

What the Tortoise has to say about diachronic rationality¹

Markos Valaris

Forthcoming in Pacific Philosophical Quarterly.

Even if you believe just what you rationally ought to believe, you may be open to rational criticism if you do so ‘for the wrong reasons’, as we say. Some have thought that this familiar observation supports the idea that there are diachronic norms of epistemic rationality — namely, norms of *good reasoning*. Partly drawing upon Carroll’s story of Achilles and the Tortoise, this paper criticises this line of thought on the grounds that it rests on a mistaken conception of inference.

1. Introduction

Does what you believe at one time affect what you should rationally believe at other times? One tempting reason to think that it does is the phenomenon of *reasoning*. Reasoning can plainly be evaluated as rational or irrational. And reasoning seems to involve changing your attitudes over time. Thus, it is tempting to think that the norms that determine whether a piece of reasoning (an inference) counts as good or bad would also serve as norms of diachronic rationality. This is an idea that has recently been defended by John Broome (2013; 2015), Ulf Hlobil (2015), and Abelard Podgorski (2016) for example.

In what follows I will argue against this view. I do not aim to show that there are *no* such things as diachronic norms of rationality.² My aim in the first instance is the relatively modest one, of showing that one particular line of thought in support of taking the rules of (deductive) inference to be diachronic norms of rationality is unsuccessful. But, as we will see, the issues here go deeper than it might seem.

The debate about the status of the rules of reasoning should lead us to raise some fundamental questions about the *nature* of the activity of reasoning itself. The view that the rules of deductive reasoning are diachronic norms of rationality requires for its plausibility a particular type of view of the activity of reasoning: specifically, a view on which reasoning is a *kind of process*, in which some of one's beliefs cause one either to acquire or to revise some others. Rules of reasoning count as diachronic norms of rationality because they are the norms governing that process.

While this view of reasoning is presupposed by much contemporary discussion, it is not obviously correct. Gilbert Ryle, for one, did not think of inferring as a process. According to him, “‘inferring’ is not used to denote either a slowish or a quickish process. ‘I begun to [infer]’, but I had no time to finish” is not something we can significantly say’ (1949, 302-303). Ryle’s appeal to evidence from ordinary language may not have much power to convince. There are, however, good systematic reasons to think that he was right on this point. Indeed, as I will argue, we can see this by revisiting Lewis Carroll’s famous (if cryptic) argument in ‘What The Tortoise Said To Achilles’ (Carroll 1895). Carroll’s story is typically read as an argument for ‘blind’ reasoning, i.e., reasoning that does not depend on the reasoner’s taking it that her conclusion follows from her premisses. I suggest a different reading. On my reading, the real target of the story is a view of beliefs (or other doxastic states) as isolated *atoms*, with no internal relations to each other. As I will argue, if we abandon this atomistic view of our doxastic states the view of reasoning as a process — and, with it, the view of the norms of reasoning as diachronic norms — loses its attraction.

The paper proceeds as follows. In Section 2 I present more clearly a line of thought that might seem to support the view that norms of good reasoning are diachronic norms of rationality, and explain why I think it fails. Then in Sections 3 and 4 I examine more carefully both why this line of thought is attractive, and why it fails. It is here that the story of Achilles and the Tortoise becomes relevant.

2. An argument for diachronic norms

It is possible for a reasoner to reach a correct conclusion, but in the wrong way. As Broome puts it, ‘incorrect reasoning may bring you to satisfy [a requirement of rationality]’ (2015, 5). You may have very good evidence for something, and yet through inattention, confusion, logical ignorance, or some other misadventure come to believe it on some other grounds which fail to support it at all. Alfred might have very good evidence that Beth committed the murder, and yet come to believe that Beth committed the murder because of a certain pattern in the tea leaves (see Turri (2010, 315–17) for a similar example). Moreover, this kind of mistake can (depending on context) leave you open to a charge of irrationality. This shows that rationality constrains not just *what* you believe, but also *why* you believe it.⁴

I have no dispute with this familiar observation, of course, but I question some recent arguments that are based on it. The type of argument I have in mind goes like this. If rationality constrains not just what you believe but also *why* you believe it, then it must impose constraints on *belief-forming processes*. These constraints, then, are the sought-after diachronic norms of rationality. Ulf Hlobil helpfully makes the underlying line of thought more explicit, as follows (for a very similar argument, see also Broome 2013, 246-247):

Suppose I am doxastically justified in believing at t_0 that *not-A*, that *if B, then A*, and that *if not-B, then not-A*. At time t_1 , I have added the belief that *not-B* to my stock of beliefs, and there is no time between t_0 and t_1 at which I am in a third belief state. Now consider two ways in which this might happen. In the first case, I fallaciously infer *not-B* from *not-A* and *if not-B, then not-A* by affirming the consequent. In the second case, I competently infer *not-B* by Modus Tollens [...]. In the first case I am less epistemically rational than in the second case. My belief states at t_0 and at t_1 are identical. Hence, I must have violated a norm of diachronic rationality in the first case but not in the second. Therefore, there are diachronic norms of rationality. (2015, 38-9)

Since in both cases the agent starts out with the same premiss-beliefs, and ends up with the same conclusion-belief, any difference in the rationality of her performance cannot have anything to do with *where she ends up*, but only with *how she got there*.

Hlobil breaks the argument down in the following way:

- (1) The two cases do not differ with respect to the doxastic states of the agent but only with respect to the transitions between these states.
- (2) The agent violates an epistemic norm only in the first case.
- (3) Therefore, there are epistemic norms that concern the transitions between states.

Regimented in this way the argument is valid; but what should we make of its premisses?

There is room to worry about Hlobil's apparent assumption that fallacious inference is *ipso facto* irrational (or at least rationally defective). If, say, intuitionism turns out to be the correct logic for mathematics then much currently accepted mathematical reasoning will turn out to have been *fallacious*; but it does not seem to follow that it will turn out to have been *irrational*. So believing through fallacious reasoning is not always irrational. For present purposes I am going to set this point aside: after all, it surely is plausible that fallacious inference *can* be irrational, and we can take Hlobil's case to be such. Thus let us grant premiss (2); the problem I want to highlight has to do with premiss (1).

For Hlobil's purposes there has to be a psychological distinction between believing *not-B* through Modus Tollens and believing *not-B* through affirming the consequent. But what exactly is this difference? On one intuitive conception of inference, inference is subject to the 'Taking Condition' — to reason from a set of premisses R to a conclusion p requires *taking it that* R supports p in a relevant sense (Boghossian 2014; Valaris 2014, 2016). If this is so, however, then Hlobil's premiss (1) is false. When the agent reasons by Modus Tollens her doxastic state includes her taking it that *not-B* follows from *not-A* and that *if B, then A*, while when she reasons by affirming the consequent her doxastic state includes her taking it that *not-B* follows from *not-A*

and that *if not-B, then not-A* instead. Thus the agent's doxastic states in the two cases do differ after all.

Could we explain the distinction without appealing to anything like the Taking Condition? It is not easy to see how. The natural alternative to the Taking Condition is a dispositional view of inference: perhaps in the good case the agent is disposed to reason by Modus Tollens, while in the bad case she is disposed to reason by affirming the consequent. But this is not good enough for present purposes. For one thing, as Hlobil (2015, 40-41) realizes, if we allow the agent to have different dispositions in the two cases then this might also count as a difference in their doxastic states, thus undermining premiss (1).

In response Hlobil stipulates that the agent has the same dispositions in the two cases; the cases differ only in which disposition gets manifested. But a deeper problem remains. For, what makes it the case that a given belief manifests one inferential disposition rather than another? As virtually all authors agree, manifesting an inferential disposition is not merely a matter of your being *caused to* believe your conclusion by some pre-existing beliefs (Pollock and Cruz 1999, 35–36; Wedgwood 2006, 661–62; Broome 2013, 225–34). You must, in some sense, come to believe your conclusion *on the strength of* your premisses. Accordingly, if we are to make sense of the claim that our agent manifests one inferential disposition rather than another, we need to make sense of the idea that she believes her conclusion on the strength of one subset of her pre-existing beliefs rather than another. But we have been given no indication of how we are to understand this distinction, without giving an explanatory role to the agent's *taking it that* her conclusion follows from one rather than another set of premisses.

These considerations suffice to show that Hlobil's argument will not do as it stands. But it is worth taking a deeper look at the basic picture of reasoning underlying the idea that rules of reasoning are diachronic norms of rationality, namely, the picture of reasoning as a *process* of belief-revision.

3. Reading Carroll's 'What The Tortoise Said To Achilles'

Our discussion in the last section suggests that the argument for taking rules of inference to be norms of diachronic rationality requires both of the following two claims:

1. For each case of reasoning, there is a determinate rule or pattern that it instantiates
2. Which rule or pattern a case of reasoning instantiates does not depend on the thinker's taking it that any particular connection obtains between her premisses and her conclusion

Claim (1) is needed so that we can, for instance, distinguish between Hlobil's agent's reasoning by Modus Tollens and her reasoning by affirming the consequent. As we saw, however, it is *also* essential to the argument that what pattern the agent's reasoning instantiates not be a matter of the agent's *beliefs* (or other doxastic states) regarding the connection between her premisses and her conclusion. Otherwise, there would be a purely synchronic failing in the cases in question, and we would have no need to appeal to diachronic norms.

As it happens, a large number of philosophers actually *do* accept these claims, for independent reasons. One major reason is a certain influential reading of Lewis Carroll's (1895) 'What the Tortoise Said To Achilles'. According to this reading, the lesson of Carroll's story is that in the most basic case reasoning is 'blind', in the sense that it does not depend on the reasoner's taking it that her conclusion follows from her premisses (see, among many others, Winters 1983; Brewer 1995; Brandom 1998; Boghossian 2003; Railton 2006).⁵ Since proponents of blind reasoning agree that reasoning follows determinate rules, the rules of blind reasoning are just the right candidates for the role of norms of diachronic rationality of the sort we have been looking for.

My aim in this section is to argue that this reading is not convincing. The considerations Carroll's story brings to light are indeed important, but they do not show that there must be such a thing as blind reasoning. On the contrary, they cast doubt on the very idea that reasoning is any

kind of *process* — thereby undermining, rather than supporting, the idea that the norms of good reasoning are diachronic norms of rationality.

Let us, then, look at Carroll's story. In the story, the Tortoise starts out accepting⁶ two propositions A and B , from which a third proposition Z logically follows.⁷ But the Tortoise does not accept Z , and she challenges Achilles to compel her to do so. In particular, the Tortoise is prepared to also accept C , the statement 'if A and B are true then Z must be true' (Carroll 1895, 692); but she does not see why accepting this should compel her to accept Z , when accepting A and B alone could not. Moreover, it quickly becomes clear that, while the Tortoise is prepared to accept *any* number of further conditionals of this sort (that if A , B and C are true then Z must be true, and so on), she still fails to see why she should accept Z . Apparently, there is nothing Achilles can do to make the Tortoise feel compelled to accept Z .

What are we supposed to make of the story? Obviously, the Tortoise's behavior is supposed to strike us as *puzzling*: there is something odd, and perhaps irrational, in the Tortoise's persistent failure to accept Z . At the same time, however, we are clearly not meant to think of the Tortoise as *simply confused*: if there is something wrong with her thinking, Carroll implies, it is something that will take some work to uncover. In what follows I will suggest one way to think about what is going on.

At the beginning of the story the Tortoise accepts A and B , but she does not accept Z . Is anything rationally amiss with the Tortoise at this point? It is not obvious that there is. It is true that her doxastic state is not closed under logical consequence. But even if we assume (controversially) that deductive closure is an ideal epistemic agents should strive towards, it does not follow that falling short of this ideal constitutes a failure of *rationality* (as opposed to, e.g., a case of mere ignorance or error). It does not seem irrational for you to remain agnostic about Goldbach's Conjecture, even if you accept Peano's axioms.⁸ How exactly to circumscribe the demands of rationality remains a contested issue; but, according to a standard view that I propose to follow here, failures of rationality characteristically involve internal *incoherence* or *conflict*

of some sort: being irrational involves going wrong *by your own lights*, in some sense.⁹ Failing to notice deductive consequences of what you believe does not seem to be a failing of *this* sort — a point that Achilles, for one, is willing to concede (Carroll 1895, 692).

Broome (2013, 157-158) argues that rationality requires that you believe obvious logical consequences (Modus Ponens is the example Broome discusses) of what you believe, on the condition that you ‘care’ about them. Since the Tortoise cares about whether *Z*, Broome might argue that the Tortoise is, in fact, irrational in not accepting *Z*. However, the caring condition does not suffice to make closure principles plausible as requirements of rationality. At any given time there is any number of things you may care about. So you have to prioritize: even if you care about all of *W*, *X*, and *Y*, surely rationality does *not* require of you to attend to all of them at the same time. So the fact that the Tortoise does not accept *Z*, even though she believes *A* and *B* and cares about whether *Z*, does not necessarily impugn her rationality: she may simply not have gotten around to considering the matter yet.

What about the even weaker principle, that if you care about obvious consequences of things that you accept, *and* you have competently considered the matter, then rationality requires that you accept them? I think that we should all agree that, when in Carroll’s story the Tortoise fails to accept *Z* even after considering it, her attitude is bizarre and likely irrational. But our problem is *explaining why* this is so; building this verdict directly into a rational requirement does not help us with that.

Stroud (1979, 191-2) points out that if someone were to fail to be sensitive to *any* of the deductive consequences of a proposition *A*, we would hesitate to ascribe to her any attitudes with the content that *A*, on the grounds that she manifestly does not understand it. This seems reasonable, but (as Stroud recognizes) it does not undermine the present point. Even if constraints on understanding or attitude attribution require that agents exhibit some degree of sensitivity to consequence relations, understanding surely is still compatible with a lot of

deductive ignorance. Thus there still need be nothing rationally amiss about the Tortoise's failing to accept Z , even though she accepts (and understands) things from which it logically follows.

So how can Achilles compel the Tortoise to accept Z ? The use of brain-washing or other forms of non-rational manipulation are obviously not at issue. What seems to be at issue, rather, is getting the Tortoise to accept that Z 'must be true' (Carroll 1895, 692). It is not immediately clear from the text why this would count as 'compelling' the Tortoise to accept Z . I think, however, that we can reasonably guess that Achilles's thought is this: if she accepts that Z *must be true*, the Tortoise *would* be in a position where she *is* incoherent in not accepting Z . Whether Achilles is right about this is worth investigating in its own right; unfortunately for him, however, the Tortoise never gets to accepting that Z must be true in the first place.

As we already saw, the Tortoise is prepared to accept C — the conditional that if A and B are true, then Z must be true — and it follows logically from A , B and C that Z must be true. But this does not help. If the Tortoise were prepared to willy-nilly accept the logical consequences of what she accepts, our characters would not have found themselves in this bind anyway. But since we have already granted that rationality does not demand accepting the logical consequences of what you accept, rationality will not demand of the Tortoise that she accept that Z must be true, even after she accepts C . And if accepting C does not suffice to compel the Tortoise to accept Z , accepting yet more conditionals of the same sort will not suffice either.

Now, most recent commentators think that the lesson of Carroll's story is that there must be such a thing as *blind reasoning*: a fundamental kind of reasoning that does not rest on the thinker's grasp of a rational connection between her premisses and her conclusion. Roughly speaking, the rationale behind this reading seems to be this. Carroll's story shows that grasping the connection between your premisses and your conclusion is not sufficient for reasoning: the Tortoise, after all, is aware of the connection between her premisses and her conclusion, and yet does not reason. This, a proponent of blind reasoning would continue, suggests that reasoning must have to do with something *other than* grasping the connection between your premisses and

your conclusion. But then it should be possible to reason *independently of* grasping the connection between your premisses and your conclusion: blind reasoning is possible.¹⁰

This reading is problematic, however. After all, if it really were true that reasoning does not have much to do with grasping the connection between one's premisses and one's conclusion, then the Tortoise's situation should look completely commonplace. There should be nothing peculiar about a thinker who refuses to accept acknowledged consequences of things that she explicitly accepts. But I take it that such cases genuinely *are* puzzling. Carroll himself clearly intended his readers to be puzzled by the Tortoise's behavior. And in real life, we generally do assume that pointing out to someone the logical consequences of things that they accept will have an impact on their thinking. For example, if you come to realize that your favorite philosophical view entails a counterintuitive conclusion, you face an uncomfortable choice: revise your view, or accept the counterintuitive consequence. Moreover, this is a fact about *your own perspective on things*: indifference of the sort displayed by the Tortoise will not seem to you to be an option in the circumstances.

Thus the blind reasoning interpretation seems to be missing something crucial about Carroll's story. What we learn from the story is not that reasoning is about something other than grasping facts about what follows from what; it is, rather, that we need *an account of how* facts about what follows from what have the first personal significance that they plainly do. This, I take it, is the real challenge that emerges from Carroll's story. As we will see, addressing this challenge holds lessons for the theory of reasoning.

4. Is reasoning a mental process?

Let us look a bit more closely at the story, to see why the puzzle arises. Carroll's story represents the things the Tortoise accepts as sentences written down in Achilles's notebook. The metaphor is significant (a point also emphasized by Stroud (1979)): it suggests that a subject's doxastic state, considered over time, can be broken down into *atoms* that, like sets of marks on a page,

stand in no internal relations to each other. Even if A and B logically entail Z , writing down *sentences expressing* A and B in no way commits you to writing down a sentence expressing Z .

This framework correctly captures a point that Carroll emphasizes in his story, namely that rationality does not require acceptance (or belief) to be closed under deductive consequence: there need be nothing irrational in the Tortoise's accepting A and B but not Z , as is the case at the outset. But there is a crucial further feature of doxastic states, which Carroll's atomistic framework does not easily capture.

Accepting a particular claim *commits* you to the world's being a certain way and not in others; so it commits you to accepting further claims and to rejecting others. I take the relevant notion of commitment to be familiar from our ordinary practices of giving and assessing arguments. In our example, the Tortoise starts out accepting A and B but not Z . And yet, since Z is logically entailed by A and B , there is clearly an intuitive sense in which the Tortoise is already *committed to* Z : given what she has already accepted, there is no consistent way for things to be that does not involve Z . Moreover, although commitment should be distinguished from explicit acceptance and belief, epistemic agents cannot (rationally) be *indifferent to* facts about what they are committed to. This is why recognizing that your favourite philosophical view entails some counterintuitive conclusion presents you with an uncomfortable choice (as was pointed out above): recognizing that your view entails a counterintuitive conclusion is tantamount to recognizing that you are *already committed to* that conclusion, so long as you continue to accept that view. Facts about commitment have first personal significance.

So how could Carroll's atomistic framework capture the role that facts about commitment play in our rational economy? As we see from the story, the Tortoise has no problem accepting claims of the form if A and B are true, then Z must be true. The problem is that when she accepts such claims, they simply become more sentences written down in a notebook. Nothing further follows about her doxastic state. But rational agents cannot be indifferent to facts about commitment. If you recognize that your favorite philosophical view

commits you to an unwelcome conclusion you face an uncomfortable choice; simply shrugging your shoulders, as the Tortoise does, is not an option for you. Facts about commitment, in other words, cannot be for you *just more sentences* written in a special notebook. Something is deeply wrong with the atomistic framework that underlies Carroll's story.

So how can we capture the role that facts about commitment play in our rational economy? If we set aside the atomistic framework suggested by the notebook metaphor, we can see an answer. As already explained, accepting a claim or set of claims commits you to the world's being in certain ways but not others. Thus at the outset of Carroll's story, the Tortoise accepts that the world is such that A and B are true, and accordingly is committed to ruling out ways for the world to be that are incompatible with that. Since we are not indulging in the idealization of deductive closure for doxastic states, however, we must allow that you can accept something *without actually ruling out* everything that you are thereby committed to ruling out. Given a broadly dispositional account of belief and acceptance, this type of conflict should not appear mysterious. For example, assuming that the Tortoise lacks any disposition to assent to sentences expressing Z (which, by assumption, she understands well enough), we may well conclude that, so far as she can see, Z could turn out to be false even if A and B are both true. What is *doxastically* possible from an agent's point of view may outstrip what is *logically* possible, given what she accepts.¹¹

Furthermore, since (as we have seen) deductive closure does not even seem to be a rational requirement for doxastic agents, the fact that you are now, given what you have already accepted, already committed to a certain further claim does not entail that you are *required to accept* that further claim (if you have not yet realized that your favorite philosophical view commits you to some counterintuitive conclusion, there is no sense in which you are required to accept that conclusion). But once the agent recognizes that a certain claim follows from other things that she accepts, then — assuming she continues to accept those things — that simply amounts to her *accepting* that claim. More broadly, once an agent rules out ways for things to be in which her

premisses are true but in which her conclusion fails to be true, her reasoning is *done* — there is no further process for her to engage in.¹²

It is important to note that, so long as we accept that rationality does not demand deductive closure, recognizing what follows from what in the relevant sense is *not* just a matter of accepting a conditional. If deductive closure is not a demand of rationality, there is no guarantee that accepting *A* and *if A then B* will compel me to accept *B*. By contrast, what I have been calling recognizing that *A* rationally commits me to *B* involves *ruling out* the possibility that *A* might turn out to be true while *B* is false. Thus recognizing that *A* commits me to *B* in this sense is not input into something else, *a process of reasoning* to *B*; in context, it just *is* my inferring *B* from *A* (See also Gibbons 2009; Valaris 2014, 2016).

This view of reasoning has, as we shall see in a moment, important implications for the debate regarding the (alleged) diachronic norms of rationality. Before we get to this, however, it is worth taking a moment to head off some possible objections and misunderstandings concerning the view itself.

To begin with, I do not deny that in many cases *working out* what follows from a given set of claims takes time, and genuinely constitutes a process. Sometimes this process can be protracted, and can even involve a whole community of scholars. Following Rumfitt's (2015, 35) usage, we may call this process 'deduction'. But it is crucial to note that such processes *are not processes of adopting or revising any* (first order¹³) *attitude*, and so they are *not* processes of reasoning in the sense that is relevant here. You may investigate what follows from a given set of premisses regardless of what attitude (if any) you have towards those premisses, and regardless of what (if any) attitude you end up adopting towards your conclusion. Indeed, you may investigate what follows from a set of *purely schematic* premisses (as in a logic class), towards which it is very doubtful you could have any attitude whatsoever.

Furthermore, there is no doubt that, even after having seen that a certain conclusion follows from things that you accept, you may still experience *resistance* to that conclusion.

Something like this would seem to be the case with the Tortoise (although of course in actual cases such resistance is unlikely to be motivated by sheer desire to make a philosophical point). And it might seem that this familiar possibility is in tension with the claim that inferring a conclusion can just consist in recognizing that it follows from things that you accept.

But, in fact, my proposal is perfectly compatible with this possibility. My claim is that reasoning consists in accepting a conclusion in virtue of recognizing that it follows from things that you accept; this clearly does not imply that *in every* case where you recognize that something follows from things that you accept you automatically accept that thing. There are several interpretive strategies we can use to make sense of mental states of this sort consistently with my proposal. Return to the Tortoise. Perhaps the problem with the Tortoise is some kind of Moorean incoherence. On the one hand, she *accepts* Z (after all, she accepts A and B , and that Z follows from A and B), while at the same time *denying that* she accepts Z . This kind of failure of self-knowledge, although bizarre, is not impossible. Or, perhaps, we can appeal to mechanisms of *division*: the Tortoise somehow manages to keep her recognition that Z follows from A and B apart from her acceptance of A and B . If this is correct, we can take the Tortoise to be speaking truthfully when she denies accepting Z . Either way, the Tortoise's state of mind clearly contains significant internal tension; and this is as it should be.¹⁴

These worries seem to be symptomatic of a more general concern, which was helpfully formulated by an anonymous referee. On the view developed here, inferring a conclusion from a set of premisses is a special way of *accepting* that conclusion — specifically, accepting that conclusion by taking it that it follows from those premisses. But, the referee points out, the study of reasoning is typically taken to be a study of attitude *revision*, of how we *change our minds*. And if reasoning is a kind of *change*, it would seem to follow that it is not any kind of *state*. Since acceptance is a state, my proposal would seem to be ruled out, almost by definition.

This argument, however, equivocates on the notion of change. To see this, notice that accepting something *can*, in a perfectly colloquial sense, be a kind of change — if what you

accept now is something you did not accept before. And we can legitimately wonder about the rationality of this change, for example by wondering why you think your newly accepted conclusion follows from other things that you accept. The current objection seems to be that, in speaking this way, we are ignoring what happens during the temporal instant or interval where you go from not accepting your conclusion to accepting it. But it is tendentious to claim that the study of reasoning is exclusively (or even at all!) concerned with what happens during that in-between time. Consider the following analogy. Arden changes out of his work clothes before going out to dinner. We may wonder about the appropriateness of Arden's change of outfit. But if we do so our concern is most likely not with the few minutes he spent pulling his work clothes off and his dinner clothes on, but rather with the states that come before and after.

For present purposes, the crucial point is that once we stop thinking of reasoning as a process the temptation to view rules like Modus Ponens or Modus Tollens as diachronic norms of rationality vanishes. Rules like Modus Ponens and Modus Tollens are, in the first instance, rules that govern rational commitment, not belief or acceptance. Modus Ponens tells you that if you are *committed to A* and *if A then B* then you are also committed to *B*; nothing directly follows about what you should accept or believe. To the extent that we think of reasoning as a matter of adopting or revising doxastic attitudes, these are not rules of reasoning at all (a point famously made by Harman 1986).

Finally, let us return to the Hlobil's original puzzle, to see how the present proposal allows us to address it without appealing to diachronic norms. Hlobil's agent starts out believing *not-A*, *if B then A*, and *if not-B then not-A*. She then goes on to infer *not-B*, in one case by affirming the consequent, in the other by Modus Tollens. From the present point of view, the difference between the two cases comes down to what relations the agent takes to hold among her beliefs. In the first case, the agent accepts *not-B* because she takes it to follow from *not-A* and *if not-B then not-A*; in the second, because she takes it to follow from *not-A* and *if B then A*.¹⁵ The first of these beliefs is (at least if deductive consequence is what is at issue) is false, while the second true.

This, of course, need not mean that the agent is *unjustified* or *irrational* in holding the first of these beliefs. However, just as we granted Hlobil the assumption that the agent's reasoning by affirming the consequent was irrational in this case, it seems we can equally assume that her *belief that not-B follows from not-A* and *if not-B then not-A* is irrational. But then, since it is precisely *in virtue of* this belief that she believes *not-B*, it is natural to think that this belief is not rationally all right either. Thus we can distinguish between the rational standing of the two cases without appealing to diachronic norms.

School of Humanities and Languages

UNSW Australia

¹ I wish to thank three referees for this journal for their constructive and generous comments. I also would like to thank the organizers, participants, and audience of the 2016 Leipzig Workshop on Reasoning for the very helpful exchanges.

² For one thing, Broome (2013) discusses at length various norms of *persistence* rather than change. I have nothing to say about such norms here. Thus, while the argument of this paper is friendly to 'time slice' theories of the sort defended by Brian Hedden (2015a; 2015b) and Sarah Moss (forthcoming), I do not aim to provide a *general* defence of such theories, and neither do I presuppose any of the arguments in their favor.

³ Ryle actually uses the verb 'deduce' instead of 'infer' here. But it is not clear that 'deduce' and 'infer' are synonyms — Rumfit (2015, 35-6), for one, draws a distinction between them, which I will adopt in Section 4 below. Relatedly, an anonymous referee suggests that Ryle's test fails if we substitute 'reason' for 'infer'. There is a *dialectical* sense of 'reason', about which the referee is undoubtedly right: consider, for instance, 'don't try to reason with Donald, you'll be wasting your time'. This sense of 'reason' can clearly refer to processes, but it seems equally clear that this is not the sense at issue here. Furthermore, 'reason', especially in technical contexts, can be used to refer to the activity of *constructing derivations* (much like 'deduction' in Rumfit's sense), and that too is a process that can be 'quickish or slowish'. But constructing derivations is not the same thing as revising your attitudes (setting aside attitudes about the derivation itself), and so this sense of 'reasoning' is not the one at issue either. In any case, none of the arguments in this paper turn on anything along the lines of Ryle's test.

⁴ This point is familiar from the long-running debate in epistemology regarding the basing relation. Looming large in this debate is the question whether the basing relation is causal or not (Harman 1970; Lehrer 1971). This is not our concern here.

⁵ A rather different (and more dismissive) reading of Carroll's paper reads it as attempting to raise problems about the notion of *logical consequence* (Thomson 1960; Smiley 1995). On that reading Carroll's supposed problem would be resolved by the distinction (already familiar to Frege) between premisses and rules of inference. Even if Carroll's text fails to raise problems for logical consequence, however, it may still succeed in raising difficulties about the role of logic in the theory of belief or reasoning, as more recent commentators have noted (see Stroud 1979).

⁶ While Hlobil frames his discussion in terms of belief, Carroll speaks of acceptance. I will follow Carroll here, but the lessons are directly applicable to Hlobil's argument as well.

⁷ The propositions are, respectively, 'Things that are equal to the same are equal to each other', 'The two sides of this triangle are things that are equal to the same', and 'The two sides of this triangle are equal'. Timothy Smiley (1995, 727) denies that the third follows *logically* from the other two, since he denies that the relation invoked in the argument is the logical relation of identity. I am not sure Smiley is correct, as the minor premiss may easily be read as elliptical for 'the two sides of this triangle are equal to each other *in length*', where length is something (a magnitude or quantity) literally shared between the two sides of the triangle. But even granting Smiley's point, this does not affect my reading of Carroll's argument. My reading would work just as well if we were to replace references to *Z's logically* following from *A* and *B* with references to its *deductively* following from *A* and *B*, where deductive consequence includes materially as well formally valid inferences (for example, from 'the roses are red' to 'the roses are colored').

⁸ There are views — e.g., Jarvis and Ichikawa (2013) — according to which failures of deductive closure count as failures of rationality. But examples like this one make it doubtful that the notion of rationality they rely on is the ordinary one.

⁹ It is often suggested that views of rationality as incoherence contrast with views according to which rationality consists in something more 'objective', such as responding correctly to reasons (for a statement of this contrast, see Broome (2007)). But, to the extent that the reasons in question are reasons *possessed by* or *easily accessible to* the agent, it is not clear how deep this contrast is. In particular, it looks like failing to respond correctly to reasons may just *be* a way of going wrong by your own lights. For a thorough discussion see Gibbons (2013).

¹⁰ My quick presentation of this line of thought is not intended to do justice to any one specific version of the argument, but only to capture something of its flavor. For a clear presentation, see Boghossian (2003). Proponents of blind reasoning also argue that it would be impossible for agents to always have the required grasp. I will not touch upon this argument here.

¹¹ For the idea of allowing epistemic space to outstrip logical possibility as a tool for thinking about deductive ignorance, see Jago (2014) and Rumfitt (2015).

¹² Note that this gives us a way of vindicating the idea, which I attributed to Achilles above, that getting the Tortoise to accept that Z ‘must be true’ amounts to *compelling her* to accept Z . If accepting that Z must be true amounts to ruling out all ways for things to be in which Z fails to be true, then that just *is* accepting Z .

¹³ Working out that B follows from A will normally result in your *accepting that* B follows from A ; the point is that working this out does not, by itself, imply anything regarding your attitudes to A or B themselves.

¹⁴ Note that resolving such tensions typically will involve further mental processes. Suppose you find yourself in a situation where you can see that some conclusion follows from premisses that you accept, but you have trouble endorsing the conclusion. How is the tension resolved? Perhaps the conclusion of your reasoning will eventually sink in, and any residual dispositions to deny that you believe it will fade away. Or, conversely, you may lose confidence in your premisses (or in your assessment of their relation to her conclusion). For present purposes, however, what matters is that there is no reason to identify such processes with the process of *reasoning*. For one thing, they have much broader application than just inference — you may struggle to believe your own eyes as much as the conclusions of your inferences, and the processes involved in resolving the tension would seem to be much the same in both cases. For another, they seem to be very different in nature from what we would normally think of as reasoning: they involve things such as coming to terms with emotionally unsettling ideas or simply getting used to novel ones. These are not likely candidates for a process of reasoning.

¹⁵ In putting matters this way, it is essential that we keep track of the *subset* of her beliefs from which the agent infers *not-B* (if we do not do this, then the set of ways for things to be ruled out to arrive at the conclusion is the same in both cases). But note that Hlobil’s appeal to rules of inference requires carving up the agent’s doxastic state in the same way as well: a rule such as Modus Tollens applies only to a proper subset of the agent’s beliefs, not to her doxastic state as a whole. This fineness of grain is necessary if we are to capture not just *what* an agent believes, but also *why*.

References

- Boghossian, Paul (2003) 'Blind Reasoning.' *Aristotelian Society Supplementary Volumes* 177: 225–48.
- . (2014) 'What Is Inference?' *Philosophical Studies* 169 (1): 1–18.
- Brandom, Robert (1998) *Making It Explicit: Reasoning, Representing, and Discursive Commitment* (Cambridge, MA: Harvard University Press).
- Broome, John (1999) 'Normative Requirements.' *Ratio* 12: 398–419.
- . (2013) *Rationality through Reasoning* (Chichester: Wiley Blackwell).
- . (2015) 'Synchronic Requirements and Diachronic Permissions.' *Canadian Journal of Philosophy*. Accessed online from <<http://dx.doi.org/10.1080/00455091.2015.1118231>>
- Carroll, Lewis (1895) 'What the Tortoise Said to Achilles.' *Mind* 104: 691–693.
- Gibbons, John (2009) 'Reason In Action'. In *Mental Actions*, Lucy O'Brien and Matthew Soteriou eds., (Oxford: Oxford University Press) 72–94.
- . 2013. *The Norm Of Belief*. Oxford: Oxford University Press.
- Harman, Gilbert (1970) 'Knowledge, Reasons and Causes'. *Journal of Philosophy* 67: 841–855.
- . (1986) *Change in View* (Cambridge, MA: MIT Press).
- Hlobil, Ulf (2015) 'There Are Diachronic Norms of Rationality.' *Thought: A Journal of Philosophy* 4: 38–45.
- Jago, Mark (2014) *The Impossible: An Essay on Hyperintensionality* (Oxford: Oxford University Press).
- Lehrer, Keith (1971) 'How Reasons Give Us Knowledge'. *Journal of Philosophy* 68: 311–313.
- Moss, Sarah (Forthcoming) 'Time Slice Epistemology And Action Under Indeterminacy'. In *Oxford Studies in Epistemology* 5, Tamar Szabó Gendler and John Hawthorne (eds.). Accessed from <<http://www-personal.umich.edu/~ssmoss/>>, 15/06/2016.
- Railton, Peter (2006) 'How to Engage Reason: the Problem of Regress'. In *Engaging Reason: Themes from the Moral Philosophy of Joseph Raz*, Ray Wallace, Philip Pettit, Stephen Scheffler and Michael Smith eds. (Oxford: Oxford University Press): 176–201.
- Podgorski, Abelard (2016) 'A Reply To The Synchronist'. *Mind*. Accessed from <<http://mind.oxfordjournals.org/content/early/2016/04/21/mind.fzv153>>
- Pollock, John and Joseph Cruz (1999) *Contemporary Theories of Knowledge* (New York: Rowman and Littlefield).
- Rumfitt, Ian (2015) *The Boundary Stones of Thought* (Oxford: Oxford University Press).
- Smiley, Timothy (1995) 'A Tale of Two Tortoises'. *Mind* 104: 725–736.
- Stroud, Barry (1979) 'Inference, Belief and Understanding'. *Mind* 88 (350): 179–196.
- Turri, John (2010) 'On The Relationship Between Propositional And Doxastic Justification'. *Philosophy and Phenomenological Research* 80: 312–326.
- Valaris, Markos (2014) 'Reasoning and Regress.' *Mind* 123 (489): 101–27.
- . (2016) 'What Reasoning Might Be'. *Synthese*, available on Early View.
- Winters, Barbara (1983) 'Inferring.' *Philosophical Studies* 44: 201–20.