giving norms for temporally extended mental processes like reasoning requires a deviation from the letter of Time-Slice Rationality, but it does not require deviating from its spirit. We can have norms for reasoning within an approach which is time-slice-first, but not time-slice-only.

References

Berker, Selim 2013: 'Epistemic Teleology and the Separateness of Propositions'. *Philosophical Review*, 122, pp. 337-93.

Broome, John 2013: *Rationality through Reasoning*. Oxford: Blackwell Press.

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

Cohen, Stewart 1984: 'Justification and Truth'. *Philosophical Studies*, 46, pp. 279–95.

Colyvan, Mark 2013: 'Idealisations in Normative Models'. *Synthese*, 190, pp. 1337–50.

Dutant, Julien and Daniel Dohrn (eds) forthcoming: *The New Evil Demon*. Oxford: Oxford University Press.

Gendler, Tamar Szabó and John Hawthorne (eds) 2015: *Oxford Studies in Epistemology, Vol. 5*. New York: Oxford University Press.

Hedden, Brian 2015a: 'Time-Slice Rationality'. *Mind*, ????

Hedden, Brian 2015b: *Reasons without Persons*. Oxford: Oxford University Press.

James, William 1896: 'The Will to Believe'. *The New World*, 5, pp. 327–47.

Kagan, Shelly 1992: 'The Structure of Normative Ethics'. *Philosophical Perspectives*, 6, pp. 223–42.

Kelly, Thomas 2003: 'Epistemic Rationality as Instrumental Rationality: A Critique'. *Philosophy and Phenomenological Research*, 66, pp. 612-40.

Lam, Barry 2007: *The Dynamic Foundations of Epistemic Rationality*. PhD Thesis, Princeton University.

Moss, Sarah 2014: 'Credal Dilemmas'. *Noûs*, 48, pp. 1-19.

Moss, Sarah 2015: 'Time-Slice Epistemology and Action Under Indeterminacy'. In Gendler and Hawthorne 2015, pp. 172-194.

Podgorski, Abelard 2015: 'A Reply to the Synchronist'. *Mind*, ????

Williamson, Timothy forthcoming: 'Justifications, Excuses, and Illusions.' In Dutant and Dohrn (eds) forthcoming.

Williamson, Timothy 1997: 'Knowledge as Evidence.' *Mind*, 106, pp. 1–25.