

Reasoning and Regress

MARKOS VALARIS

University of New South Wales

m.valaris@unsw.edu.au

Published in *Mind*. Please cite published version.

Regress arguments have convinced many that reasoning cannot require beliefs about what follows from what. In this paper I argue that this is a mistake. Regress arguments rest on dubious (although deeply entrenched) assumptions about the nature of reasoning—most prominently, the assumption that believing p by reasoning is simply a matter of having a belief in p with the right causal ancestry. I propose an alternative account, according to which beliefs about what follows from what play a constitutive role in reasoning.

1. Introduction

Reasoning is a way of bringing one's beliefs in line with one's evidence. Good reasoning is a way of knowing. But what exactly does good reasoning consist in? More specifically, under what conditions does one know a proposition p by reasoning from a set of premisses R ?

Consider, for example, how Mary may gain knowledge by reasoning through the following simple argument:

- (1) If this snake is a *vipera aspis* then it is venomous
- (2) This snake is a *vipera aspis*
- (3) Therefore, this snake is venomous

Supposing that Mary knows (1) and (2) independently of (3), what else is required for Mary to know (3) in this way?

One further necessary condition would seem to be that Mary must *know that* (3) follows from (1) and (2): she must recognize that (1) and (2) constitute evidence (indeed, conclusive evidence) for (3). Otherwise, it is hard to see how her believing (3) could be properly *based on* her knowledge of (1) and (2).

In general, it is clearly possible to have a true belief, and to possess good evidence for that belief, but nevertheless lack knowledge because one's belief is not properly based on one's evidence. The suggestion is that the following knowledge requirement (KR) is a necessary condition for knowing p on the basis of R through reasoning:

(KR) If one knows p by reasoning from R , one knows that p follows from R

My aim in this paper is to defend (KR) by developing an account of reasoning that entails it. As I will argue, (KR) holds because of the following belief requirement (BR):

(BR) If one believes p by reasoning from R , one believes that p follows from R

The relationship between the two principles is straightforward: if *believing* p by reasoning from R requires believing that p follows from R , it is natural to think that *knowing* p by reasoning from R requires knowing that p follows from R .¹

Intuitive though BR and (KR) may be, they have proven unpopular in contemporary epistemology and philosophy of mind. Many philosophers, inspired by Lewis Carroll's (1895) famous regress argument, have been persuaded that trying to find a role for beliefs about what follows from what in reasoning invites a vicious infinite regress. As a result, they have tended to reject (BR) and (KR).²

¹ How should we understand the proposition that p follows from R ? This relation holds between R and p just in case R constitutes *conclusive evidence* for p , i.e. just in case R justifies *outright belief* in—and potentially knowledge of— p . I believe that most of what I have to say in this paper has analogues in the case of partial belief and merely probabilifying evidence, but I cannot go into this here.

² Note that this is not the only regress argument to be used against (KR) and (BR). The even more venerable *epistemic regress* argument—which begins by asking how, if not by further reasoning, the beliefs about what follows from what required by (BR) might be *justified*—has been very influential too (see, e.g., the discussion of 'inferential internalism' and 'inferential externalism' in Fumerton 1995). Although my account of reasoning is designed to avoid the epistemic regress argument, my focus in this paper is not going to be on it: the

In what follows I will argue that this response to the regress argument is mistaken. The belief and knowledge requirements only generate an infinite regress in the context of further assumptions, which are not beyond question. Most recent philosophers conceive of reasoning as a *causal process*—a process in which some of one’s existing beliefs (or other reason-giving mental states, if there are any) cause the acquisition of a new belief.³ On this view, believing p by reasoning is just a matter of having a belief in p with the right causal ancestry. From this starting point, it is natural to conclude that, if believing that p follows from R is necessary for reasoning from R to p , then that belief has to figure as a link in the causal chain that connects one’s beliefs in the members of R to one’s belief in p . As we will see in section 3, with these two assumptions in place (BR) and (KR) do indeed threaten a vicious regress.⁴

The response to the regress argument I will recommend is rejecting this framework. On the view I will suggest, reasoning is not to be identified with a causal process (although, of course,

epistemic regress argument concerns the conditions of justification *in general* (at least insofar as it involves beliefs resting on other beliefs), whereas I am interested in the psychological phenomenon of *reasoning* in particular.

³ Of course, it is recognized that sometimes reasoning may instead reinforce an existing belief, or lead one to abandon it. Furthermore, sometimes we reason hypothetically, i.e., without commitment to the truth either of our premisses or of our conclusion; and sometimes we trace inferential relations *via* schematic arguments, which lack determinate premisses or conclusion altogether. It is unlikely that any of these cases introduce significant new issues, however, so most authors do not discuss them directly. I will follow the same practice.

⁴ Work that rejects (BR) and (KR), apparently under the influence of the causal process view of reasoning, includes Winters 1983, Boghossian 2003, Broome 2006, Railton 2006, and Wedgwood 2006. Johnston 1988, 87–8 employs the regress argument to argue against some views that entail (BR), but he does not seem to take it to rule out all such views. Taking an approach rather different from any of the above, Brewer 1995 and Rödl 2007 argue that standard versions of the causal process view go wrong in failing to incorporate awareness of inferential relations *in* the relevant causal relation. (Brewer nonetheless rejects (BR), on the basis of regress worries; his preferred view seems to involve *sub-doxastic* awareness of inferential relations.)

reasoning may *involve* such processes). Believing something by reasoning is not just a matter of having a belief with the right causal ancestry. Accordingly, it is a mistake to think that beliefs about what follows from what can be involved in reasoning only as links in a causal chain. Indeed, as I will argue, such beliefs play a *constitutive*, rather than *causal*, role in reasoning.

The paper proceeds as follows. In section 2 I clarify the notion of reasoning that is at issue, and I address some methodological worries. In section 3 I present the regress argument, which is typically taken to show that beliefs about what follows from what cannot be involved in reasoning. As I argue, in order to respond to this argument proponents of (BR) and (KR) need to develop an alternative to the causal process view of reasoning.⁵ This is my task in section 4. In section 5 I clarify my account and I respond to some objections.

Finally, before moving on, let me note that instead of (KR) and (BR) one might prefer the following, more cautious, formulations of the recognition and belief requirements:

(KR*) If one knows p by reasoning from R , one is *in a position to* know that p follows from R

⁵ Millar 2004, 192–202 argues for something along the lines of (BR), on the grounds that otherwise it would remain obscure how facts about what follows from what could have explanatory relevance. But Millar does not attempt to give an account of reasoning that explains how beliefs about what follows from what figure in it. This, I believe, is a shortcoming. After all, the problem with (BR) and (KR) is not that they lack intuitive appeal; the problem is that the dominant conception of reasoning has made it hard to see how they could be true. So long as this conception of reasoning is not challenged, arguments like Millar's will simply spur philosophers to find strategies to account for the intuitions that seem to support (BR) and (KR) in alternative ways. Wedgwood 2006 seems to do just that in relation to Millar's explanatory worries, by developing a version of the causal process view of reasoning in which *normative facts* play a direct causal role, unmediated by beliefs about what follows from what.

(BR*) If one believes p by reasoning from R , one *is disposed to* believe that p follows from R ⁶

Indeed, as will become clear later on, there may be good reasons to prefer (KR*) and (BR*) over their un-starred cousins. For the main purposes of this paper, however, the differences between the two formulations do not matter much. If, as the causal process view of reasoning would have it, believing p by reasoning from R is just a matter of having a belief in p with the right causal ancestry, a mere disposition to believe that p follows from R cannot be required: even if dispositions sometimes are causes (a controversial claim), *this* disposition would seem to be the wrong type of thing to figure in the aetiology of belief. Thus it is plausible that, from the perspective of the causal process view of reasoning, the disposition to believe that p follows from R could be required for believing p by reasoning from R only if *the belief that is its manifestation* were required. In other words, assuming the causal process view of reasoning, (BR*) could be true only if (BR) were true. Since my principal aim in this paper is to bring out the shortcomings of the causal process view and to suggest an alternative, I will concentrate on the simpler claims (BR) and (KR) rather than (BR*) and (KR*).

⁶ What is it to be *in a position to know* a proposition, and what is it to be *disposed to believe* a proposition? Being in a position to know p does not entail that one knows p , but neither is it entailed by one's merely being able to know p (perhaps as the result of lengthy investigations). Being in a position to know p requires having done all the epistemic work necessary to know p , so that, if one is in a position to know p and one considers the matter, one thereby comes to know p (see Williamson 2000, 95). Being disposed to believe p similarly does not entail believing p , but also is not entailed by its being the case that there are some circumstances that would lead one to believe p . Being disposed to believe p requires being in a state such that, were one simply to consider the matter in the right way, one would believe p (see Audi 1994).

2. Explicit Reasoning and Believing for a Reason

My goal in this section is to bring our topic into sharper focus and to address some methodological questions this sharpening raises.

So what is reasoning? If one believes p by reasoning from R , one's belief in p is *based on* one's belief in the members of R ; equivalently, R comprises one's *reasons for* believing p .⁷ However, not every case of believing for a reason counts as reasoning. Reasoning, as I intend the term to be understood here, is a personal-level, conscious activity; and it seems clear that many of our beliefs are based on other beliefs without such explicit reasoning. For example, suppose that while watching a football game I come to believe that a goal was just scored. It is plausible that this belief of mine is based on more basic perceptual knowledge, as well as my knowledge of the rules of the game. Nevertheless, to the extent that I never explicitly considered the matter, this is not a case of reasoning in our sense. Reasoning is a sophisticated species of the broader phenomenon of *believing for a reason*.

This fact is bound to raise some methodological worries. Does explicit reasoning deserve to be taken as a self-standing philosophical topic at all? And, even if it does, shouldn't an investigation into the nature of explicit reasoning follow on the heels of an investigation of the prior topic of believing for a reason?

To begin with, note that by restricting our attention to explicit reasoning we are not begging any questions: it does not immediately follow from the claim that reasoning is a conscious activity that it involves knowledge or beliefs about what follows from what. In fact, opponents of (BR) and

⁷ I am assuming, somewhat controversially, that reasons for belief are propositions rather than mental states. Nothing hangs on this.

(KR) often restrict their discussion to explicit reasoning too. For example, here is how John Broome characterizes his topic:

Reasoning is an activity—something we do—through which we can satisfy some requirements [of rationality] ... Some unconscious processes could be called unconscious reasoning. But in this paper I am interested only in conscious processes, and I shall give the name ‘reasoning’ to those ones only. Unconscious processes are not activities, and I am interested in reasoning as an activity. (Broome 2006, 184)

Despite this, Broome goes on to reject (BR) by appealing to Carroll’s regress argument. Ralph Wedgwood 2006 makes very similar remarks as well. Moreover, even those of the critics of (BR) who don’t explicitly *restrict* their discussion to explicit reasoning, clearly wish to *include* explicit reasoning within the scope of their discussion (see, for instance, the cases Winters 1983 and Boghossian 2003 pick to illustrate their accounts); so focusing on explicit reasoning does not beg the question against them either.

But I also think we should reject the inference from the premiss that explicit reasoning is just a species of believing for a reason to the conclusion that the latter is the philosophically prior topic. In general, it does not follow from the fact that *A* is a species of *B* that *B* is the philosophically prior topic. This is because it might be the case that, although not all *Bs* are *As*, *As* are the *paradigm cases* of *B*—that is, the cases by reference to which we explain what it is to be a *B*. This, I believe, is the structure that we find here. Unreflectively believing for a reason—like explicit reasoning but unlike, say, a conditioned reflex—is an exercise of rationality; and we should expect to get a clearer picture of what exercises of rationality involve by looking into the explicit ones first. In other words, the broader phenomenon of believing for a reason might be understood by reference to the special, but also paradigmatic, case of explicit reasoning.

My argument in the rest of this paper will be only that (BR) and (KR) are true of explicit reasoning. I will not attempt to explore the consequences of this fact for the broader phenomenon of believing for a reason.

3. The Regress Argument

As I pointed out at the beginning of this paper, (KR) and (BR) seem intuitively attractive. But if beliefs about what follows from what are really required reasoning, then it seems fair to request an account of reasoning which shows *why* such beliefs are necessary: what indispensable role do beliefs about what follows from what play in reasoning? According to contemporary orthodoxy, no coherent account of reasoning can meet this demand: as regress arguments are supposed to have shown, reasoning would be impossible if it had to incorporate beliefs about what follows from what.

My overarching aim in this paper, of course, is to argue that this claim is mistaken. We can construct a coherent and plausible account of reasoning in which beliefs about what follows from what play an indispensable role. I do think, however, that regress arguments show that there are difficulties in combining the currently dominant conception of reasoning—namely, the conception of reasoning as a causal process—with (BR). This is what I will try to show in this section. In coming sections I am going to suggest that it is the causal conception of reasoning that should be rejected, rather than (BR).

The relevant regress argument traces back at least to Carroll 1895, and focuses on the idea that the belief that p follows from R provides a subject who reasons from R to p with a further *premiss* for her reasoning.⁸ Remember Mary's argument from section 1:

⁸ For more recent versions of the regress argument see Johnston 1988, Boghossian 2003, and Railton 2006. I do not follow the letter of these expositions here, as they involve detail that is not relevant for our purposes.

- (1) If this snake is a *vipera aspis* then it is venomous
- (2) This snake is a *vipera aspis*
- (3) Therefore, this snake is venomous

We would intuitively describe her as reasoning to (3) from (1) and (2). According to the view under consideration here, however, this is not so: Mary really reasons to (3) from (1), (2), *and* the proposition that (3) follows from (1) and (2). This, however, is clearly disastrous. If there were any reason to appeal to Mary's belief that her conclusion follows from her premisses when her premisses included just (1) and (2), it seems that that reason has not lost any of its force now that her premisses also include the proposition that (3) follows from (1) and (2). Thus, if we felt that in the former case we had to augment Mary's set of premisses with the proposition that her conclusion follows from her premisses, it seems that we should feel the same way about the latter case too. Should we add yet another premiss, then? It is clear that doing so would only push the question further back one step. It seems that we have embarked on an infinite regress.

Now, the general structure of a successful response to the regress argument is clear: we should simply *deny* that the role of the belief that one's conclusion follows from one's premisses is to provide one with an extra premiss. That belief must play a role in one's reasoning distinct from the role that one's beliefs in one's premisses play.⁹

So what role can that be? According to current orthodoxy, reasoning is simply a causal process in which certain beliefs cause the acquisition of a further belief. On the most straightforward version of this view, believing *p* by reasoning from a set of premisses *R* consists in having a belief in

⁹ This distinction may be thought to correspond to the distinction in formal systems between *premisses* and *rules of inference*. Indeed, Carroll's regress argument is often used in introductory logic classes to motivate this distinction. But our topic here is different: our topic is the psychological phenomenon of reasoning, not the structure of formal systems.

p which is caused by one's believing R . What role could one's belief that p follows from R play on this model? It will clearly not do to say that one's belief that follows from R is *also* one of the contributing causes of one's belief in p . To do so would simply amount to saying that one's belief that p follows from R provides one with another premiss—so that one does not really reason to p from R , but rather from R *and* the proposition that p follows from R . But to say this is obviously to take the first step on the regress.

Could one say that the belief that p follows from R is, rather than another contributing cause, a necessary *enabling condition* of one's believing R causing one to believe p ? But why should that be the case? After all, it is surely *possible* that one's believing R might cause one to believe p , even if one does not believe that p follows from R . Causal relations among beliefs do not, in general, need to be accompanied by beliefs about what follows from what.¹⁰

Could a more elaborate version of the causal story do better? Perhaps one might suggest that the belief that one's conclusion follows from one's premisses does not directly contribute causally to the concluding belief. Perhaps it figures in one's reasoning by contributing causally to a higher-order belief that one *ought to* believe one's conclusion, which then causes one to believe one's conclusion through some further process. Thus, in order to get to believe (3), Mary would first need to get to believe that she ought to believe (3)—and her belief that (3) follows from (1) and (2) would play a role in getting her to have this normative belief.¹¹

¹⁰ Couldn't we simply say that, although causal relations among beliefs do not *in general* need to be accompanied by beliefs about what follows from what, we only call 'reasoning' the ones that *are* so accompanied? But notice that saying this would amount to giving up on the idea that there is some deep explanation, in terms of the nature of reasoning, for why (BR) holds. We would be holding on to (BR) by *fiat*.

¹¹ Such higher-order manoeuvres are considered in Johnston 1988 and Broome 2006.

But now, questions arise concerning both the process through which one's belief that one's conclusion follows from one's premisses contributes to the formation of this higher-order belief, and the process through which the higher-order belief contributes to the formation of a belief in one's original conclusion. Are those processes to be construed as instances of reasoning? Does Mary reason to the conclusion that she ought to believe (3) from (1), (2), and the proposition that (3) follows from (1) and (2), for example? If yes, then we have to contend with regress problems all over again, for we will need to find a place for Mary's belief that the proposition that she ought to believe (3) follows from (1), (2) and the proposition that (3) follows from (1) and (2).

But if these processes are not processes of reasoning, then what are they? And, most importantly, why are such non-reasoning (and thus non-regress threatening) processes available to get one from one's premisses to this higher-order belief, and then from that higher-order belief back again to one's conclusion, but are not available to get one directly from one's premisses to one's conclusion? Why is it that Mary's belief that she ought to believe (3) *can* be caused without her believing that the proposition that she ought to believe (3) follows from (1), (2) and the proposition that (3) follows from (1) and (2), while her belief in (3) cannot similarly be caused without her believing that (3) follows from (1) and (2)? It is not easy to see how there could be satisfying answers to these questions.

It seems, therefore, that there is a fundamental difficulty with combining (BR) with the causal view of reasoning, even in its more elaborate, higher-order forms. I think, however, that we should not conclude from this that (BR) is false: on the contrary, I think we should take this as a reason to look for alternatives to the causal view of reasoning. This is what I aim to do in the rest of this paper.

4. An Account of Reasoning

Let us begin by drawing a distinction between *basic* and *non-basic* reasoning. Suppose that Fred is a police detective investigating a murder. Fred finds a bloody knife in the garden and concludes that the butler committed the murder. According to (BR), this entails that Fred believes that it follows from the presence of the bloody knife in the garden that the butler did it. But, of course, we can go on to ask: *why* does Fred believe this? How did he arrive at the belief that it follows from the presence of the bloody knife in the garden that the butler committed the murder? We expect that this question will have a substantive answer, which, most likely, will involve further reasoning. That is what makes this a case of non-basic reasoning. Non-basic reasoning will be the topic of section 4.1.

Now, it is clear that not all reasoning can take this form: if all reasoning relied on further reasoning, reasoning could never get started. Some reasoning must be basic, in the sense that it does not rely on further reasoning. It has sometimes been thought that this fact alone suffices to rule out (BR) and (KR), at least so long as we are not willing to postulate a special-purpose faculty of non-inferential insight into facts about what follows from what (see, e.g., Fumerton 1995, 199-200 and Wedgwood 2006, 675). As I will argue in section 4.2, however, this is a mistake: (BR) and (KR) can be satisfied without further reasoning, and without the benefit of a special-purpose faculty of insight.

4.1 Non-Basic Reasoning

Suppose that one believes R and that p follows from R . What else might it take for one to count as believing p by reasoning from R ? The crucial point here is that, if one believes both R and that p follows from R , then—barring inattention or irrationality—one thereby believes p . If Fred believes that there is a bloody knife in the garden, and he comes to realize that it follows from this that the butler did it, his reasoning is done—no further process is required to get him to believe that the

butler did it. In general, the relation between believing that one has conclusive evidence for a proposition and believing that proposition is constitutive, not merely causal.¹²

We can find independent support for this claim in the observation that instances of the schema ‘ R and p follows from R , but I do not believe p ’ are as unassertable as instances of the (so-called ‘omissive’¹³) Moorean schema ‘ p and I do not believe p ’, and apparently for similar reasons. Although the truth-conditions of sentences of either of these forms need involve no incoherence, asserting them does seem to involve incoherence.¹⁴ Now, although the debate over the deeper sources of the incoherence of Moorean assertions is on-going, the surface phenomenon seems clear enough. What one conveys about one’s state of mind by asserting p in an intuitive sense *conflicts with* or *contradicts* what one conveys by asserting ‘I do not believe p ’ (why asserting p has this effect is a difficult question, which does not concern us here).¹⁵ Similarly, then, we should expect what one conveys about one’s state of mind by asserting ‘ R and p follows from R ’ to conflict with what one conveys by asserting ‘I do not believe p ’. If we assume that by asserting ‘ R and p follows from R ’ one

¹² A similar view is suggested—though not explicitly defended—in Hieronymi 2006, 51: ‘if you take certain reasons to show that p , you *therein* believe that p ’.

¹³ For the distinction between ‘commissive’ and ‘omissive’ versions of Moore’s paradox see Williams 1979.

¹⁴ It is worth noting that it is doubtful whether either of these forms is *absolutely* unassertable. For example, Mary might acknowledge that the evidence for her son’s guilt is conclusive, while also feeling that she cannot bring herself to really believe that he is. In such circumstances she might say something along the lines of ‘the evidence for Bob’s guilt is conclusive, but I just can’t bring myself to believe that he is’, or even ‘I know that Bob is guilty, but I just cannot believe that he is’. However, such cases do not detract from the point here, because it is very plausible that the subject in such cases really is in an incoherent, or conflicted, state—indeed on a natural reading of the above statements this is precisely what Mary is trying to communicate. I will discuss such cases in more detail in sect. 5.2.

¹⁵ Moore himself speaks of a conflict between what one’s assertion ‘means’ and what it ‘implies’ (Moore 1993, 209-210). By using the generic term ‘convey’, I hope to sidestep all the controversial—and for present purposes irrelevant—issues in pragmatics and semantics in this area.

conveys that one believes R and that p follows from R , and that the latter state in a rational and attentive mind constitutes believing p , the expectation is borne out: it is readily intelligible why this conflicts with what one conveys by asserting ‘I do not believe p ’.

On the other hand, the expectation would not be borne out if the state of mind one conveys by asserting ‘ R and p follows from R ’ were *only contingently* related to believing p . There is no incoherence in denying that one is in a mental state M while also ascribing to oneself a mental state that is contingently associated with being in M . Believing that a cake is delicious is a state that tends to contribute to having a desire to eat some of the cake in question. But this connection is a merely contingent causal one, and so we have no trouble making sense of a person who asserts ‘this cake is delicious’ while denying having a desire to eat any—perhaps by going on to assert ‘but I have no desire to eat any’. By contrast, it is hard to make sense of someone’s asserting an instance of ‘ R , and p follows from R , but I do not believe p ’. The difference, I suggest, is due to the fact that there is normally no gap between believing that one has conclusive evidence for a proposition and believing that proposition—the relation between these states is constitutive.¹⁶

I suggest that we take this thought as the core of our conception of reasoning, at least in the non-basic case: non-basic reasoning *consists in* determining that (by one’s own lights, of course) one’s conclusion follows from one’s premisses, and thereby believing one’s conclusion. This conception of non-basic reasoning obviously entails (BR): non-basic reasoning just *is* believing that one’s conclusion follows from one’s premisses, and thereby believing one’s conclusion. Moreover, this

¹⁶ It might be suggested that, rather than outright *belief* in p , believing R and that p follows from R constitute some sort of *commitment* to p . It is hard, however, to see what exactly this talk of commitment comes down to. It cannot be a merely normative commitment to p , for one is already normatively committed to p by believing R alone (since R is, by hypothesis, conclusive evidence for p). But if the sort of commitment at issue also involves, say, dispositions to act and to reason on the basis of p , how does it differ from a belief in p ?

fact pre-empts the threat of regress: it is simply not the case that beliefs about what follows from what must figure in reasoning as links in a causal chain, and thus no question arises of how they could possibly do so.

Furthermore, as I already suggested in section 1, if (BR) is true then (KR) naturally follows: if beliefs about what follows from what are necessary for believing by reasoning, then knowledge of what follows from what is necessary for knowing by reasoning. Supposing that (BR) is true, and that one believes p by reasoning from R , it is very hard to see how that belief could amount to knowledge unless one's belief that p follows from R amounted to knowledge as well. Thus (KR) seems like a reasonable requirement, at least in the case of non-basic reasoning.

4.2 Basic Reasoning

We have so far seen how, in the case of non-basic reasoning, it is possible to hold on to (BR) and (KR) while avoiding regress. Our business is not done yet, however, because non-basic reasoning cannot be all the reasoning that there is. Some reasoning must be basic, in the sense that it does not rely on further reasoning. If (BR) is true, then this entails that the beliefs about what follows from what involved in basic reasoning must not be arrived at through further reasoning.

Let me start with an outline of how my view is supposed to work and how I plan to argue for it. Consider again our example from section 1: Mary comes across a vividly patterned snake and reasons as follows:

- (1) If this snake is a *vipera aspis* then it is venomous
- (2) This snake is a *vipera aspis*
- (3) Therefore, this snake is venomous

Since this argument is an instance of an elementary inference rule, Mary's performance is plausibly an instance of basic reasoning. My argument in what follows is going to be that, given certain

assumptions about our capacity for self-knowledge,¹⁷ Mary's believing (3) by reasoning from (1) and (2) entails that she *believes that* (3) follows from (1) and (2). Her reasoning still qualifies as basic, because this belief of hers is not itself the product of further reasoning. Indeed, I will argue that there is no independent explanation for why she believes that (3) follows from (1) and (2). Mary believes this simply because she believes (3) by reasoning from (1) and (2), and—given certain assumptions about our capacity for self-knowledge, which I will explain and defend in what follows—this constitutively entails that she believes that (3) follows from (1) and (2).

In a nutshell, the point is this. No-one would deny that, in inferring (3) from (1) and (2), Mary shows that in *some sense* she takes it that (3) follows from (1) and (2). My argument will be that, assuming that Mary is equipped with a capacity for self-knowledge, we should simply understand this to mean that she *believes that* (3) follows from (1) and (2). This is why (BR) holds in basic reasoning.

But why should we accept that Mary believes (3) on the basis of (1) and (2) at all? I assume that Mary possesses ordinary logical competence, and that this entails that her beliefs about matters she is actively considering tend to be closed under *modus ponens*. I will simply assume that we have some such capacities to conform our beliefs to general patterns. Such capacities operate largely unconsciously, and might be either hard-wired or the result of training. The assumption that we have such capacities is not controversial for present purposes: it is certainly not denied by the opponents of (BR) and (KR). The point of contention is whether such capacities involve beliefs about what follows from what. Accordingly, this is the issue I focus on here.

¹⁷ Note that although I use the term 'self-knowledge' to characterize the relevant capacity, I do so only for reasons of economy, and because it is conventional in the literature. What matters for present purposes is only that, in the relevant cases, we have *true* beliefs about what we believe and why; whether such higher-order beliefs also qualify as knowledge is not important here, and I am not going to argue for it.

Let us now put some flesh on these bare bones. In what follows I will rely upon the following claim (which I call (HOB), for ‘higher-order belief’):

(HOB) If one believes p by reasoning from R , one believes, without observation or inference, that one believes p for reason R

Some comments about (HOB) and its role in my argument are in order. To begin with, one might wonder whether it might not be better to replace (HOB) with a weaker principle (HOB*) which, instead of requiring actual higher-order beliefs, would only require that one be disposed *to* have such beliefs. Perhaps, one might think, even explicit reasoning does not require of one to have any beliefs about one’s own beliefs. I will remain agnostic on this. Adopting (HOB*) instead of (HOB) would require moving from (BR) and (KR) to the weaker claims (BR*) and (KR*) introduced at the end of section 1, which only require that one be disposed to believe that p follows from R , and that one be in a position to know that p follows from R . But, as I already argued in section 1, adopting (BR*) and (KR*) would make no substantive difference to the main line of argument of this paper. For this reason, I will continue to employ the simpler formulations (HOB), (BR) and (KR).

More threateningly, one might worry that appealing to (HOB) at this juncture would beg important questions. For my argument in defence of (BR) and (KR) to have force it must take as its starting point a conception of basic reasoning that can be agreed upon by all sides. But many arguments for claims like (HOB) start from premisses about the nature of reasoning along lines similar to (BR) and (KR) (e.g., Shoemaker 1988 and Burge 1996). If, however, accepting (HOB) depends upon already finding these requirements plausible, we cannot appeal to (HOB) in an argument for the requirements themselves. The issue, then, is whether we can reverse this order of explanation—i.e., whether we can start from an independent defence of (HOB), and then proceed to use (HOB) in arguing for a conception of reasoning that entails (BR) and (KR).

I believe that we can. (HOB) is just a special case of a more general claim about self-knowledge: it is a necessary truth about us that we are often (though not always) in a position to know *what* we believe and *why*, without observation or inference. Although I cannot offer a full defence of this claim here, in the next few paragraphs I want to suggest that it can be defended, without relying on anything along the lines of (BR) and (KR).

To begin with, we can argue independently of (BR) and (KR) that we are often (though not always) able to know both what we believe and why, without relying on observation or inference. This, I think, is one lesson of the phenomenon of *transparency*. As several philosophers have pointed out, many of our beliefs are transparent to the world, in that we can self-ascribe them simply by answering questions about their subject matter, rather than on the basis of special psychological or behavioural evidence. As Gareth Evans famously puts it:¹⁸

[I]n making a self-ascription of belief, one's eyes are [...] directed outward—upon the world. If someone asks me 'Do you think there is going to be a third world war?' I must attend, in answering him, to precisely the same outward phenomena as I would attend to if I were answering the question 'Will there be a third world war?' I get myself in a position to answer the question whether I believe that *p* by putting into operation whatever procedure I have for answering the question whether *p*. (Evans 1982, 225)

Transposing Evans's talk of self-ascriptions to talk about higher-order beliefs, the point seems to be that one can form a higher-order belief that one believes *p* simply by considering the first-order question whether *p*, and answering it in the affirmative. Moreover (although Evans does not say this), it seems clear that one can also form beliefs about one's reasons for believing *p* in this way, since they will be just the considerations that move one to answer in the affirmative.

¹⁸ See also Edgley 1969, Gallois 1996, Moran 2001, and Byrne 2005, 2011.

Now, in arguing for the above we have made no use of anything along the lines of (BR) or (KR). Still, this does not take us all the way to (HOB), for although the phenomenon of transparency suggests that we have a capacity for higher-order beliefs which does not involve reasoning from psychological or behavioural evidence, it does not show that such higher-order beliefs do not rest on reasoning from evidence of other kinds. Indeed, according to some views, transparency involves reasoning according to a novel pattern, which allows one to infer ‘I believe p ’ from p (see Gallois 1996, Byrne 2005, 2011). If such views were correct, the phenomenon of transparency could offer us no help in an account of basic reasoning—since basic reasoning is supposed to be reasoning that depends on no other reasoning. But I believe that these views are not correct.

If one believes one thing by reasoning from another, then one’s *reasons for* believing the former must at least include the latter. This is not the case, however, for the reasoning allegedly involved in transparency. Suppose Evans really did come to believe that he believes that there will be a third world war just by answering in the affirmative the first-order question whether there will be one. Still, it seems false that Evans’s reason for believing that he believes that there will be a third world war is that there will be one. Certainly, if one asked Evans, ‘what is your reason for believing that you believe that there will be a third world war?’, it would sound very odd for him to reply: ‘that there will be a third world war’. This, I believe, strongly suggests that the inferential view of transparency is not correct.¹⁹

So what does transparency involve? This is not the place to go into details, but I suggest that we may simply accept that first-order beliefs constitutively involve a disposition to form a corresponding higher-order belief, without observation or inference. That disposition is triggered

¹⁹ I discuss the inferential account of transparency in more detail in [omitted for review].

when the first-order belief is actively entertained in the right way—including, for example, considering the question whether p and answering it in the affirmative, in the context of the transparency procedure.

Now of course the above is not meant to be a complete account of our ability to form higher-order beliefs about our own beliefs. All I have tried to show is that, by taking transparency as our starting point, we can argue for the possibility of non-observational, non-inferential higher-order beliefs without relying on anything along the lines of (BR) or (KR). If this is accepted, then on the assumption that beliefs had by reasoning are among those for which transparency holds, (HOB) itself follows²⁰—still without any need to presuppose anything along the lines of (BR) and (KR).

How, then, can we use (HOB) to argue that (BR) holds in basic reasoning? The sense of ‘reason’ in (HOB) is explanatory, rather than normative: R comprises the considerations that *move* one to believe p , even though they might have no tendency to support p . Moreover, most of us are aware that we sometimes hold beliefs for *bad* reasons—i.e., reasons that do not support those beliefs. Thus (HOB) alone is not enough to explain even the beliefs required by (BR), let alone the knowledge required by (KR). In order to get from beliefs about our reasons to beliefs about what follows from what we need an extra premiss. But that premiss is readily available. It is often remarked that, even though one may suspect that *some* of one’s beliefs are held for bad reasons, there is a difficulty with, for any particular belief, simultaneously believing that one has it for reason R and also failing to believe that R is a *good* reason for it. Thus, something along the lines of the following principle seems attractive (I call this principle (GR), for ‘good reasons’):

²⁰ One might wonder whether it is really (HOB) that follows, rather than the weaker (HOB*). This will depend on whether we take it that reasoning from R to p is the sort of thing that would trigger the disposition to believe that one believes that p . As before, I will not take sides on this.

(GR) If one believes that R is one's reason for believing p , then one believes that p follows from R ²¹

Now, although this principle seems attractive, it is unlikely to hold with full generality. Consider the case of Fred, who is superstitious and knows himself to be so. Fred recognizes that the fact that a black cat crossed his path is not a good reason for him to believe that he is going to have a bad day—and yet later on he finds himself unable to shake the conviction that he will have a bad day, while also acknowledging that his reason for that conviction is that a black cat crossed his path.²² Cases like Fred's, however, although perfectly possible, do not seem to be relevant here. This is because it seems clear that they are not cases of *reasoning*. It seems wrong to take Fred to be guilty of what is, by his own lights, a gross *non sequitur*. There surely are such things as compulsive patterns of belief, or patterns that one exemplifies even despite oneself. I suggest that such patterns do not count as instances of reasoning, precisely because of their imperviousness to one's recognition of their irrationality.²³ Thus, I believe, we can take (GR) to be true of cases of reasoning, which is what

²¹ I move here from a formulation in terms of what is a *good reason* for what to a formulation in terms of what *follows from* what. This is because, as I explained earlier (fn. 1), for the purposes of this paper I am setting aside cases of partial belief and merely probabilifying evidence.

²² The example is borrowed from Winters 1983, 210. Winters takes it to be a counterexample to (BR). This assessment seems wrong to me, for the reasons described in the text.

²³ To say as much is not to beg the question against opponents of (BR). It is one thing to claim that reasoning requires the belief that one's conclusion follows from one's premisses, and another to say that a pattern of belief cannot count as reasoning if it is entirely impervious to one's assessment of its rationality. Winters 1983 seems willing to deny even the latter claim, but I have not seen other opponents of (BR) take an explicit stand against it. I believe that Winters's rather extreme stance reflects a failure to acknowledge that, as I emphasized in sect. 2, reasoning is just one species of the phenomenon of *believing for a reason*; thus, denying that cases of compulsive belief are cases of reasoning does not entail denying that they are genuine cases of believing for a reason.

matters for present purposes. Crucially, now, by combining (HOB) and (GR) we get the conclusion that if one believes p by basic reasoning from R , then one believes that p follows from R . But this is just what (BR) requires.

The belief that p follows from R just arrived at is not the product of reasoning. One believes that p follows from R because one believes—through one’s capacity for non-observational, non-inferential higher-order belief—that one believes p for reason R , and one cannot believe that R is one’s reason for believing p without also believing that p follows from R . But it is not plausible to take this belief that p follows from R to be the product of further reasoning—for what could it have been inferred from? Clearly, it would be a gross mistake to infer that p follows from R from the fact that one’s reason for believing p is R . I suggest that we should simply accept that the beliefs about what follows from what involved in basic reasoning are based on no independent grounds.²⁴ Explicitly reasoning from R to p just is, in part, believing that p follows from R . But if the beliefs about what follows from what required by (BR) are not themselves the products of further reasoning, then (BR) does not threaten the status of such reasoning as basic.

Moreover, there is clearly no threat of regress in this account. There is no suggestion that the belief that p follows from R plays a causal role in one’s reasoning from R to p , and so no room to wonder about how it could possibly play that role. Thus, there is no obstacle to taking (BR) to hold in basic reasoning.

Now, if we take (BR) to hold in basic reasoning, I believe we should take (KR) to hold too. We can see the intuitive pressure to accept (KR) by returning to our example. Given what we have said so

²⁴ But then, what is the rational basis of such beliefs? Epistemological questions in this area are taken up below, and again in sect. 5.4.

far, Mary's believing (3) by reasoning from (1) and (2) entails her believing that (3) follows from (1) and (2). But, supposing that Mary *already knew* (1) and (2) independently of (3), it seems natural to take Mary *to know* (3) by reasoning from (1) and (2). If so, can we plausibly deny that her belief that (3) follows from (1) and (2) constitutes knowledge? This would seem very hard to reconcile with her really *knowing* (3) by reasoning from (1) and (2). The point is simply that, if we accept (BR), and we are prepared to attribute to a subject knowledge of some proposition by reasoning, it seems very odd to deny that her belief that her conclusion follows from her premisses constitutes knowledge. But this is exactly what (KR) requires.

This step, however, raises some worries of its own. As we saw, in the case of basic reasoning one's belief that one's conclusion follows from one's premisses is based on no independent grounds. So how can such beliefs ever be justified, and even count as knowledge, as (KR) requires?

First, notice that since in the cases of interest the subject by hypothesis is in a position to know p by reasoning from R , p really does follow from R ; so there is no worry about the subject's belief failing to be true in the relevant cases. The worry has to do with whether the relevant belief should count as *rational*, or *justified*.

This question is threatening only if we assume that the knowledge required by (KR) must be certifiable as knowledge *prior to* the epistemic status of the piece of reasoning in question. But this assumption can be doubted. In cases of basic reasoning, whether the corresponding beliefs about what follows from what count as knowledge, I suggest, simply depends upon the epistemic status of the inference in question. This should not seem surprising: after all, this belief is, in a sense, simply a 'by-product' believing p by basic reasoning from R . It has no independent grounds.

Now, it has sometimes been thought that, unless we have an independent way of knowing what follows from what, inferential knowledge is impossible (see, e.g., Fumerton 1995). Certainly, my account of basic reasoning does nothing to respond to sceptical worries of this sort. This is, I

think, not a serious problem for my view: there should be no presumption that an account of reasoning should, as such, provide a response to scepticism. Still, one might worry that my account gets things backwards, epistemologically speaking: isn't there a sense in which one's believing p by reasoning from R should count as knowledge in part *in virtue of* one's knowing that p follows from R ? Shouldn't beliefs about what follows from what pull some weight in an account of inferential justification, even in the case of basic reasoning? I respond to this worry in section 5.4.

5. Questions and Objections

The argument of the last section was long and complex, and so it may help to begin with a brief summary of the picture of reasoning I am defending. Start with a distinction between basic reasoning (i.e., reasoning that presupposes no other reasoning) and non-basic reasoning (i.e., reasoning that may depend on prior reasoning). Basic reasoning is the product of the interaction of two basic capacities: our capacity to conform our beliefs to general patterns and our capacity for non-observational, non-inferential higher-order belief. For example, Mary's logical competence means that faced with the premisses that this snake is a *vipera aspis*, and that if this snake is a *vipera aspis* then it is venomous, she will tend to also believe that the snake is indeed venomous.

Opponents of (BR) would simply stop here: what I have described so far is, according to them, a complete account of basic reasoning. But this, as I argued, is incorrect: since Mary is endowed with the capacity for non-observational, non-inferential higher-order belief, what we have said about her so far *constitutively entails* that she also believes—indeed, plausibly *knows*—that her conclusion follows from her premisses. For subjects equipped with our capacity for higher-order belief, reasoning from R to p just is, in part, believing that p follows from R . This is why (BR) and (KR) hold, even in the case of basic reasoning, without threat of Lewis Carroll-style regress.

Of course, sophisticated subjects often reason about what follows from what, too. This makes *non-basic* reasoning possible. Intuitively, in non-basic reasoning one comes to believe p as a result of determining that one has conclusive evidence for p . As I suggested, we should take this relation to be constitutive, not causal: believing p by non-basic reasoning simply consists in having determined that one has conclusive evidence for p . Therefore (BR), and consequently (KR) too, hold in the case of non-basic reasoning as well—and, once again, without threat of regress.

This, then, is the picture of reasoning that emerges from the argument of section 4. In the rest of this section I will attempt to clarify my account and to respond to possible objections.

5.1 Constitutive relations among beliefs

My account relies on the claim that there are constitutive relations among beliefs. But what does this talk of constitutive relations among beliefs come down to?

The details will depend on the metaphysics of belief, and in this paper I want to remain somewhat non-committal on this. What I hope to do here is show that my claim is compatible with various familiar accounts of belief.

Pamela Hieronymi (2009) has recently proposed that believing p consists in having settled the question whether p in the affirmative. This conception of belief is clearly very congenial to my claim: if one has determined that one has conclusive evidence for p , then the question whether p is settled for one.

Belief is most commonly thought of in functional or dispositional terms. For example, it might be held that, among other things, believing p entails a disposition to use p as a premiss in further reasoning. Then to say that one's believing p consists in believing R and that p follows from R means that in virtue of believing R and that p follows from R a rational and attentive subject is in a

state with a causal role which includes the causal role characteristic of the state of believing p —for example, it includes the disposition to use p as a premiss in further reasoning.

If belief also has a normative dimension, that could be accommodated in a similar way. For example, several philosophers argue that belief is governed by a ‘norm of truth’, in that it is a conceptual truth about belief that a belief is correct only if the proposition that is its content is true. But one’s beliefs in R and in the proposition that p follows from R can both be true only if p is true; thus in believing R and that p follows from R one is *already* subject to the norm that corresponds to believing p .²⁵

Note also that the idea that there may be constitutive relations among beliefs is entirely compatible with physicalism: the physical realization of a belief may well include the physical realization of another.

5.2 *Believing R and that p follows from R without believing p*

One might be tempted to object that my account of non-basic reasoning links believing that one’s conclusion follows from one’s premisses and believing one’s conclusion too tightly. It surely is possible that one might believe R and that p follows from R but *fail* to believe p . Doesn’t this show that believing p by non-basic reasoning from R cannot simply consist in believing R and that p follows from R ?

It does not. The claim that *sometimes* one’s belief in p might consist in believing R and that p follows from R does not entail that one cannot believe R and that p follows from R without

²⁵ For recent defences of the claim that belief is governed by a ‘norm of truth’, see Wedgwood 2002 and Shah 2003. Millar 2004 argues that believing p incurs a normative commitment to believe what follows from p in light of one’s other beliefs. Since everything that follows from p already follows from R and the proposition that p follows from R , in having the latter beliefs one already incurs the commitments entailed by the former.

believing p . Consider an analogy. Some shootings are killings. For example, Lee Harvey Oswald's killing of John Kennedy consists in his shooting John Kennedy—Oswald didn't have to do anything else, over and above shooting, to kill Kennedy.²⁶ This, however, does not entail either that all shootings are killings, or that Oswald himself might not have shot at Kennedy and missed. Similarly, my view entails only that *if* one believes p by non-basic reasoning from R , that belief consists in believing R and that p follows from R . It does not entail that every case of believing R and that p follows from R is a case of believing p , any more than the fact that some shootings are killings entails that all shootings are.

So how would my account handle such cases? Broadly speaking, we can distinguish between two types of case in which the constitutive relation seems to fail: one of them involves a failure of rationality, the other not. I will consider each in turn.

Suppose that Mary recognizes that she has conclusive evidence of her son's guilt, but—as she might put it—she is not able to bring herself to believe that he is guilty. She simply cannot think of him in this way. What should we say about such cases? *Prima facie*, at least, I can see no special reason not to apply my account to such a case. If Mary really believes that the evidence for her son's guilt is conclusive, we should also take her to believe that her son is guilty. Thus, the constitutive relation holds in this case; its failure is only apparent. The difficulty in this case lies elsewhere: it lies in the fact that Mary also seems to be in a conflicting state, perhaps the state of being *agnostic about* her son's guilt. This is why she claims—falsely, on the present view—that she does not believe that her son is guilty. Now, giving a detailed account of a conflicted psychology such as Mary's might be a difficult task, but this challenge is in no way special to my account. Incoherence and irrationality

²⁶ See Davidson 1971 for this point. Note that we can accept this point without accepting the further claim that the shooting is identical with the relevant movements of Oswald's body.

seem to be unfortunate facts about our mental lives, which we have to account for regardless of our account of reasoning.

We can, however, also find cases in which, without failure of rationality, a subject believes R and that p follows from R , while genuinely failing to believe p . Suppose Fred today figures out that p follows from R . Tomorrow, when this piece of information has receded from the forefront of his consciousness, he learns R . In such a case Fred might fail to come to believe p , simply because he fails to combine the two pieces of knowledge—because he fails to put two and two together, as we say. Such cases are cases of *inattention*, and they need not involve any irrationality. Beliefs that are not properly attended to may simply fail to combine in the way necessary for reasoning.

Now, giving a general account of the distinction between beliefs that are attended to in the relevant sense and beliefs that are not is a difficult issue, but fortunately we do not need such an account. My claim is just that once Fred has done *whatever it takes* to bring both beliefs into proper focus together, *then* those beliefs suffice for him to believe p . At that point, nothing *else* is required of him: at that point, Fred believes p by non-basic reasoning from R . Does the fact that I lack an independent account of what it is to attend to one's beliefs in the relevant sense threaten to trivialize my account? I do not believe so. The challenge was to give an account of the role of beliefs about what follows from what in reasoning that does not threaten regress. My response consists of two claims: first, that sometimes believing R and that p follows from R constitutes believing p , and second, that all and only cases of believing p that are so constituted are cases of believing p by non-basic reasoning from R . These are both substantive theses, even without an independent way of specifying when believing R and that p follows from R suffices for believing p .

5.3 *Is reasoning a causal process?*

What exactly are the implications of my account for the causal process view of reasoning? To begin with, notice that on my view, *non-basic* reasoning consists simply in determining that one's conclusion follows from one's premisses, and thereby believing one's conclusion. Thus, once all the evidence is in, and one has determined that one's conclusion follows from it, one's reasoning is *done*—no further process is needed. On the causal process view, by contrast, even after all the antecedent beliefs are in place, one is not done: one still needs to engage in a distinctive belief-forming process—the process of reasoning itself.

More generally, according to the causal process view of reasoning, believing *p* by reasoning is a matter of having a belief in *p* with the right causal ancestry. On my view, by contrast, this is not correct: believing *p* by reasoning from *R* constitutively entails that one believes that *p* follows from *R*. The crucial relation here is constitutive, not causal. Thus, although reasoning clearly may *involve* processes—such as collecting or examining evidence, or checking the validity of arguments—it is not to be *identified* with a causal process.

5.4 *Basic Reasoning and Epistemic Priority*

One distinctive feature of my account is that it denies that the beliefs about what follows from what involved in reasoning are always *epistemically prior* to the reasoning itself. In the case of basic reasoning, as I argued, the belief that *p* follows from *R* is constitutively entailed by believing *p* by reasoning from *R*; it has no independent grounds. Indeed, I suggested that it counts as knowledge (when it does) just because one's belief in *p* by reasoning from *R* counts as knowledge. This, however, might make my account of basic reasoning appear irrelevant to an important motivation for claims like (BR) and (KR)—namely, the idea that it is at least in part *in virtue of* recognizing that one's conclusion follows from one's premisses that one comes to know one's conclusion on the

basis of those premisses. Wouldn't that require that beliefs about what follows from what be epistemically prior to the corresponding piece of reasoning?

To address this question we need to be a bit more precise about what exactly the intuitive epistemological role of beliefs about what follows from what is supposed to be. On the one hand, there are familiar sceptical arguments which begin with the idea that, in order to have knowledge through reasoning at all, one must have *prior* knowledge that one's conclusion follows from one's premisses. It is clear that my account will not satisfy such a sceptic, or anyone who thinks that the sceptic succeeds in capturing our intuitions about the role of beliefs about what follows from what in reasoning. My account of basic reasoning *presupposes* that sometimes we know things by reasoning, and attempts to tease out the implications of this fact.

On the other hand, however, I believe that our intuitive sense that beliefs about what follow from what play an epistemological role in reasoning does not *have* to be spelled out in the sceptic's way. There are other ways to secure an epistemological role for such beliefs in reasoning, which do not require that they rest on independent grounds. As I have argued, if one believes *p* by reasoning from *R* one recognizes that one's reasons for believing *p* consist in *R* (setting aside, as before, the question whether we should prefer (HOB*) to (HOB)). But then, if one does not also believe that *p* follows from *R*, one would seem to be guilty of some kind of *incoherence*. One cannot coherently *and knowingly* go on believing *p* for reason *R* while disbelieving, or even merely being agnostic about, the proposition that *p* follows from *R*. In such a case, one's belief in *p* would not be properly grounded in one's belief in *R*, in an entirely straightforward sense: even if one's belief in *R* were perfectly

justified, its justification would not be simply reflected in the epistemic status of one's belief in p . The latter belief is epistemically problematic, in a way that the former is not.²⁷

Thus, on my view (BR) and (KR) should be understood as constraints of coherence. This secures an epistemological role in reasoning for beliefs about what follows from what, without requiring that such beliefs always rest on independent grounds. This, I suggest, should be enough to satisfy our intuitive motivations for (BR) and (KR).

6. Conclusion

In this paper I have argued that, contrary to contemporary orthodoxy, knowledge by reasoning requires knowledge of what follows from what. This is an interesting result in itself, but it also promises to have important consequences for our approach to other questions as well.

For example, much of the debate in the epistemology of basic logic has focused on so-called 'blind reasoning'—i.e., reasoning whose epistemic status is allegedly independent of its subject's recognizing its validity (Boghossian 2000, 2003). This approach is motivated in part by regress arguments. If the present reflections are on the right lines, however, this approach is baseless. There is no such thing as blind reasoning, and—given the failure of the regress argument—neither is there any need for it. On the contrary, the present reflections suggest that our most basic logical knowledge may be derived from the knowledge of particular entailments required by (KR) in cases of basic reasoning, and thus ultimately based on exercises of our capacity for non-inferential, non-

²⁷ We saw an example of this earlier, in superstitious Fred who recognizes that his reason for believing that he is likely to have a bad day is that a black cat crossed his path, without believing that black cats are evidence of bad luck. As I pointed out earlier, on my view (and on any view that accepts (BR)) such compulsive patterns of belief cannot count as reasoning: reasoning rules out certain types of incoherence.

observational self-knowledge. This may allow us to find an element of truth in the traditional connection between a priori knowledge and knowledge of our own minds.²⁸

MARKOS VALARIS

School of Humanities

University of New South Wales

Kensington 2052

NSW, Australia

m.valaris@unsw.edu.au

References

- Audi, Robert 1994: 'Dispositional Beliefs and Dispositions to Believe'. *Noûs*, 28, pp. 419-34.
- Binkley, Robert, Richard Bronaugh, and Ausonia Marras (eds.) 1970: *Agent, Action, and Reason*. Toronto: University of Toronto Press.
- Baldwin, Thomas (ed.) 1993: *G. E. Moore: Selected Writings*. London, Routledge.
- Boghossian, Paul 2000: 'Knowledge of Logic'. In Boghossian and Peacocke 200, 229-254.
- 2003: 'Blind Reasoning'. *Proceedings of the Aristotelian Society, Supplementary Volume*, 77, pp. 225-248.
- Boghossian, Paul and Christopher Peacocke (eds.) 2000: *New Essays on the A Priori*. Oxford: Oxford University Press.
- Brewer, Bill 1995: 'Mental Causation: Compulsion by Reason'. *Proceedings of the Aristotelian Society*
-

²⁸ Earlier versions of this paper were presented at ANU in Canberra and at the 2010 AAP conference at UNSW. I thank Annette Baier, David Chalmers, John Maier and John McDowell for helpful comments on those occasions. I am also very grateful to Melissa Merritt, for reading and commenting on several versions of the paper. But most of all I must acknowledge the help of the editor and referees for *Mind*, whose comments led to several significant revisions of this paper.

Supplementary Volumes, 69, pp. 237-253.

- Broome, John 2006: 'Reasoning with Preferences?'. In Olsaretti 2006, pp. 183-238.
- Burge, Tyler 1996: 'Our Entitlement to Self-Knowledge'. *Proceedings of the Aristotelian Society*, 96, pp. 91-116.
- Byrne, Alex 2005: 'Introspection'. *Philosophical Topics*, 33, pp. 79-104.
- 2011: 'Transparency, Belief, Intention.' *Proceedings of the Aristotelian Society Supplementary Volume*, 85, pp. 201-221.
- Carroll, Lewis 1895: 'What the Tortoise Said to Achilles'. *Mind*, 4, p. 278.
- Davidson, Donald 1971: 'Agency'. In his 1980, 43-62. Originally published in Binkley, Bronaugh, and Marras 1971.
- Edgley, Roy 1969: *Reason in Theory and Practice*. London: Hutchinson.
- Evans, Gareth 1982: *The Varieties of Reference*. New York: Oxford University Press.
- Fumerton, Richard 1995: *Metaepistemology and Skepticism*. Lanham, MD: Rowman & Littlefield.
- Gallois, André 1996: *The World Without, the Mind Within*. Cambridge: Cambridge University Press.
- Hieronymi, Pamela 2006: 'Controlling Attitudes'. *Pacific Philosophical Quarterly*, 87, pp. 45-74.
- 2009: 'Two Kinds of Agency'. In O'Brien and Soteriou 2009, pp. 138-162.
- Johnston, Mark 1988: 'Self-Deception and the Nature of Mind'. In McLaughlin and Rorty 1988, pp. 63-91.
- McLaughlin, Brian and Amelie Rorty (eds.) 1988: *Perspectives on Self-Deception*. Berkeley and Los Angeles: University of California Press.
- Millar, Alan 2004: *Understanding People: Normativity and Rationalizing Explanation*. Oxford: Oxford University Press.
- Moore, G. E. 1944: 'Moore's Paradox'. In Baldwin 1993, pp. 207-212. Previously unpublished manuscript.
- Moran, Richard 2001: *Authority and Estrangement*. Princeton: Princeton University Press.
- O'Brien, Lucy and Matthew Soteriou (eds.) 2009: *Mental Actions* Oxford: Oxford University Press.
- Olsaretti, Serena (ed.) 2006: *Preferences and Well Being*. Cambridge: Cambridge University Press.

- Railton, Peter 2006: 'How to Engage Reason: the Problem of Regress'. In Wallace, Pettit, Scheffler, and Smith 2006, pp. 176-202.
- Rödl, Sebastian 2007: *Self-Consciousness*. Cambridge, MA: Harvard University Press.
- Shah, Nishi 2003: 'How Truth Governs Belief'. *The Philosophical Review*, 112, pp. 447-482.
- Shoemaker, Sydney. 1988. 'On Knowing One's Own Mind'. *Philosophical Perspectives*, 2, pp. 183-209.
- Wallace, Ray, Philip Pettit, Stephen Scheffler, and Michael Smith (eds) 2006. Oxford: *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*. Oxford: Oxford University Press.
- Wedgwood, Ralph 2002: 'The Aim of Belief'. *Philosophical Perspectives*, 16, pp. 267-97.
——— 2006: 'The Normative Force of Reasoning'. *Noûs*, 40, pp. 660-86.
- Williams, John 1979: 'Moore's Paradox-One or Two?'. *Analysis*, 39, pp. 141-2.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Winters, Barbara 1983: 'Inferring'. *Philosophical Studies*, 44, pp. 201-20.