Philosophical Issues
A Supplement to NOÛS

Philosophical Issues, 28, Philosophy of Logic and Inferential Reasoning, 2018
doi: 10.1111/phis.12115

DELIMITING THE BOUNDARIES OF INFERENCE

Paul Boghossian
NYU

1. Introduction

In this short essay, I tackle, yet again, the question of the nature of inference and elaborate somewhat on the agential conception of inference that I've been pursuing (Boghossian 2014, 2016 and forthcoming). What's new in this essay is a better way of setting up the issue about the nature of inference; a better identification of the concerns that lie at the back of this way of thinking about the topic; and a response to some important criticisms that have been made of the agential view I've been advocating.

Inference, of course, is to be distinguished from *argument* (see Harman 1986). An argument is just *a set of propositions*, some of which are designated 'premises,' and one (or more) of which is designated the 'conclusion.'

(Rain-Argument)

1. It rained last night
2. If it rained last night, the streets are wet
3. The streets are wet

By contrast, an inference is not a set of propositions, but a *movement* of thought.

Theoretical inference is movement between beliefs, or suppositions. Here you start out with some beliefs, and either end up adding a new belief, or losing some beliefs you already had, or modifying the credence with which you hold some belief, or changing the basis on which you hold some belief.

(Rain-Inference)

1. Bel: It rained last night
   So,
2. New Bel: The streets are wet

Practical inference is reasoning with a combination of intentions and beliefs:

1. Intention: I shall go to Paris
2. Belief: If I go to Paris, I shall buy a ticket
   So,
3. Intention: I shall buy a ticket

We will be concentrating on theoretical inference.

What is the nature of this movement of thought I am calling 'inference'?

And how do we specify our subject matter without prejudging important issues?

Here is one promising approach: Start by looking at the entire range of phenomena that have been *labeled* 'inference' and ask later whether they form a natural epistemic or mental kind, or whether they should be carved up into separate subgroups.

2. The Many Types of Phenomena Labeled 'Inference'

2.1. Inference 2.0: Fully Explicit Reasoning (FER)

I start with what to my mind is the paradigm example of inference, what we may call 'fully explicit inference.' In fully explicit inference, we have the following elements:

(1) You consider explicitly some proposition that you believe, for example p.
(2) (Meta Q) You wonder, in the context of some particular inquiry, what other relevant proposition you have reason to believe on this basis?

    It then strikes you that q follows from p. Hence,
    (Taking) You take it that q follows from p
    [Optionally, you might ask yourself:

(3) Is q plausible? Is it less plausible than the negation of p?
(4) You conclude that q is not less plausible than not-p.]

(5) So, you judge q. (You add q to your stock of beliefs.)

Leaving aside the optional bit, Inference 2.0 looks as though it calls for the following model: The thinker

a. Explicitly judges the premises of the inference;
b. Explicitly wonders, in the context of some particular inquiry, what else she has reason to believe, what other proposition the premises support.[1]
c. Explicitly *takes* the premises to *support* the conclusion.
d. Knows the properties of the premises in virtue of which they support the conclusion (knows the *epistemic principle* that validates moving from the premises to the conclusion).
e. *Believes* the conclusion *because* she believes the premises.
f. *and* believes the conclusion *because* she *takes* the conclusion to be supported by the premises.

Call this the Intellectualized Model of Fully Explicit Reasoning, or IMFER for short.

Inference 2.0 is admittedly not the most common type of inference. But it is probably not as rare as it is fashionable to claim nowadays. Inference 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 philosophy paper. But it also happens in a host of other cases that are much more mundane. Basically, anyone engaging in 'critical reasoning,' would be engaging in inference 2.0.

2.2. Inference 1.75

Inference 1.75 is just like 2.0 except that (d) does not obtain: the thinker doesn't know which epistemic principle validates the transition; he just knows that it seems to him that the premises support the conclusion.

2.3. Inference 1.5: Quick, Effortless Inference

The Rain Inference outlined earlier is an example of such reasoning. Here, the premise and conclusion are both things of which you are aware. But, it would seem, there is no explicit meta-question; nor, it would seem, is there, a taking state, although the adult thinker would certainly be capable of thinking both types of thoughts. Rather, the conclusion comes seemingly immediately and automatically. I will call this an example of *inference 1.5*.

In this case, IMFER's (a) and (e) are both true; but (b), (d) and (f) all seem false.

2.4. Inference 1.25: Inference in Children

Consider an example drawn from Boghossian (2001). Jack and Jill are playing hide and seek. Jack sees Jill's bicycle leaning against a tree and reasons as follows.

1. If Jill were hiding behind that tree she wouldn't have leaned her bicycle against it.
2. So, she must be hiding behind some other tree.

Here, one might think, this is just like inference 1.5 except that, unlike the adult, the child may not even be capable of thinking the (Meta Q) or the Taking thought.

2.5. Inferences in Non-Human Animals

We are prone to describe various non-human animals as capable of inference. Kornblith (2012) offers us a contrast between the tongue-flicking behavior of a frog and the broken-wing display of a piping plover. The first is said to illustrate mere stimulus-bound behavior; but the latter is said to furnish a clear case of reasoning.

> The frog's response to flies is . . . [stimulus bound]. If one rolls a piece of birdshot (BB) in front of the frog, it will suck it up as if it were a fly. Well, we all make mistakes. But the frog will do it a second time if a second BB is rolled in front of it; and a third; . . . The frog will not learn from its experience. (p. 51)

By contrast, the piping plover seems to show a greater capacity for adapting productively to its experience. When a nesting piping plover detects a potential predator approaching its nest, it feigns a broken wing, luring the predator away from the nest; once the predator is far enough away, the bird flies off. However, if it becomes familiar with a particular potential predator, and establishes though experience that it is non-threatening, the plover ceases to engage in the broken-wing display when it comes to that particular creature. Kornblith writes:

> This is not some stimulus-bound performance. In order to explain it, we need to see the plover as picking up information about its environment and integrating it with stored information about the past. (pp. 51–2)

> 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. (p. 55)

2.6. Unconscious and Sub-personal Inferences of, e.g., the Human Visual System

The word 'inference' is also often used in psychology and cognitive science to stand for sub-personal processes and transitions. Helmholtz 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. In some important sense, our visual system must conduct a multitude of 'computations' and 'inferences' if we are to enjoy the 3D visual image it presents us with.

2.7. Inferences and Computations in Computers, Calculators and 'Intelligent' Machines

Finally, it is very common for us to talk about computers and other types of 'intelligent' machines as performing computations and inferences.

3. Inference as a Central Epistemological Notion

As far as I can see, this is a fairly comprehensive list of all the types of phenomena that we have been tempted to apply the label 'inference' to. It is time to ask, though, whether all these instances form a mental or epistemic *kind*? Do they belong together in some important sense, or do we have here a bit of a jumble, an accident of terminology? And how do we decide such a question?

It's up to us to say what we're interested in. What are the features that we care about when, in epistemology, we invoke the notion of inference?

I propose that the notion we care about in epistemology must have the following three features.

B. (Basing) When you infer from P to C, you *establish* the premises P as *your reason* for believing the conclusion C, you establish them as the *basis* on which you believe B.
Q. (Quality) Given this basing fact, your belief can be assessed as resting on good or bad reasons.
R. (Responsibility) And given this fact, too, you can be held *responsible* for having reasoned well or badly. The quality of your reasoning will enter into an assessment of your rationality.

Let me make a few brief comments about each of these conditions, B, Q and R.

Basing seems to me to be the minimal core of inference in the sense in which it plays a central role in epistemology. When you infer from P to C you establish P as your reason for believing C and it is this fact that makes it true that your belief has an inferential reason and not an observational one.

On Q: It is in virtue of the fact recorded in B that your belief is assessable as having a good or bad inferential reason supporting it.

What makes some reason a good or bad reason for believing C? Well, for a reason P for believing C to be good, C must, in some broad sense of 'follow,' follow from P. But that's not enough. FLT follows from the Peano axioms, but it wouldn't do just to infer the one from the other.

It looks as though what's also needed is that the conclusion not be at too far a distance from the premise. But what does that mean?

The only good answer that I can think of is that the step from premise to conclusion be such that the thinker have some *appreciation* that the conclusion does indeed follow from the premises. Of course, unless this condition is to generate a super-task, it had better be that, for a wide range of basic inferences, this appreciation is *non-inferential* in character.[2]

Although I will have occasion to return to this thought, it's worth remarking that considerations based *solely* upon the notion of *good reasoning* seem to have led us to the following view: Inferences require that the thinker *take* her premises to support her conclusion and that, in good reasoning, this taking is backed by an intuition to the effect that the taking is true.

On R: What's needed to make it apt to hold you responsible for the quality of your reasoning?

It's important to realize that this condition is not redundant on the previous one. Your digestion can be assessed as good or bad; but you can't be held responsible for how well you digest (except indirectly) since it is not under your control.

Inference, as I have characterized it, is mental behavior and, so, for it to make sense to hold you responsible for your inferences, inferring has to be something you *do*, and not just something that happens to you. It has to be a mental *action* of yours, something you have control over, and which you could have done differently, had you thought it desirable to do so.

So much for my brief comments on features B, Q and R. Why should a mental process with features B, Q and R be important in epistemology?

Well, for better or worse, in epistemology we are obsessed with the notion of rationality. We are obsessed with the idea that there are better and worse

ways for you to manage your beliefs; and that these ways reflect on your virtues as an epistemic agent.

Given all this, it is not a mystery that we should be interested in mental transitions with features B, Q and R. Reasoning in this sense is a recognizable epistemic kind. It is a personal-level establishment of your reasons for belief, in a way that grounds both the assessment of the rationality of what you believe and of your virtues as an epistemic agent.

Now, given that it is inference in this sense that I am stipulating I am interested in, let us look back at the list of things that have historically been labeled 'inference' and ask which of them belongs to this kind?

Clearly, 'machine inferences' are out—we don't hold machines responsible for their reasoning; at best, we might hold their designers, builders or users responsible.

Sub-personal visual system inferences are also out. We don't hold the person responsible for the 'inferences' made by his sub-personal cognitive mechanisms. Those, we think, are just programmed and not under the thinker's rational control.[3]

What about the piping plover or other non-human animals? Once again, if we keep in mind that the notion of control, in the sense of its being possible for the thinker to have done otherwise, is important to our conception of mental action, then it is hard to see the non-human animals as qualifying.

Perhaps unsurprisingly, then, the concept of inference that is central to epistemology encompasses *only*, but also *all*, the examples of human inference on our original list, since even 1.25 inferences seem to have satisfy all of B, Q and R. What we need to do, therefore, is provide an account that captures inferences 1.25 and up, from the reasoning in children all the way to fully explicit reasoning.

4. The Puzzles We Confront

Even this more limited task, however, is quite daunting. We need to provide not only a satisfactory picture of *each* of these examples of human inference, but, in addition, a satisfactory picture of how they all belong to the same natural epistemic kind.

Let me begin with the puzzles that attend accounting for Inference 2.0, fully explicit reasoning:

As I've previously noted, there's no puzzle about how I2.0 could be mental action (except for the general problem of explaining how any action could be genuinely free). The agent has an aim; she has a view about a way

of accomplishing that aim; and she performs an action as a result of that combination.

However, even with respect to something as relatively clear as FER, there's, of course, a potential Carrollian regress, both about how one could *enter* into a taking state, and about how that taking state could rationally *control* the formation of the conclusion.

I used to think these regress problems were insuperable; but I am now inclined to believe that they have adequate solutions.

There are two regress worries: The first one concerns *entry* into the taking state.

> (Ingress Regress) How can we rationally get into the taking state? If that state has general content, then it looks as though we would have needed to have got there by inference. But that looks like a super task. If each taking state is arrived at by inference, and each inference involves its own proprietary taking state, then how is it possible for there to be any inference at all?

The second problem concerns how we could bring the taking state to rationally bear on the drawing of the conclusion?

> (Egress Regress) The conclusion is said to be drawn in part as a result of the taking state. However, if the taking state has a general, typically conditional content, then it looks as though inference from it will be required to explain how it leads to the thinker's drawing the conclusion that he draws.

A proper treatment of these two problems would require a much longer discussion than I am able to give here; here I will have to settle for briefly indicating the type of answer that I favor.

With regards to the Ingress problem, as I've already indicated, the answer is that taking states have an intuitional epistemology, so the regress problem is nipped in the bud. Why does the thinker take his MP premises to support his MP conclusion? Because he has the vivid intellectual impression that whenever MP premises are true a MP conclusion must also be true.

Many questions come crowding in but it's not part of my brief in this essay to provide a comprehensive defense of an intuitional epistemology (for more details see Boghossian forthcoming). I will have to make do with a few observations.

As I've already indicated, something like this view is already indicated by a plausible account of what makes a reason a good reason for believing something.

A further point: Taking states can seem like beliefs; but it's important that, although they are belief-like, they are distinct from beliefs. One of the ways in which they are distinct is that underived taking states, that is, taking states not derived from other taking states, can only be entered into via intuitions (and not by testimony or inference). This fits in well with the

view I expressed earlier, that in a certain range of basic, underived cases the appreciation that the premises support the conclusion must be non-inferential.[4]

The short answer to the Egress problem is that we know of many examples of intentional states with general, conditional contents rationally controlling behavior without the benefit of inference. For example: the tennis player's intention to respond with a backhand if the ball comes to his left, and with a forehand if it comes to his right. That general intention can rationally control the tennis player's behavior without his having to conduct person-level inferences in order to act upon it.[5]

It stands to reason that there should be solutions to the problems that seem to afflict I2.0, because it is intuitively entirely clear that there are inferences 2.0 and that such inferences are well captured by IMFER.

Let us turn our attention to Inference 1.75. If what I've been saying about I2.0 is correct, I1.75 should pose no new problems. The only difference between I2.0 and I1.75 is that in the former, but not in the latter, the thinker believes not only that the premises support the conclusion but knows in virtue of what they do. However, we understand perfectly how someone can rationally perform an action that seems right to them even if they don't know exactly *why* it's right.

5. Inferences 1.5

The big problem for theorizing in this area is presented by inferences of levels 1.5 and lower. The reason is that, starting with inferences of level 1.5, there seems to be no explicit aim to guide the mental transitions, and no explicit taking state to capture the thinker's view of how the aim is to be achieved.[6] In the absence of such goal-directedness, how are we to explain why 1.5 style transitions count as inferences in our sense—that is, inferences that possess features B, Q and R?

The answer is that I don't believe we can. Sticking for now with level 1.5 (I will come to 1.25 later), I believe the only way we can account for features B, Q and R requires 1.5 transitions to be regarded to be goal-directed actions over which the thinker has rational control.

To see why, try doing without an aim and a taking state. Imagine someone who thinks the premise judgment, and then finds himself thinking the conclusion judgment, and that's it. What about this black-box process would distinguish it from a mere *association* of judgments, something which could not be said to have features B, Q and R? A mere association of judgments wouldn't establish anything as the basis for believing something; it couldn't be assessed as done well or badly; and it isn't something that we could hold the person responsible for.

What, however, do we do about the fact that, phenomenologically speaking, there seem to be no explicit aim states or taking states involved in ordinary, quick and automatic 1.5-style inferences?

The answer has to be that, although such states are not present explicitly, they are present *tacitly.* Later on, I will say a bit more to elucidate this answer; for now, let me note that this is hardly the first context where an appeal to guidance by a tacit state is called for. Most of the mental phenomena that are examples of guidance—by rules, long-term plans, standing beliefs, and so forth—involve guidance by tacit states.

6. Attempting to Make Do Without Aims and Takings

Hilary Kornblith and Mark Richard (forthcoming) have both tried to deny this; they have both tried to present pictures of inference which do not invoke an aim or taking state.

Kornblith's picture, as we've already seen, is that inferences are not mere causal transitions; rather, they are *content-sensitive* transitions. He says:

> But, as I see it, there is now every reason to regard [the plover's] 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. To cite an example I've used in the past, 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 that is sensitive to their contents. However, he is not thereby inferring from the one proposition to the other: he is not establishing the judgment about fun as an epistemic basis for his judgment about suffering and the latter cannot be assessed as a poorly held judgment in virtue of its association with the former. (For more discussion of Kornblith see Boghossian 2016).

I think that the notion of a content-sensitive transition (with a specified functional role) may well account for what we are tempted to call 'inferences' in non-human animals, sub-personal visual systems and machines. But it doesn't capture the notion of inference that I have isolated as forming our subject matter and which plays a central role in epistemology.

Mark Richard proposes a different way of avoiding a reliance on my Mental Action picture of inference.[7] Unlike Kornblith, Richard agrees that responsibility is an important part of the phenomenon of human inference. But he disputes that one can only be held responsible for things one *does directly*. You don't perform your reasoning; but you can still be held responsible for it.

> 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 endorse it. I can accept the implicatum, or reject the premise, or try to consciously check to see if the implication in fact holds. 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—I'm tempted to say all—adult human inference.

Richard's picture is this: you consider some proposition A (the light won't go on). Although he doesn't say this, you presumably ask yourself: I wonder why. This already establishes, in contrast with what Richard would like us to believe, that what happens next has an acknowledged aim, that of figuring out why.

You then find yourself judging B (I probably pulled out the wire). You're not aware of reasoning from A to B. You're not clear what happened. You just find a new belief in your belief box.

In a passage I didn't quote, Richard says that, given that I have the concept of 'consequence,' I will be able to *interpret* the appearance of my new belief as a sign of the conclusion's following from the premises. I suppose his picture is that you *infer* that you did some inferring by noticing that you considered some question and a new belief occurred to you.

This is quite a strange picture of our mental lives. I don't think we usually infer that we have done some inferring. But never mind. What about the proffered account of responsibility for one's inferences? On Richard's account, you don't count as responsible because you *performed* your inference; rather, it's because, in retrospect, you can be held responsible for whether you endorse the belief that it results in.[8]

I ignore here that, at best, this simply accounts for the fact that you can be held responsible for the *output* of your reasoning, not for the reasoning itself. Suppose you do consider whether you should keep belief C, having found yourself with that belief, as a result of wondering about P.

When you do ask yourself this question, in order to answer it, what you have to do is some person-level fully explicit reasoning: Does C really follow from P? If it does, should I accept C as well as P, or should I reject both? Is C so implausible that it would be more plausible to reject P than to accept C? If so, I ought to reject P, rather than accept C. If not, then I ought to accept.

But this thing that you now need to do, in order to assess whether to keep the output of the reasoning, is just some person-level fully explicit reasoning.

So, even on Richard's own account, what makes it the case that you can be held responsible for some unreflective reasoning is that you can, if called

upon, perform some fully explicit reasoning. So, it can't be true, even on his account, that *all* reasoning is the sort of black box process that he describes.

7. Guidance by Tacit States

Neither Richard's nor Kornblith's attempt to explain inference without aims and takings succeeds; and I can't think of how to execute their agenda better than they do it. As I've indicated without proving, it's impossible to explain how a mental transition could have features B, Q and R without seeing it as a mental action—that is, without seeing it as being directed towards a goal and embodying a view about that goal is to be achieved.

If this is true, though, how are we to reconcile the fact that 1.5 inferences don't seem to involve an explicit aim and an explicit taking state?

The answer, of course, is that while they may not involve these things explicitly, they involve them *tacitly*. The key is to show not only that this picture provides a satisfying account of inference 1.5 but also that it is not an ad hoc maneuver, but rather relies only on resources that are independently plausible.

I've already expanded at length elsewhere about how such a picture provides a satisfactory account of 1.5-level inferences (see Boghossian 2014, 2016 and forthcoming). Here I will only list them briefly. Even with quick and automatic inferences of the 1.5 kind, it is always appropriate to precede the conclusion with a 'So' or 'Therefore.' What are those words supposed to signify if not that the agent is taking it that her conclusion is justified by her premises?[9] Postulating a taking state accounts well for the Inferential Moorean phenomenon, which applies equally to 1.5 inferences. Furthermore, taking-based accounts offer a neat explanation of how there could be two kinds of inference—deductive and inductive.[10] Finally, such accounts explain why certain inferences—such as ones that jump directly from the Peano axioms to Fermat's Last Theorem—would not only be unjustified but also, in some sense, impossible.

But does the idea of guidance by tacit states make sense, or is it just an ad hoc response to our problem about inference?

There is no question that we need the idea of guidance by tacit states (tacit guidance) in a wide variety of contexts, quite apart from the present one. Ordinary life teems with examples of such guidance.

The idea of a tacit state is the idea of a mental state that seems not to be an occurrent state of consciousness yet is still present and plays a role in guiding behavior. When we first learn to use an iPhone, the rules for operating the machine are learned explicitly, and followed explicitly. However, after a certain period of use, the procedure, at least for basic 'startup' functions, may become automatic, unlabored and unreflective: you rest 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.

It's not as though your grasp of the rules has disappeared; if it had, you wouldn't be able to operate the phone, and when something goes wrong you rehearse them to yourself to make sure that they have been correctly followed. However, the grasp is no longer an explicit item of consciousness, but rather serves to guide the behavior tacitly. 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. The same goes for playing the piano, using a language, operating the elevator and countless other ordinary activities.

Nevertheless, it may still be felt that there is a problem about invoking aims and takings for 1.5-level inferences. For even if a taking state is said to be tacit, it is still a state with intentional content. But do typical adults, with no prior training in philosophy, possess the concept of *consequence*?

Well, what's at issue is a generic notion of 'follows from' or 'supports,' one that you have to learn when you learn how to form beliefs critically, that is by interrogating the evidence said to be in their favor. I don't know the empirical literature on this topic, but I would conjecture that this happens pretty early on. Indeed, wouldn't something like the notion of 'following from' and 'not following from' come just with a grip on the appearance/reality distinction, something which occurs in children around the ages of 3 or 4? I am, thus, not inclined to worry about whether the notion of taking makes too rich a conceptual demand on the ordinary adult thinker.

8. Inferences in Children

What about inferences 1.25?

If what I have been saying just now is correct, children do not pose a problem for the picture I'm defending. Children certainly seem to understand 'so' early on. (Indeed, they seem to begin saying "so what?" rather early in life!)

And even in the case of children, we adopt 'participant reactive attitudes.' We might say to a child: "you mustn't jump to conclusions like that."[11]

In any case, I am reluctant to place too much weight on alleged facts about the mental lives of creatures—whether they be children or animals—that we don't understand very well.

9. Conclusion

Many different processes have been labeled 'inference,' but these don't all form a natural epistemic or mental kind. The notion of

inference that plays a central role in epistemology is the notion of a process with certain constitutive features: it's a process that establishes a proposition the thinker believes as her basis for believing something else, with consequential normative implications. This process cannot be adequately explained without construing it as goal-directed mental action in which the thinker takes herself to be trying to figure out what she has reason to believe. If this is correct, it has at least three important implications.

First, it gives us a handle on the difference between human inference and the various other transitions that we are tempted to call by that label, in sub-personal systems, non-human animals and machines.

Second, it gives us much better prospects for giving a satisfying internalist account of inferential justification, unlike the accounts in terms of 'blind inference' which I had explored in previous work.

Finally, charges that accounts such as mine present impossibly over-intellectualized accounts of our mental lives can be satisfactorily addressed.[12]

Notes

1. That inference is goal-directed in this way was always part of my view, although I used to think that emphasizing that inference always involved *taking* the premises to support the conclusion was enough to capture that. Comments by Joshua Schechter led me to believe that emphasizing the presence of a goal explicitly was important.
2. For criticism of this answer and for an alternative account see Schechter (forthcoming).
3. This is far from uncontroversial—for an opposing view, see Siegel (2017).
4. Thanks to Josh Schechter for making me clarify this.
5. I believe that, some time ago, Rory Madden urged this view on me. More recently, Crispin Wright (2014) has also expressed sympathy for a similar view.
6. And the problem only gets worse with inferences 1.25, since in those cases not only does there seem to be no such explicit states, there seems to be no capacity to have such states.
7. I also discuss Richard's account in Boghossian forthcoming.
8. This is the analogy with the VP of Quality Control. She didn't make the widgets herself; but she is responsible for not letting any defective widgets get through.
9. Pettit made this observation in his (2007, p. 500).
10. The following two points were already mentioned in Boghossian (2014); I mention them here again for the sake of completeness.
11. Thanks to Crispin Wright for this observation.
12. I am grateful to Joshua Schechter, Crispin Wright, participants at the UNC Workshop on Inference in March 2018 and at my joint seminar with Wright at NYU in Spring 2018 for comments that helped improve this paper.

References

Boghossian, Paul. (2001). How Are Objective Epistemic Reasons Possible? *Philosophical Studies* 106(1/2), 1–40. https://doi.org/10.1023/A:1013141719930

Boghossian, Paul. (2014). What is Inference? *Philosophical Studies*, 169(1), 1–18. https://doi.org/10.1007/s11098-012-9903-x.

Boghossian, Paul. (2016). Reasoning and Reflection: A Reply to Kornblith. *Analysis*, 76(1), 41–54. https://doi.org/10.1093/analys/anv031

Boghossian, Paul. (forthcoming). Inference, Agency and Responsibility. In Magdalena Balcerak Jackson and Brendan Balcerak Jackson (eds.), *Reasoning: Essays on Theoretical and Practical Thinking*. Oxford University Press.

Harman, Gilbert. (1986). *Change in View*. MIT Press.

Von Helmholtz, Heinrich. (1867). *Treatise on Physiological Optics Vol. III*. Dover Publications.

Kornblith, Hilary. (2012). *On Reflection*. Oxford University Press.

Pettit, Philip. (2007). Rationality, Reasoning and Group Agency. *Dialectica*, 61(4), 495–519. https://doi.org/10.1111/j.1746-8361.2007.01115.x

Richard, Mark. (forthcoming). Is Reasoning a Form of Agency? In Magdalena Balcerak Jackson and Brendan Balcerak Jackson (eds.), *Reasoning: Essays on Theoretical and Practical Thinking*. Oxford University Press.

Schechter, Joshua. (forthcoming). Small Steps and Great Leaps in Thought: The Epistemology of Basic Deductive Rules. In Magdalena Balcerak Jackson and Brendan Balcerak Jackson (eds.), *Reasoning: Essays on Theoretical and Practical Thinking*. Oxford University Press.

Siegel, Susanna. (2017). *The Rationality of Perception*. Oxford University Press.

Wright, Crispin. (2014). Comment on Paul Boghossian, "What is Inference?" *Philosophical Studies*, 169(1), 27–37. https://doi.org/10.1007/s11098-012-9892-9.