

Inference, Agency, and Responsibility

Paul Boghossian

1. Introduction

What happens when we *reason* our way from one proposition to another? This process is usually called “inference” and I shall be interested in its nature.¹

There has been a strong tendency among philosophers to say that this psychological process, while perhaps real, is of no great interest to epistemology. As one prominent philosopher (who shall remain nameless) put it to me (in conversation): “The process you are calling ‘inference’ belongs to what Reichenbach called the ‘context of discovery’; it does not belong to the ‘context of justification,’ which is all that really matters to epistemology.”²

I believe this view to be misguided. I believe there is no avoiding giving a central role to the psychological process of inference in epistemology, if we are to adequately explain the responsibility that we have, qua epistemic agents, for the rational management of our beliefs. I will try to explain how this works in this chapter.

In the course of doing so, I will trace some unexpected connections between our topic and the distinction between a priori and a posteriori justification, and I will

Earlier versions of some of the material in this chapter were presented at a Workshop on Inference at the CSMN in Oslo, at the Conference on Reasoning at the University of Konstanz, both in 2014, at Susanna Siegel’s Seminar at Harvard in February 2015, and at a Workshop on Inference and Logic at the Institute of Philosophy at the University of London in April 2015. A later version was given as a keynote address at the Midwest Epistemology Workshop in Madison, WI in September 2016. I am very grateful to the members of the audiences at those various occasions, and to Mark Richard, Susanna Siegel, David J. Barnett, Magdalena Balcerak Jackson, and an anonymous referee for OUP for detailed written comments.

¹ Importantly, then, I am not here talking about inference as *argument*: that is as a set of propositions, with some designated as “premises” and one designated as the “conclusion.” I am talking about inference as reasoning, as the psychological transition from one (for example) belief to another. Nor am I, in the first instance, talking about *justified* inference. I am interested in the nature of inference, even when it is unjustified.

² For another example of a philosopher who downplays the importance of the psychological process of inference see Audi (1986).

draw some general methodological morals about the role of phenomenology in the philosophy of mind.

In addition, I will revisit my earlier attempts to explain the nature of the process of inference (Boghossian 2014, 2016) and further clarify why we need the type of “intellectualist” account of that process that I have been pursuing.

2. Beliefs and Occurrent Judgments

I will begin by looking at beliefs. As we all know, many of our beliefs are *standing beliefs* in the sense that they are not items in occurrent consciousness, but reside in the background, ready to be manifest when the appropriate occasion arises.

Among these standing beliefs, some philosophers distinguish between those that are *explicit* and those that are *implicit*.

Your explicit beliefs *originated* in some *occurrent judgment* of yours: at some point in the past, you occurrently affirmed or judged them (or, more precisely, you occurrently affirmed or judged their subtended propositions). At that point, those judgments went into *memory storage*.³

In addition to these occurrent-judgment-originating beliefs, some philosophers maintain that there are also *implicit* standing beliefs, beliefs that did not originate in any occurrent judgment, but which you may still be said to believe. For example, some philosophers think that you already believed that Bach didn’t write down his compositions on watermelon rinds, prior to having encountered this proposition here (as we may presume) for the first time. Others dispute this. I won’t take a stand on the existence of implicit beliefs for the purposes of this chapter.

Let me focus instead on your *explicit* standing beliefs. Many of these beliefs of yours (we may also presume) are *justified*. There is an epistemic *basis* on which they are *maintained*, and it is in virtue of the fact that they have that basis that you are justified in having them, and justified in relying on them in coming to have other beliefs. In Pryor’s terminology (2005: 182), bases are justifiers.

3. Epistemic Bases for Standing Beliefs

What are the bases or justifiers for these explicit standing beliefs? There are (at least) two distinct questions here.

First, what *sorts of things* are such bases?

Second, how do these bases behave over time? More specifically, once a basis for a belief is established, does it retain that basis (putting aside cases where the matter is explicitly reconsidered), or can that basis somehow shift over time, even without reconsideration?

³ Later on, I shall argue that this should be regarded as *preservative memory storage*, in the sense of Burge (1993) and Huemer (1999).

Let me start with the first question.

Many epistemologists have become increasingly sympathetic to the view that justifiers are *exclusively propositions*, rather than *mental states* (albeit, states with propositional content). I want to buck this particular trend.

In the case of a perceptual belief that *p*, I think, like Pryor (2000), that it can be the case that the thing in virtue of which your belief that there is a cat on the mat in this room is justified is that you have a *visual state* as of seeing that there is a cat on the mat in this room, and that you base your belief on that state. Call this the *Statist View* (see Pryor 2007).

On the alternative view, the justifier for your belief that there is a cat on the mat in this room is not your visual state as of seeing that there is a cat on the mat in this room, but, rather, just the proposition itself that there is a cat on the mat in this room. Call this the *Propositional View* (see, for example, Williamson 2000 and Littlejohn 2012).

Now, of course, even on a Propositional View, there must be something about you in virtue of which the proposition that there is a cat on the mat can serve as a justification for *your* belief that there is a cat on the mat, even while it does not serve as such a justification for me who, not having seen the cat, have no such justification. What is that explanation, given that the proposition that there is cat on the mat is, in some sense, available to both of us?

The answer presumably is that, while the proposition as such is available to both of us, your visual experience makes it available to you *as evidence*, while it is not in this way available to me. Your visual experience, but not mine, gives you *access* to the relevant proposition as evidence on which you base your beliefs.

This is no doubt a subtle difference. Still, there seems to be something importantly different between the two views.

On the first view, mental states themselves play an epistemic role, they do some justifying. On the alternative view, the mental states don't do any justifying themselves; their role is to give you *access* to the things—the propositions—that do the justifying.

I wasn't inclined to worry much about this distinction between alternative views of justifiers until I came to think about its interaction with the *a priori/a posteriori* distinction. I am now inclined to think that the *a priori/a posteriori* distinction requires rejecting the view that *only* propositions can serve as justifiers.

Let me try to spell this out.

The tension between the Propositional View and the *a priori/a posteriori* distinction might be thought to emerge fairly immediately from the way in which the distinction between the *a priori* and the *a posteriori* is typically drawn. We say that a belief is *a priori* justified just in case its justification does not rely on perceptual experience. This seems to presuppose that perceptual experiences have justifying roles.

But this is, of course, far from decisive. As we have already seen, even on the Propositional View, experiences will need to be invoked in order to explain why some propositions are available to a thinker as evidence and others aren't.

Given that fact, couldn't the Propositionalist make sense of the a priori/a posteriori distinction by saying that perceptual experiences give one access to a posteriori evidence, while non-perceptual experiences, such as intuition or understanding, give one access to a priori evidence?

The problem with this reply on behalf of the Propositionalist is that we know that *sometimes* perceptual experience will be needed to give one access to a proposition that is then believed on a priori grounds.

For example, perceptual experience may be needed to give one access to the proposition "If it's sunny, then it's sunny." But once we have access to that proposition, we may come to justifiably believe it on a priori grounds (for example, via the understanding).

In consequence, implementing the response given on behalf of the Propositionalist will require us to distinguish between those uses of perceptual experience that merely give us access to thinking the proposition in question, versus those that give us access to it *as evidence*.

But how are we to draw this distinction without invoking the classic distinction between a merely *enabling* use of perceptual experience, and an *epistemic* use of such experience, a distinction that appears to presuppose the Statist View. For what would it be to make a proposition accessible as evidence, if not to experience it in a way that justifies belief in it?

To sum up. Taking the a priori/a posteriori distinction seriously requires thinking of mental states as sometimes playing a justificatory role; it appears not to be consistent with the view that it is *only* propositions that do any justifying.⁴

I don't now say that the Statist View is *comprehensively* correct, that all reasons for judgment are always mental states, and never propositions. I only insist that, if we are to take the notion of a priori justification seriously, mental states, and, in particular, experiences, must sometimes be able to play a justificatory role.

4. Epistemic Bases for Standing Beliefs: Then and Now

Going forward, and since it won't matter for present purposes, I will assume that the Statist View is *comprehensively* correct. However, the reader should bear in mind that this is done merely for ease of exposition and because the issues to be addressed here don't depend on that assumption.

⁴ Once we have successfully defined a distinction between distinct ways of believing a proposition, we can introduce a derivative distinction between types of proposition: if a proposition can be reasonably believed in an a priori way, we can say that it is an a priori proposition; and if it can't be reasonably believed in an a priori way, but can only be reasonably believed in an a posteriori way, then we can say that it is an a posteriori proposition. But this distinction between types of proposition would be dependent upon, and derive from, the prior distinction between distinct ways of reasonably believing a proposition, a distinction which depends on construing epistemic bases as mental states, rather than propositions.

Here is a further question. Assuming that we are talking about an explicit standing belief, does that belief always have as its basis the basis on which it was *originally* formed; or could the basis have somehow shifted over time, even if the belief in question is never reconsidered?

It is natural to think that the first option is correct, that explicit standing beliefs have *whatever bases they had* when they were formed as occurrent judgments. At their point of origin, those judgments, having been arrived at on a certain basis, and not having been reconsidered or rejected, go into memory storage, *along* with their epistemic bases.

Is the natural answer correct? It might seem that it is not obviously correct. For sometimes, perhaps often, you *forget* on what basis you formed the original occurrent judgment that gave you this standing belief. If you now can't recall what that basis was, is it still the case that your belief *has* that basis? Does it carry that basis around with it, whether you can recall it or not?

I want to argue that the answer to this question is "yes." I think this answer is intuitive. But once again there is an unexpected interaction with the topic of the *a priori*. Any friend of the *a priori* should believe that the natural answer is correct. Let me say a little more about this interaction.

5. Proof and Memory

Think here about the classic problem about how a lengthy proof might be able to deliver *a priori* knowledge of its conclusion. In the case of some lengthy proofs, it is not possible for creatures like us to carry the whole thing out in our minds. Unable to keep all the steps of the proof in mind, we need to write them down and look them over. In such cases, memory of the earlier steps in the proof, and perceptual experience of what has been written down, enter into the full explanation of how we arrive at our belief in the conclusion of the proof.

And these facts raise a puzzle that has long worried theorists of the *a priori*: How could belief arrived at on the basis of this sort of lengthy proof be *a priori* warranted? Won't the essential role of memory and perceptual experience in the process of carrying out the proof undermine the conclusion's alleged *a priori* status?⁵

To acquiesce in a positive answer to this question would be counterintuitive. By intuitive standards, lengthy proof can deliver *a priori* warrant for its conclusion just as well as a short proof can. The theoretical puzzle is to explain how it can do so, given the psychologically necessary role of perceptual and memory experiences in the process of proof.

The natural way to respond to this puzzle is to appeal to an *expanded* conception of the distinction between an enabling and an epistemic use of experience. Tyler

⁵ See, for example, Chisholm (1989).

Burge has developed just such a view (I shall look at his application of it to memory). Burge distinguishes between a merely enabling use of memory, which Burge calls “preservative memory,” and an epistemic use, which Burge calls “substantive memory”:

[Preservative] memory does not supply for the demonstration propositions about memory, the reasoner, or past events. It supplies the propositions that serve as links in the demonstration itself. Or rather, it *preserves* them, together with their judgmental force, and makes them available for use at later times. Normally, the content of the knowledge of a longer demonstration is no more about memory, the reasoner, or contingent events than that of a shorter demonstration. One does not justify the demonstration by appeals to memory. One justifies it by appeals to the steps and the inferential transitions of the demonstration.... In a deduction, reasoning processes’ working properly depends on memory’s preserving the results of previous reasoning. But memory’s preserving such results does not add to the justificational force of the reasoning. It is rather a background condition for the reasoning’s success.⁶

Given this distinction, we can say that the reason why a long proof is able to provide a priori justification for its conclusion is that the only use of memory that is essential in a long proof is preservative memory, rather than substantive memory. If substantive memory of the act of writing down a proposition were required to arrive at justified belief in the conclusion of the proof, that might well compromise the conclusion’s a priori status. But it is plausible that in most of the relevant cases all that’s needed is preservative memory.

Now, against the background of the Propositional View of justifiers, all that would need to be preserved is, as Burge says, just the propositions (along perhaps with their judgmental force).

However, against the background of the Statist View, we would need to think that preservative memory can preserve a proposition not only with its judgmental force, but *also* along with its mental state justifier.⁷

Now, just like an earlier step in a proof can be invoked later on without this requiring the use of substantive memory, so, too, can an a priori justified standing belief be invoked later on without this compromising its ability to deliver a priori justification. To account for this, we must think that when an occurrent judgment goes into memory storage, it goes into *preservative* memory storage.

⁶ Burge (1993, pp. 462–3).

⁷ A competing hypothesis is that preservative memory need only retain the proposition, along with its judgmental force and *epistemic status*, but without needing to preserve its mental state justifier. The problem with this competing hypothesis is that we will want to retain the status of the proposition as either a priori or a posteriori justified. However, a subject can be a priori justified in believing a given proposition, without himself having the concept of a priori justification. If we are to ensure that its status as a priori is preserved, without making undue conceptual demands on him, we must require that the mental state justifier that determines it as a priori justified is preserved. I am grateful to the anonymous referee for pressing me on this point.

Thus, it follows that a standing belief will have whatever basis it originally had, whether or not one recalls that basis later on. And so, we arrive at the conclusion we were after: a belief's original basis is the basis on which it is maintained, unless the matter is explicitly reconsidered.

Of course, both of these claims about bases are premised on the importance of preserving a robust use for the *a priori/a posteriori* distinction. But as I've argued elsewhere (Boghossian forthcoming), we have every reason to accord that distinction the importance it has traditionally had.

6. The Basis for Occurrent Judgments

We come, then, to the question: What are the bases for those original occurrent judgments, the ones that become our standing beliefs by being frozen in preservative memory? On what bases do we tend to arrive at particular occurrent judgments?

Well, here, as we are prone to say, the bases may be either *inferential* or *non-inferential*.

When a judgment's basis is non-inferential, it will typically consist in a perceptual state, such as a visual or auditory state. It may also consist in some state of introspection. Some philosophers allow that non-inferential bases may also consist in such experiential states as intuitions, and such non-experiential states as the understanding of concepts. For our purposes, here, though, we may leave such controversies aside.

I will be interested, instead, in what it is for an occurrent judgment to have an *inferential* basis.

Given our working assumption that the Statist View is comprehensively correct, the inferential basis for an occurrent judgment will always be some other occurrent judgment.⁸ The question is: How does one judgment get established as the basis for another judgment?⁹

When we ask what the justification for a judgment is, we need to be asking what that justification is for a particular person who makes the judgment. We are asking what *your* basis is for making the judgment, what *your* reason is.

And so, naturally, there looks to be no evading the fact that, at the end of the day, it will be some sort of *psychological* fact about you that establishes whether something

⁸ On the Propositional View, the basis would always be a proposition that is the object of some occurrent judgment. As I say, this particular distinction won't matter for present purposes.

⁹ Of course, judgments are not the only sorts of propositional attitude that inference can relate. One can infer *from* suppositions and from imperatives (for example) and one can infer *to* those propositional attitudes as well. Let us call this broader class of propositional attitudes that may be related by inference, *acceptances* (see Wright 2014 and Boghossian 2014). Some philosophers further claim that even such non-attitudinal states as *perceptions* could equally be the conclusions of inferences (Siegel 2017). In the sense intended (more on this below), this is a surprising claim and it is to be hoped that getting clearer on the nature of inference will eventually help us adjudicate on it. For the moment, I will restrict my attention to judgments.

is a basis for you, whether it is *the reason for which* you came to make a certain judgment.

Thus, it is some sort of psychological fact about you that establishes that this perception of yours serves as your basis for believing this observable fact. And it is some sort of psychological fact about you that establishes that it is these judgments that serve as your basis for making this new judgment, via an inference from those other judgments.

I am interested in the nature of this process of arriving at an occurrent judgment that *q* by inference from the judgment that *p*, a process which establishes *p* as your reason for judging *q*.

7. Types of Inference

What is an example of the sort of process I am talking about?

One example, that I will call, for reasons that I will explain later, reasoning 2.0, would go like this:

- (1) I consider explicitly some proposition that I believe, for example *p*.

And I ask myself explicitly:

- (2) (Meta) What follows from *p*?

And then it strikes me that *q* follows from *p*. Hence,

- (3) (Taking) I take it that *q* follows from *p*.

At this point I ask myself

- (4) Is *q* plausible? Is it less plausible than the negation of *p*?
- (5) I conclude that *q* is not less plausible than not-*p*.
- (6) So, I judge *q*.

I add *q* to my stock of beliefs.

Reasoning 2.0 is admittedly not the most common type of reasoning. But it is probably not as rare as it is fashionable to claim nowadays. In philosophy, as in other disciplines, there is a tendency to overlearn a good lesson. Wittgenstein liked to emphasize that many philosophical theories overly intellectualize cognitive phenomena. Perhaps so. But we should not forget that there are many phenomena that call for precisely such intellectualized descriptions.

Reasoning 2.0 happens in a wide variety of contexts. Some of these, to be sure, are rarified intellectual contexts, as, for example, when you are working out a proof, or formulating an argument in a paper.

But it also happens in a host of other cases that are much more mundane. Prior to the 2016 presidential election, the conventional wisdom was that Donald Trump was extremely unlikely to win it. Prior to that election, many people are likely to have

taken a critical stance on this conventional wisdom, asking themselves: “Is there really evidence that shows what conventional wisdom believes?” Anyone taking such a stance would be engaging in reasoning 2.0.

Having said that, it does seem true that there are many cases where our reasoning does not proceed via an *explicit* meta-question about what follows from other things we believe. Most cases of inference are seemingly much more automatic and unreflective than that. Here is one:

On waking up one morning you recall that

(Rain Inference)

(1) It rained heavily through the night.

You conclude that

(2) The streets are filled with puddles (and so you should wear your boots rather than your sandals).

Here, the premise and conclusion are both things of which you are aware. But, it would seem, there is no explicit meta-question that prompts the conclusion. Rather, the conclusion comes seemingly immediately and automatically. I will call this an example of *reasoning 1.5*.

The allusion here, of course, is to the increasingly influential distinction between two kinds of reasoning, dubbed “System 1” and “System 2” by Daniel Kahneman. As Kahneman (2011, pp. 20–1) characterizes them,

System 1 operates automatically and quickly, with little or no effort and no sense of voluntary control.

System 2 allocates attention to the effortful mental activities that demand it, including complex computations. The operations of System 2 are often associated with the subjective experience of agency, choice, and concentration.

As examples of System 1 thinking, Kahneman gives detecting that one object is more distant than another, orienting to the source of a sudden sound, and responding to a thought experiment with an intuitive verdict. Examples of System 2 thinking are searching memory to identify a surprising sound, monitoring your behavior in a social setting, and checking the validity of a complex logical argument.

Kahneman intends this not just as a distinction between distinct kinds of reasoning, but of thinking more broadly.

Applied to the case of reasoning, it seems to me to entail that a lot of reasoning falls somewhere in-between these two extremes.

The (Rain) inference, for example, is not effortful or attention-hogging. On the other hand, it seems wrong to say that it is not under my voluntary control, or that there is no sense of agency associated with it. It still seems to be something that I *do*. That is why I have labeled it “System 1.5 reasoning.”

The main difference between reasoning 2.0 and reasoning 1.5 is not agency *per se*, but rather the fact that in reasoning 2.0, but not in 1.5, there is an *explicit* (Meta)

question, an explicit state of *taking* the conclusion to follow from the premises, and, finally, an explicit *drawing* of the conclusion as a result of that taking. All three of these important elements seem to be present in reasoning 2.0, but missing from reasoning 1.5.

8. Inference versus Association

Now, one of the main claims that I have made in previous work on this topic is that we need to acknowledge a state of *taking* even in the 1.5 cases, even in those cases where, since the reasoning seems so immediate and unreflective, a taking state appears to be absent.

This way of thinking about inferring, as I've previously noted, echoes a remark of Frege's (1979, p. 3):

To make a judgment because we are cognisant of other truths as providing a justification for it is known as inferring.

My own preferred version of Frege's view, for reasons that, since I have explained elsewhere, I won't rehearse here, I would put like this:

(Inferring) S's inferring from p to q is for S to judge q *because S takes* (the accepted truth of) p to provide (contextual) support for (the acceptance of) q.

Let us call this insistence that an account of inference must in this way incorporate a notion of "taking" the Taking Condition on inference:

(Taking Condition): Inferring from p to q necessarily involves the thinker *taking* p to support q and drawing q *because* of that fact.

As so formulated, this is not so much a view, as a schema for a view. It signals the need for something to play the role of "taking," but without saying exactly what it is that plays that role, nor how it plays it.

In other work, I have tried to say more about this taking state and how it might play this role (see Boghossian 2014, 2016).

In an important sense, however, it was probably premature to attempt to do that. Before trying to explain the nature of the type of state in question, we need to satisfy ourselves that a correct account of inference has to assume the shape that the (Taking) condition describes. Why should we impose a phenomenologically counter-intuitive and theoretically treacherous-looking condition such as (Taking) on any adequate account of inference?

In this chapter, then, instead of digging into the nitty-gritty details of the taking state, I want to explain in fairly general terms why an account of inference should assume this specific shape, what is at stake in the debate about (Taking), and why it should be resolved one way rather than another.

9. What is at Stake in this Debate?

You might think it's quite easy to say what's at stake in this debate. After all, making claims that involve the notion of inference are common and central to philosophy.

Many philosophers, for example, think that it's an important question about a belief whether its basis is inferential or not. If it's inferential, then its justification is further dependent on that of the belief on which it rests; if it's not, the chain of justification might have reached its end.

To give another example, Siegel (2017) claims that the *conclusion* of an inference need not be an acceptance but might itself be a *perceptual* state. This is linked to the question whether one may be faulted for having certain perceptions, even if, in some sense, one can't help but have the perceptions one has.

Both of these claims depend on our being able to say what inference is and how to recognize it when it occurs.

There is reason, then, to think that the question of the nature of inference is important. The question, though, is to say how to distinguish this substantive dispute about the nature of inference, from a *verbal* dispute about how the word "inference" is or ought to be applied.

Here, as elsewhere in philosophy, the best way to get at a substantive dispute about the nature of a concept or phenomenon is to specify *what work* one needs that concept or phenomenon to do. What is expected of it?

An initial thought about what work we want the concept of reasoning to do is that we need it to help us distinguish reasoning from the *mere association of thoughts* with one another. In giving an account of inferring a *q* from a *p*, we need to distinguish that from a case in which *p* merely gives rise to the judgment that *q* in some associative way.

That's fine as far as it goes, but it now invites the question: Why does it matter for us to capture the distinction between association and inference?

The start of an answer is that inferring from the judgment that *p* to the judgment that *q* *establishes p as the epistemic basis for judging q*, whereas associating *q* with *p* does not.

The fuller answer is that we want to give an account of this idea of a thinker establishing *p* as a basis for believing *q* that is subject to a certain constraint: namely, that it be intelligible how thinkers can be held *responsible* for the quality of their reasoning.

I can be held responsible for the way I reason, but not for what associations occur to me. I can be held responsible for what I establish as a good reason for believing something, but not for what thoughts are prompted in me by other thoughts.

These, then, are some of the substantive issues that I take to be at stake in the debate I wish to engage in, and they are issues of a normative nature. In genuine reasoning, you establish one judgment as your basis for making another; and you can be held responsible for whether you did that well or not.

All of this, of course, is part and parcel of the larger debate in epistemology about the extent to which foundational notions in epistemology are normative and deontological in nature. The hope is that, by playing out this debate in the special case of reasoning, we will shed light both on the nature of reasoning and on the larger debate.

With these points in mind, let's turn to our question about the nature of inference and to whether it should in general be thought of as involving a "taking" state.

10. Inference Requires Taking

One point on which everyone is agreed is that for you to infer from *p* to *q* your acceptance of *p* must *cause* your acceptance of *q*.

Another point on which everyone is agreed is that while such causation may be necessary, it is not sufficient, since it would fail to distinguish inference from mere association. The question is: What else should we add to the transition from *p* to *q* for it to count as inferring *q* from *p*?

Hilary Kornblith has claimed that, so long as we insist that the transitions between *p* and *q* "[involve] the interaction among representational states on the basis of their content," then we will have moved from mere causal transitions to full-blooded cases of reasoning.

Thus, he says:

But, as I see it, there is now every reason to regard these informational interactions as cases of reasoning: they are, after all, transitions involving the interaction among representational states on the basis of their content. (2012, p. 55)

But this can't be right. Mere *associations* could involve the interaction of representational states on the basis of their content. The Woody Allenesque depressive who, on thinking "I am having so much fun" always then thinks "But there is so much suffering in the world," is having an association of judgments on the basis of their content. But he is not thereby inferring from the one proposition to the other. (For more discussion of Kornblith, see Boghossian 2016.)

If mere sensitivity to content isn't enough to distinguish reasoning from association, then perhaps what's missing is *support*: the depressive's thinking doesn't count as reasoning because his first judgment—that he is having so much fun—doesn't *support* his second—that there is so much suffering in the world. Indeed, the first judgment might be thought to slightly undermine the second, since at least the depressive is having fun. By contrast, in the (Rain) inference the premise does support the conclusion.

But, of course, that can't be a good proposal either. Sometimes I reason from a *p* to a *q* where *p* does not support *q*. That makes the reasoning *bad*, but it is reasoning nonetheless. Indeed, it is precisely because it is reasoning that we can say it's bad. The very same transition would be just fine, or at any rate, acceptable, if it were a mere association.

What, then, should we say?

At this point in the dialectic, something like a taking-based account seems not only natural, but forced: the depressive's thinking doesn't count as reasoning *not* because his first judgment doesn't *support* his second, but, rather, it would seem, because he doesn't *take* his first judgment to support his second. The first judgment simply *causes* the second one in him; he doesn't draw the second one because he takes it to be supported by the first.

On the other hand, in the (Rain) inference, it doesn't *prima facie* strain credulity to say that I take its having rained to support the claim that the streets would be wet and that that is why I came to believe that the streets will be wet.

At least *prima facie*, then, there looks to be a good case for the Taking Condition: it seems to distinguish correctly between mere association and inferring. And there doesn't seem any other obvious way to capture that crucial distinction.

11. Further Support for the Taking Condition: Responsibility and Control

These is, however, a lot more to be said in favor of the Taking Condition, in addition to these initial considerations. I will review some of the more central considerations in this section.

To begin with, let's note that even reasoning 1.5, even reasoning that happens effortlessly and seemingly automatically, could intuitively be said to be something that you *do*—a mental action of yours—rather than simply something that happens to you.

And this fact is crucially connected to the fact that we can (a) not only assess whether you reasoned well, but (b) hold you responsible for whether you reasoned well, and allow that assessment to enter into an assessment of your rationality. For on the picture on offer, you take your premises to support your conclusion and actively draw your conclusion as a result.

I don't now assert that the Taking Condition is the only way to make your reasoning count as agential. I do assert that it is one clear way of doing so and that no other ways appear obviously available.

As a result, the picture on offer satisfies one of the principal desiderata that I outlined: that your reasoning be a process for which you could intelligibly be held rationally *responsible*.

Second, talk of taking fits in well with the way in which, with person-level inference, it is always appropriate to precede the drawing of the conclusion with a "so" or a "therefore." What are those words supposed to signify if not that the agent is taking it that her conclusion is justified by her premises?¹⁰

¹⁰ Pettit made this observation in his (2007, p. 500).

Third, taking appears to account well for how inference could be subject to a Moore-style paradox. That inference is subject to such a style of paradox has been well described by Ulf Hlöbil (although he disputes that my explanation of it is adequate). As Hlöbil (2014, pp. 420–1) puts it:

(IMP) It is either impossible or seriously irrational to infer P from Q and to judge, at the same time, that the inference from Q to P is not a good inference.

... [It] would be very odd for someone to assert (without a change in context) an instance of the following schema,

(IMA) Q; therefore, P. But the inference from Q to P is not a good inference (in my context).

... it seems puzzling that if someone asserts an instance of (IMA), this seems self-defeating. The speaker seems to contradict herself... Such a person is irrational in the sense that her state of mind seems self-defeating or incoherent. However, we typically don't think of inferences as contentful acts or attitudes... Thus, the question arises how an inferring can generate the kind of irrationality exhibited by someone who asserts an instance of (IMA). Or, to put it differently: How can a doing that seems to have no content be in rational tension with a judgment or a belief?

Hlöbil is right that there is a *prima facie* mystery here: how could a *doing* be in rational tension with a *judgment*?

The Taking Condition, however, seems to supply a good answer: there can be a tension between the doing and the judgment because the doing is the result of taking the premises to provide good support for the conclusion, a taking that the judgment then denies.

Fourth, taking offers a neat explanation of how there could be two kinds of inference—deductive and inductive.¹¹

Of course, in some inferences the premises logically entail the conclusion and in others they merely make the conclusion more probable than it might otherwise be. That means that there are two sets of standards that we can apply to any given inference. But that only gives us two standards that we can apply to an inference, not two different kinds of inference.

Intuitively, though, it's not only that there are two standards that we can apply to any inference, but two different types of inference. And, intuitively once more, that distinction involves taking: it is a distinction between an inference in which the thinker takes his premises to deductively warrant his conclusion versus one in which he takes them merely to inductively warrant it.

Finally, some inferences seem not only obviously unjustified, and so not ones that rational people would perform; more strongly, they seem *impossible*. Even if you were

¹¹ The following two points were already mentioned in Boghossian (2014); I mention them here again for the sake of completeness.

willing to run the risk of irrationality, they don't seem like inferences that one *could* perform.

Consider someone who claims to infer Fermat's Last Theorem (FLT) directly from the Peano axioms, without the benefit of any intervening deductions, or knowledge of Andrew Wiles's proof of that theorem. No doubt such a person would be unjustified in performing such an inference, if he could somehow get himself to perform it.

But more than that, we feel that no such transition could be an *inference* to begin with, at least for creatures like ourselves. What could explain this?

The Taking Condition provides an answer. For the transition from the Peano axioms to FLT to be a real inference, the thinker would have to be taking it that the Peano axioms support FLT's being true. And no ordinary person could so take it, at least not in a way that's unmediated by the proof of FLT from the Peano axioms. (The qualification is there to make room for extraordinary people, like Ramanujan, for whom many more number-theoretic propositions were obvious than they are for the rest of us.)¹²

We, see, then, that there are a large number of considerations, both intuitive and theoretical, for imposing a Taking Condition on inference.

12. Helmholtz and Sub-personal Inference

But what about the fact that the word 'inference' is used, for example in psychology and cognitive science, to stand for processes that have nothing to do with taking?

Helmholtz is said to have started the trend by talking about unconscious and sub-personal inferences that are employed by our visual system (see his 1867); but, by now, the trend is a ubiquitous one. Who are we, armchair philosophers, to say that inference must involve taking, if the science of the mind (psychology) happily assumes otherwise?

Should we regard Helmholtz's use of "inference" as a source of potential counter-examples to our taking-based accounts of inference, or is Helmholtz simply using a different (though possibly related) concept of "inference," one that indifferently covers both non-inferential sub-personal transitions that may be found in simple creatures, as well as our examples of 2.0 and 1.5 cases of reasoning?

How should we decide this question? How should we settle whether this is a merely terminological dispute or a substantive one?

¹² I'm inclined to think that this sort of example actually shows something stronger than that taking must be involved in inference. I'm inclined to think that it shows that the taking must be backed by an *intuition* or insight that the premises support the conclusion. For the fact that an ordinary person can't take it that FLT follows from the Peano axioms directly isn't a brute fact. It is presumably to be explained by the fact that an ordinary person can't simply "see" that the FLT follows from the Peano axioms. I hope to develop this point further elsewhere.

We need to ask the following: What would we miss if we only had Helmholtz's use of the word "inference" to work with? What important seam in epistemology would be obscured if we only had his word?

I have already tried to indicate what substantive issue is at stake. I don't mind if you use the same word "inference" to cover both Helmholtz's sub-personal transitions and adult human reasoning 2.0 and 1.5. The crucial point is not to let that linguistic decision obscure the normative landscape: unlike the latter, the sub-personal transitions of a person's visual system are not ones that we can hold the person *responsible* for, and they are not ones whose goodness or badness enters into our assessments of her *rationality*. I can't be held responsible for the hard-wired transitions that make my visual system liable to the Müller-Lyer illusion. Of course, once I find out that it is liable to such illusion, I am responsible for not falling for it, so to say, but that's a different matter.

Obviously, here we are in quite treacherous territory, the analogue to the question of free will and responsibility within the cognitive domain. Ill-understood as this issue is in general, it is even less well understood in the cognitive domain, in part because we are far from having a satisfactory conception of mental action.

But unless you are a skeptic about responsibility, you will think that there are some conditions that distinguish between mere mechanical transitions, and those cognitive transitions towards which you may adopt a participant reactive attitude, to use Strawson's famous expression (see also Smithies 2016).

And what we know from reflection on cases, such as those of the habitual depressive, is that mere associative transitions between mental states—no matter how conscious and content-sensitive they may be—are *not* necessarily processes for which one can be held responsible.

It is only if there is a substantial sense in which the transitions are ones that a thinker *performed* that she can be held responsible for them. That is the fundamental reason why Helmholtz-style transitions cannot, in and of themselves, amount to reasoning in the intended sense.

We can get at the same point from a somewhat different angle.

Any transition whatsoever could be hard-wired in, in Helmholtz's sense. One could even imagine a creature in which the transition from the Peano axioms to FLT is hard-wired in as a basic transition.

What explains the perceived discrepancy between a mere transition that is hard-wired in versus one that is the result of inference?

The answer I'm offering is that merely wired-in transitions can be anything you like because there is no requirement that the thinker *approve* the transition, and perform that transition as a result of that approval; they can just be programmed in. By contrast, I'm claiming, inferential transitions must be driven by an impression on the thinker's part that his premises support his conclusion.

To sum up the argument so far: a person's inferential behavior, in the intended sense, is part and parcel of his constitution as a rational agent. A person can be held

responsible for inferring well or poorly, and such assessments can enter into an overall evaluation of his virtues as a rational agent. Making sense of this seems to require imposing the Taking Condition on inference. Helmholtz-style sub-personal “inferences” can be called that, if one wishes, but they lie on the other side of the bright line that separates cognitive transitions for which one can be held responsible from those for which one cannot.

13. Non-reflective Reasoning in Humans

Helmholtz-style sub-personal transitions, however, are not the only potential source of counterexamples to taking-based accounts. For what about cases of reasoning 1.5, such as the (Rain) inference? Those, we have allowed, are cases of very common, person-level reasoning for which we can be held responsible. And yet, we have conceded that, at least phenomenologically, a taking state doesn’t seem to be involved in them. Why, then, do they not refute taking-based accounts?

There are at least two important points that need to be made in response.

The first is that the power of an example like that of (Rain) to persuade us that taking states are not in general involved in garden-variety cases of inference stems purely from the *phenomenology* of such cases: they simply don’t *seem* to be involved in such effortless and fleeting cases of reasoning. When we reflect on such cases, we find it phenomenologically plausible that there was a succession of judgments, but not that there was a mediating taking state.

The trouble, though, is that resting so much weight on phenomenology, in arriving at the correct description of our mental lives, is a demonstrably flawed methodology. There are many conscious states and events that we have reason to believe in, but which have no distinctive qualitative phenomenology, and whose existence could not be settled by phenomenological considerations alone. This point, it seems to me, is of great importance; so I shall pause on it.

Consider the controversy about *intuitions* understood as *sui generis* states of intellectual seeming.

Many philosophers find it natural to say that, when they are presented with a thought experiment—for example, Gettier’s famous thought experiment about knowledge—they end up having an intuition to the effect that Mr. Smith has a justified true belief but does not know.

Intuition skeptics question whether these philosophers really do experience such states of intuition, as opposed to merely experiencing some sort of temptation or disposition to judge that Mr. Smith has a justified true belief but does not know. Timothy Williamson, for example, writes:

Although mathematical intuition can have a rich phenomenology, even a quasi-perceptual one, for instance in geometry, the intellectual appearance of the Gettier proposition is not like that. Any accompanying imagery is irrelevant. For myself, I am aware of no intellectual

seeming beyond my conscious inclination to believe the Gettier proposition. Similarly, I am aware of no intellectual seeming beyond my conscious inclination to believe Naïve Comprehension, which I resist because I know better. (2007, p. 217)¹³

There is no denying that a vivid qualitative phenomenology is not ordinarily associated with an intellectual seeming. However, it is hard to see how to use this observation, as the intuition skeptic proposes to do, to cast doubt on the existence of intuitions, while blithely accepting the existence of occurrent judgments.

After all, ordinary occurrent judgments (or their associated dispositions) have just as little distinctive qualitative phenomenology as intuitions do. For example, right now, as I visually survey the scene in front of me, I am in the process of *accepting* a large number of propositions, wordlessly and without any other distinctive phenomenology. These acceptances are unquestionably real. But they have no distinctive phenomenology.¹⁴

To put the point in a slogan: phenomenology is often useless as a guide to the contents of our conscious mental lives. Many items of consciousness have phenomenal characteristics; however, many others don't.^{15,16}

That is why I am not deterred from proposing, by mere phenomenological considerations, that taking states are involved even in cases, like that of the (Rain) inference, where they may not phenomenologically appear to be present.

The second point that limits the anti-taking effectiveness of (Rain)-style examples is that much of the way in which our mental activities are guided is *tacit*, not explicit.

When you first learn how to operate your iPhone, you learn a bunch of rules and, for a while, you follow them explicitly. After a while, though, operating the phone, at least for basic "startup" functions, may become automatic, unlabored, and unreflective: you place your finger on the home button, let it rest there for a second to activate the fingerprint recognition device, press through and carry on as desired, all the while with your mind on other things.

That doesn't mean that your previous grasp of the relevant rules isn't playing a role in guiding your behavior. It's just that the guidance has gone tacit. In cases where, for whatever reason, the automatic behavior doesn't achieve the desired results, you will find yourself trying to retrieve the rule that is guiding you and to formulate its requirements explicitly once again.

¹³ See also Cappelen (2012) and Deutsch (2015).

¹⁴ It won't do to argue that these acceptances are "unconscious" as they are so easily retrieved.

¹⁵ I hesitate to characterize this distinction in terms of Ned Block's famous distinction between "phenomenal" and "access" consciousness, since this may carry connotations that I wouldn't want to endorse. But it is obviously in the neighborhood of that distinction. See Block (1995).

¹⁶ An anonymous referee suggested that an alternative hypothesis is that not all aspects of phenomenal life are open to naïve introspection, that theory, or other tools, are needed to guide one's introspective search for those phenomenal characteristics that are present. As I explain below (and also in Boghossian (forthcoming) these two hypotheses are not necessarily in competition with one another.

To put these two points together, I believe that our mental lives are tacitly guided to a large extent by phenomenologically inert conscious states that do their guiding tacitly.

One of the interesting lessons for the philosophy of mind that is implicit in all this is that you can't just tell by the introspection of qualitative phenomenology what the basic elements of your conscious mental life are, especially when those are intentional or cognitive elements. You need a theory to guide you. Going by introspection of phenomenology alone, you may never have seen the need to recognize states of intuition or intellectual seeming; you may never have seen the need to recognize fleeting occurrent judgments, made while surveying a scene; and you may never have seen the need to postulate states of taking.

I think that part of the problem here, as I've already noted, stems from over-learning the good lesson that Wittgenstein taught us, that in philosophy there has been a tendency to give overly intellectualized descriptions of cognitive phenomena.

My own view is that conscious life is shot through with states and events that play important, traditionally rationalistic roles, which have no vivid qualitative phenomenology, but which can be recognized through their indispensable role in providing adequate accounts of central cognitive phenomena. In the particular case of inference, the fact that we need a subject's inferential behavior to be something for which he can be held rationally responsible is a consideration in favor of the Taking Condition that no purely phenomenological consideration can override.

14. Richard's Objections

But is it really true that the Taking Condition is required for you to be responsible for your inferences? Mark Richard has argued at length that it is not (Chapter 6, this volume).

Richard agrees that our reasoning is something we can be held responsible for; however, he disputes that responsibility requires agency:

One can be responsible for things that one does not directly do. The Under Assistant Vice-President for Quality Control is responsible for what the people on the assembly line do, but of course she is not down on the floor assembling the widgets. Why shouldn't my relation to much of my reasoning be somewhat like the VP's relation to widget assembly?

Suppose I move abductively from *the light won't go on* to *I probably pulled the wire out of the fixture changing the bulb*. Some process of which I am not aware occurs. It involves mechanisms that typically lead to my being conscious of accepting a claim. I do not observe them; they are quick, more or less automatic, and not demanding of attention. Once the mechanisms do their thing, the conclusion is, as they say, sitting in the belief box. But given a putative implication, I am not forced to mutely endorse it. If I'm aware that I think *q* and that it was thinking *p* that led to this, I can, if it seems worth the effort, try to consciously check to see if the implication in fact holds. And once I do that, I can on the basis of my review continue to accept the implacatum, reject the premise, or even suspend judgment on the whole shebang.

In this sense, it is up to me as to whether I preserve the belief. It thus makes sense to hold me responsible for the result of the process. I say that something like this story characterizes a great deal of adult human inference. Indeed, it is tempting to say that all inference—at least adult inference in which we are conscious of making an inference—is like this.

(Chapter 6, this volume, p. x)

Richard is trying to show how you can be responsible for your reasoning, even if you are not aware of it and did not perform it. But all he shows, at best, is that you can be held responsible for the *output* of some reasoning, rather than for the reasoning itself, if the output of that reasoning is a belief.

But it is a platitude that one can be held responsible for one's beliefs. The point has nothing to do with reasoning and does not show that we can be held responsible for our reasoning, which is the process by which we sometimes arrive at beliefs, and not the beliefs at which we arrive. You could be held responsible for any of your beliefs that you find sitting in your "belief box," even if it weren't the product of any reasoning, but merely the product of association.

Once you become aware that you have that belief, you are responsible for making sure that you keep it if and only if you have a good reason for keeping it.

If it just popped into your head, it isn't yet clear that it has an epistemic basis, let alone a good one.

To figure out whether you have a good basis for maintaining it, you would have to engage in some reasoning. So, we would be right back where we started.

Richard (Chapter 6, this volume, p. 93) comes around to considering this objection. He says:

I've been arguing that the fact that we hold the reasoner responsible for the product of her inference—we criticize her for a belief that is unwarranted, for example—doesn't imply that in making the inference the reasoner exercises a (particularly interesting) form of agency. Now, it might be said that we hold she who reasons responsible not just for the product of her inference, but for the process itself. When a student writes a paper that argues invalidly to a true conclusion, the student gets no credit for having blundered onto *the truth*; he loses credit for having *blundered* onto the truth. But, it might be said, it makes no sense to hold someone responsible for a process if they aren't the, or at least an, agent of the process.

Let us grant for the moment that when there is inference, both its product and the process itself are subject to normative evaluation. What exactly does this show? We hold adults responsible for such things as implicit bias. To hold someone responsible for implicit bias is not *just* to hold them responsible for whatever beliefs they end up with as a result of the underlying bias. It is to hold the adult responsible for the mechanisms that generate those beliefs, in the sense that we think that if those mechanisms deliver faulty beliefs, then the adult ought to try to alter those mechanisms if he can. (And if he cannot, he ought to be vigilant for those mechanisms' effects.)

Implicit bias is, of course, a huge and complicated topic, but, even so, it seems to me to be misapplied here.

We are certainly responsible for faulty mechanisms if we know that they exist and are faulty.

Since we are all now rightly convinced that we suffer from all sorts of implicit bias, without knowing exactly how they operate, we all have a responsibility to uncover those mechanisms within us that are delivering faulty beliefs about other people and to modify them; or, at the very least, if that is not possible, to neutralize their effects.

But what Richard needs is not an obvious point like that. What he needs to argue is that a person can be responsible for mechanisms that deliver faulty judgments, even if he doesn't know anything about them: doesn't know that they exist, doesn't know that their deliverances are faulty, and doesn't know how they operate.

Consider a person who is, relative to the rest of the community of which he is a part, color blind, but who is utterly unaware of this attribute of his. Such a person would have systematically erroneous views about the sameness or difference of certain colored objects in his environment. His color judgments would be faulty.

But do we really want to hold him responsible for those faulty judgments and allow those to enter into assessments of his rationality? I don't believe that would be right.

Once he discovers that he is color-blind relative to his peers, then he will have some responsibility to qualify his judgments about colors and look out for those circumstances in which his judgments may be faulty. But until such time as he is brought to awareness, we can find his judgments faulty without impugning his rationality.

The confounding element in the implicit bias case is that the faulty judgments are often morally reprehensible and so suggest that perhaps a certain openness to morally reprehensible thoughts lies at the root of one's susceptibility to implicit bias. And that would bring a sense of moral responsibility for that bias with it. But if so, that is not a feature that Richard can count on employing for the case of inference in general.

Richard (Chapter 6, this volume, p. 92) also brings up the case where we might be tempted to say that we inferred but where it is not clear what premises we inferred from:

More significantly, there are cases that certainly seem to be inferences in which I simply don't know what my premises were. I know Joe and Jerome; I see them at conventions, singly and in pairs, sometimes with their significant others, sometimes just with each other. One day it simply comes to me: *they are sleeping together*. I could not say what bits of evidence buried in memory led me to this conclusion, but I—well, as one sometimes says, I just *know*. Perhaps I could by dwelling on the matter at least conjecture as to what led me to the conclusion. But I may simply be unable to.

Granted, not *every* case like this need be a case of inference. But one doesn't want to say that *no* such case is. So if taking is something that is at least in principle accessible to consciousness, one thinks that in some such cases we will have inference without taking.

There is little doubt that there are cases somewhat like the ones that Richard describes. There are cases where beliefs simply come to you. And sometimes they come with a great deal of conviction, so you are tempted to say: *I just know*. I don't know exactly *how* I know, but I just know.

What is not obvious is (a) that they are cases of knowledge or (b) that when they are, that they are cases of (unconscious) inference from premises that you're not aware of.

Of course, a Reliabilist would have no difficulty making sense of the claim that there could be such cases of knowledge, but I am no Reliabilist.

But even setting that point to one side, the most natural description of the case where *p* suddenly strikes you as true is that you suddenly have the intuition that *p* is true. No doubt there is a causal explanation for why you suddenly have that intuition. And no doubt that causal explanation has something to do with your prior experiences with *p*.

But all of that is a far cry from saying that you *inferred* to *p* from premises that you are not aware of.

Richard's most compelling argument against requiring taking for inference stems, interestingly enough, from the case of perception:

If we had reason to think that it was only in such explicit cases [he means reasoning 2.0] that justification could be transmitted from premises to conclusion, then perhaps we could agree that such cases should be given pride of place. But we have no reason to think that. I see a face; I immediately think *that's Paul*. My perceptual experience—which I would take to be a belief or at least a belief-like state that I see a person who looks *so*—justifies my belief that I see Paul. It is implausible that in order for justification to be transmitted I must take the one to justify the other. (Chapter 6, this volume, p. 96)

I agree that no taking state is involved in perceptual justification. I wouldn't say, with Richard, that perceptual experience *is* a belief: I don't see how the belief *That's Paul* could justify the belief *That's Paul*.

But I would say that the perceptual experience is a visual seeming with the content *That's Paul* and that this can justify the belief *That's Paul* without an intervening taking state.

But there is a reason for this that is rooted in the nature of perception (one could say something similar about intuition). The visual seeming *That's Paul* presents the world (as John Bengson (2015) has rightly emphasized) as this being Paul in front of one. When you then believe *That's Paul* on that basis, there is no need to take its seeming to be *p* to support believing its being *p*. You are already there (that's why Richard finds it so natural to say that perceptual experience is *itself* a belief, although that is going too far).

All that can happen is that a doubt about the deliverances of one's visual system can intervene, to block you from endorsing the belief that is right there on the tip of your mind, so to speak.

But that is the abnormal case, not the default one. Not to believe what perception presents you with, unless you have reason to not believe it, would be the mistake.

But the inference from p to q is not like that. A belief that p , which is the input into the inferential process, is not a seeming that q . And while the transition to believing that q may be familiar and well supported, it is not simply like acquiescing in something that is already the proto-belief that q .

15. Conclusion

Well, there are many things that I have not done in this essay. I have not developed a positive account of the taking state. I have not discussed possible regress worries. But one can't do everything.

What I've tried to do is explain why the shape of an account of reasoning that gives a central role to the Taking Condition has a lot to be said for it, especially if we are to retain the traditional connections between reasoning, responsibility for reasoning, and assessments of a person's rationality.

References

- Audi, Robert. (1986). Belief Reason, and Inference. *Philosophical Topics*, 14 (1), 27–65.
- Bengson, John. (2015). The Intellectual Given. *Mind*, 124 (495), 707–60.
- Block, Ned. (1995). On a Confusion about a Function of Consciousness. *Behavioral and Brain Sciences*, 18 (2), 227–87.
- Boghossian, Paul. (2014). What is Inference? *Philosophical Studies*, 169 (1), 1–18.
- Boghossian, Paul. (2016). Reasoning and Reflection: A Reply to Kornblith. *Analysis*, 76 (1), 41–54.
- Boghossian, Paul. (forthcoming). Do We Have Reason to Doubt the Importance of the Distinction Between A Priori and A Posteriori Knowledge? A Reply to Williamson. In P. Boghossian and T. Williamson, *Debating the A Priori*. Oxford University Press.
- Burge, Tyler. (1993). Content Preservation. *The Philosophical Review*, 102 (4), 457–88.
- Cappelen, Herman. (2012). *Philosophy Without Intuitions*. Oxford University Press.
- Chisholm, Roderick M. (1989). *Theory of Knowledge*, 3rd edition. Prentice-Hall.
- Deutsch, Max. (2015). *The Myth of the Intuitive*. MIT Press.
- Frege, Gottlob. (1979). Logic. In *Posthumous Writings*. Blackwell.
- Helmholtz, Heinrich von. (1867). *Treatise on Physiological Optics, Vol. III*. Dover Publications.
- Hlöbil, Ulf. (2014). Against Boghossian, Wright and Broome on Inference. *Philosophical Studies*, 167 (2), 419–29.
- Huemer, Michael. (1999). The Problem of Memory Knowledge. *Pacific Philosophical Quarterly* (80), 346–57.
- Kahneman, Daniel. (2011). *Thinking, Fast and Slow*. Macmillan.
- Kornblith, Hilary. (2012). *On Reflection*. Oxford University Press.
- Littlejohn, Clayton. (2012). *Justification and the Truth-Connection*. Cambridge University Press.

- Pettit, Philip. (2007). Rationality, Reasoning and Group Agency. *Dialectica*, 61 (4), 495–519.
- Pryor, James. (2000). The Skeptic and the Dogmatist. *Noûs*, 34 (4), 517–49.
- Pryor, James. (2005). There is Immediate Justification. In M. Steup and E. Sosa (eds.), *Contemporary Debates in Epistemology*. Blackwell.
- Pryor, James. (2007). Reasons and That-clauses. *Philosophical Issues*, 17 (1), 217–44.
- Siegel, Susanna. (2017). *The Rationality of Perception*. Oxford University Press.
- Smithies, Declan. (2016). Reflection On: On Reflection. *Analysis*, 76 (1), 55–69.
- Williamson, Timothy. (2000). *Knowledge and Its Limits*. Oxford University Press.
- Williamson, Timothy. (2007). *The Philosophy of Philosophy*. John Wiley & Sons.
- Wright, Crispin. (2014). Comment on Paul Boghossian, “What is Inference.” *Philosophical Studies*, 169 (1), 27–37.