

## Reasoning and Reflection: A Reply to Kornblith

### 1. Introduction

Hilary Kornblith's book is motivated by the conviction that philosophers have tended to overvalue and overemphasize reflection in their accounts of central philosophical phenomena. He seeks to pinpoint this tendency and to correct it.

Kornblith's claim is not without precedent. It is an oft-repeated theme of 20th-century philosophy that philosophers have tended to give 'overly intellectualized' accounts of important phenomena. One thinks here of Wittgenstein, Ryle and many others.

One version of this charge is that philosophers have tended to appeal to higher-order thoughts when first-order thoughts about the world are all that's needed.

A more specific version of this charge is that philosophers have tended to appeal to second-order thoughts with *normative*, or quasi-normative, contents when all that's needed are first-order thoughts with factual contents.

It is this second version of the charge that Kornblith is particularly interested in pressing. Although he doesn't spell it out, the connection between this project and Kornblith's previous work on naturalistic conceptions of epistemology should be fairly obvious. Very roughly, if you want humans to look a lot closer to the lower animals, then you'd better think that most central human abilities can be explained without appeal to reflection and without appeal to normative thought.

What's good and important about Kornblith's book is that he gives this charge a sustained and illuminating treatment. He looks in detail at accounts of knowledge, reasoning, epistemic agency, free will and normativity; he identifies sympathetically some of the temptations to think that we must resort to second-order resources to explain these phenomena; and he attempts to show that the appeal never works and is, in any case, not needed, since first-order accounts manage very well.

### 2. Causality + ? = reasoning

I will look at the case of reasoning in some detail. I pick this case partly because it interests me, but partly also because it provides Kornblith with his strongest hand: reflection-based accounts of free will, or normativity, or justification, may appeal to varying degrees; but reflection-based accounts of reasoning look downright treacherous. As we will see, though, even in this most unpromising of cases, naturalist accounts that give no role at all to reflection seem utterly implausible.<sup>2</sup>

What is it for you to reason from  $p$  to  $q$ ? All parties are agreed that for such reasoning to take place your acceptance of  $p$  must cause your acceptance of  $q$ . But, clearly, not every such causal transition would be reasoning.<sup>3</sup>

Kornblith agrees in setting up the issue in this way:

If I am to reason from one belief to another – if the first belief is to be part of my reason for believing the second – then the first belief must have played some causal role in producing (or sustaining) the second belief.

But surely this is not enough. There are, after all, lots of ways in which one belief might play a role in producing or sustaining a second belief which would not constitute either

reasoning or believing for a reason. ... So what else is required if one's belief is to be one's reason, or part of one's reason, for holding a second? (42)

In response (as we will further see below), some philosophers have suggested that it is enough to add that the transition from the first belief to the second be sensitive to the judgements' *contents*. It must be in virtue of the belief's contents that the first causes the second.

But that is clearly not enough. In a habitually depressive person, the judgement 'I am having so much fun' may routinely cause and explain his judging 'But there is so much suffering in the world'; this transition would be content-sensitive, but would not be reasoning. It would be content-sensitive *associative thinking*.

At this point, you might be attracted to the suggestion that the depressive's thinking cannot count as reasoning because his first judgement – that he is having so much fun – doesn't *support* his second – that there is so much suffering in the world. Indeed, the first judgement might be thought to (somewhat) undermine the second.

But that can't be a good suggestion – it conflates reasoning per se with *good* reasoning. Although this conflation is quite common in the literature, it must be avoided: our theory of reasoning must allow for bad reasoning. And one way that reasoning can be bad is for its premise to fail to support its conclusion.

What, then, *should* we say about the difference between a merely causal transition and reasoning?

If we think about our depressive's associative thinking, the answer seems to leap out at us. The reason that his thinking doesn't count as reasoning is *not* (as we've said) because his first judgement doesn't support his second, but, rather, because he doesn't *take* it that his first judgement supports his second. The first judgement simply *suggests* the second one to him; he doesn't draw the conclusion because he takes it to be supported by the first.

This way of thinking about reasoning echoes a remark of Frege's ( 1979 ):

To make a judgment because we are cognisant of other truths as providing a justification for it is known as *inferring*. (3)

Frege's characterization needs some tweaking. We needn't always infer from truths, if we are to count as inferring. We may reason from falsehoods, or even from imperatives, which are incapable of truth or falsehood. Furthermore, 'cognisant' has a success grammar built into it, and so Frege's characterization implies that one cannot reason badly. But, as I've been emphasizing, it is possible to reason badly.

So, I would rather put it like this: But this is clearly a view in the same spirit as Frege's. And it is one that gives a crucial role to some sort of reflective element – the taking.

- (Taking) A transition from some acceptances to a conclusion counts as inference only if the thinker *takes* her conclusion to be supported by the accepted truth of those other acceptances, and draws her conclusion *because* of that fact.

Now, there are, of course, many different ways of elaborating such a view. As formulated so far, it is just a schema for a view, one which signals that something has to play the role of 'taking' but without specifying exactly what it is that plays that role, nor how it plays it. In some of my other work, I have looked at different versions of such a view in some detail ( Boghossian 2014 ). I will return to this issue further below.

### 3. The adequacy of first-order views of reasoning

Before doing so, however, I want to look at Kornblith's arguments for the claim that no account of reasoning requiring a reflective element could be correct. Kornblith has both a positive and a negative argument for this claim.

His positive argument proceeds via the claim that: His negative argument, on the other hand, depends on claiming that:

- (A) We know of clear-cut cases of reasoning that incorporate no reflective element, hence no such condition could be necessary.
- (B) No view of reasoning incorporating such a second-order element can be made to work.

I will look at each argument in turn.

The supposedly clear-cut cases of non-reflective reasoning include both non-human and human instances. The non-human instances are of special interest since, if they were convincing, they would clinch the case, its being widely accepted (it's not entirely clear why) that non-human animals are incapable of reflective thought.

### 4. The case of the clever plover

What are these allegedly clear cases of non-human reasoning? Kornblith offers us a contrast between the tongue-flicking behaviour of a frog and the broken-wing display of a piping plover. The first is said to illustrate mere stimulus-bound behaviour; 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. (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. Kornblith writes:

One may reasonably ask whether this behavior is simply stimulus-bound. After all, there are behaviors like this in the animal world which are simply wired in, like the frog's tongue-flicking, and show no sign of sensitivity to other information which the animal may, in some sense, possess.

But the piping plover's broken wing display is not like this. Thus, for example, if a human being walks on a path which would come close to the nest, the plover will approach the individual and perform the broken wing display. But if that person continues on her path without disturbing the nest, and follows the same path on a daily basis, then the plover will stop performing the display. When a different individual approaches, however, on the very same path, the plover goes back to the broken wing performance. 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. (51–52)

There is no disputing that there is an important difference between the behaviour of the frog and that of

the plover. The frog is programmed to perform a particular action (flick its tongue) whenever it registers a certain bit of information about its environment – the presence of a fly-like object. The plover is capable not only of registering information about its environment, but of having those informational states interact with one another, thus allowing it to have more sensitive interactions with its environment.

Kornblith, however, goes on to interpret this distinction between the frog and the plover as a distinction between behaviour that is merely ‘wired in’, versus behaviour that manifests a capacity to *reason*. He says:

When the plover discovers, after repeated interactions with this particular individual, that no threat is present, she ceases to respond by engaging in the display. This requires the integration of new information with old, and this is why we need to posit a representational system to account for such sophisticated behavior. 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. (55, emphasis added)

This, of course, is the point at which, if we are to get clear on the philosophical issue about reasoning, we need to go slowly. Why is there ‘every reason’ to regard the interactions between the plover’s informational states as reasoning?

Kornblith could just be thinking, I suppose, that it is intuitively clear that the plover reasons. But that really doesn’t seem right. Intuitively, one recoils not only from the claim that the plover *reasons*, but even from what looks to be the more modest claim that it has *beliefs*. Rather, the plover looks to have only informational states and algorithms governing the transitions between them.

And *prima facie*, anyway, the difference between the frog and the plover is not a difference between behaviour that’s ‘wired in’ as opposed to behaviour that’s the result of reasoning; rather it appears to be the difference between one sort of behaviour that’s wired in and a slightly more complex form of behaviour that’s wired in. Whereas the frog’s behaviour looks to be governed by a rule with a single type of input and a single type of output,

If C, stick out your tongue and suck it up!

the plover’s broken wing behaviour is governed by a more complex rule:

If C and not D, display broken wing behaviour; otherwise, carry on!

Or something like that.

I don’t see that a slight increase in the complexity of a hard-wired algorithm governing a creature’s behaviour can, all by itself, make for the difference between reasoning and mere mechanism, no matter how much the more complex algorithms might contribute to a ‘more sensitive’ interaction with the environment.

In fact, Kornblith doesn’t seem to be primarily relying on intuition at all here. Instead, he seems to think that his case is secured by an *argument*, one of whose premises is that reasoning is nothing other than transitions among informational states on the basis of their content:

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*. (55, emphasis added)

This is clearly not a sound argument. An alarm system has the capacity to make transitions between informational states on the basis of their content, without its being true that it is capable of reasoning.

Kornblith could try insisting that we restrict our attention to processes that occur in living organisms; but that still wouldn't help. The process of *breathing*, for example, would count as reasoning on that criterion. Breathing is controlled automatically and subconsciously by the respiratory centre at the base of the brain. Sensory organs in the brain and in the aorta and carotid arteries monitor the blood and sense its oxygen and carbon dioxide levels. When these sensors detect an increased concentration of carbon dioxide in the blood, the brain stimulates the lungs to breathe more deeply and frequently. These transitions 'involve interaction among representational states on the basis of their content', but do not count as reasoning.

Kornblith could also try insisting that we should restrict our attention to transitions between informational states that are *beliefs*, thereby ruling out both the alarm systems example and the breathing example.

This would be a problematic move in at least two respects. First, we don't want to rule out practical reasoning or reasoning with imperatives or suppositions. But, secondly, and more importantly, why should we be so confident that the plover *has* beliefs?

Kornblith is just as confident that the plover has beliefs as that it reasons. He says:

I want to agree ... that having beliefs requires a sensitivity to reasons, and I believe that the behavior of the plovers ... shows that they are sensitive to reason. When the plover encounters a reason to stop engaging in the broken wing display, it stops; ... As I see it, meta-cognitive evaluation is only one way to show a sensitivity to reasons. (56)

What does Kornblith mean when he talks about 'responsiveness to reasons'?

He doesn't really say, but it's clear that he *doesn't* mean responsiveness to reasons *qua* reasons, but, rather, just that the behaviour in question conforms to, or is consistent with, the reasons one has. That is why he can be so confident that the plover can be responsive to reason.

However, we don't typically think that beliefs are so easily had. Once again, the example of breathing explains why. When my carbon dioxide levels are low and I respond algorithmically by breathing more deeply and frequently, I am, in the relevant sense, responding to the reasons I have. But something can be a breathing, living organism, without having any beliefs at all (consider a kitten born in a persistent vegetative state).

However, even if we ignored this difficulty and gave Kornblith the claim that the plover has not merely 'informational states' about its environment, but beliefs about it, it still wouldn't be enough. As we have already seen, with our example of the habitual depressive, content-sensitive transitions between judgements are not sufficient for reasoning.

## 5. Non-reflective reasoning in humans

So, we haven't been given a characterization of reasoning that, while being purely first-order, predicts the intuitively correct extension for the concept. And we haven't been shown a clear-cut instance of reasoning in non-human animals. But what about reasoning in adult humans?

Clearly, human beings are capable of reasoning in a way that the Taking account captures perfectly; and they often do. In such reasoning, everything is fully explicit and self-conscious: the thinker knows what premises he is reasoning from; he takes those premises to support a given conclusion, C; he takes them to do so because they bear a certain relation to C; and, finally, he infers C because of all these facts.

Call this a *fully reflective* case of reasoning.

Much human reasoning isn't fully reflective in this sense. It seems automatic, unlaboured and unreflective. Doesn't that show that although Taking might correctly characterize *some* special instances of human reasoning, it doesn't correctly characterize its *essence*: even while Taking might be contingently involved in some reasoning, it can't be *required* for reasoning. Kornblith writes:

Ben believed, when he first got into the car that he could get home by turning left on Main Street. When he got to Main Street, however, he found out that there was construction blocking the route. Seeing the construction had the effect of changing his belief about how he might get home: he no longer believed that turning left on Main would do the trick. Is Ben's change in belief rational only if he also, first, thinks about what he believes, and second, thinks about his desire for consistent and coherent belief, and finally, in light of all this, changes his beliefs in light of the previous reflections? (46)

Two considerations blunt the effectiveness of this sort of point.

Consider, first, a case in which you learn some set of rules explicitly – how to operate the elevator in your building, for example – and, for a while, follow them explicitly. After a while, though, operating the elevator may become automatic, unlaboured and unreflective: you step in, press the appropriate buttons all the while with your mind on other things. That doesn't mean that your previous grasp of the relevant rules is playing no role in guiding your behaviour. It's just that the guidance has gone '*tacit*'. In cases where, for whatever reason, the automatic behaviour gets disrupted, you will find yourself trying to retrieve the rule that is guiding you and to formulate its requirements explicitly once again.

If you ask me to give you an account of tacit guidance, I could say a few things, but perhaps not much that's very helpful. What we can't doubt, it seems to me, is the reality and ubiquity of the phenomenon, which is why we need to be thinking a lot more about tacit guidance than we have.

Once we recognize the reality of tacit guidance, I see no reason to insist that whenever a rule provides tacit guidance, it must also at some point have provided explicit guidance. In some cases it is possible, indeed likely, that a rule's guidance should always have been tacit. That would appear to be the case with the rules that guide our linguistic behaviour, our moral behaviour or, for that matter, our reasoning.

Because we know about automation and about tacit guidance, no example of the type that Kornblith invokes can do much to refute reflective accounts of reasoning.

For all that's been shown so far, then, Taking may still be correct as an account of reasoning, provided we allow that states of taking may guide our inferential behaviour tacitly.

It is natural to ask how exactly the relation between the fully reflective case and the relatively automatic ordinary case should be conceived.

What I want to say is that the fully reflective case is the Platonic Form of reasoning. It is the *ideal* form of reasoning, one that is only approximated by ordinary reasoning, in much the way in which the Platonic form of a circle is only approximated by ordinary circles.

We depart from ideal reasoning because (i) it is a costly, laborious process that we often can't cognitively afford; (ii) in various cases we don't know exactly which principles of reasoning we are operating with and it would be a matter of some difficulty to uncover them; and, finally, (iii) for a wide variety of purposes, it works just fine to use automated procedures and shortcuts.

Why should the fully reflective case be thought of as the ideal, with the tacit case a mere approximation? Why shouldn't we rather think that the tacit case is one sort of perfectly good kind of reasoning, and the fully reflective case a different kind of reasoning, each with its own virtues and

defects? Something like this alternative picture of the relation between the two types of reasoning is becoming entrenched in the suggestion that we have two different sorts of *system* for reasoning, the so-called System 1 and System 2 (see Kahneman 2011 ).

The reason for resisting this bifurcated picture derives from the fact that the aim of all reasoning is always the same: to help us figure out what we have most reason to believe. Furthermore, on internalist ways of thinking about what you have most reason to believe, you have most reason to believe something only if it is possible for you to figure out, by reflection alone, that you have most reason to believe it. And that in turn requires that, in the ideal case, all of the factors that are relevant to assessing your reasoning should be open to reflective view.

## 6. Normativity, responsibility and control

This completes my discussion of Kornblith's positive case for claiming that no correct account of reasoning need have reflective elements. Despite what he claims, he has not shown that there is some purely first-order condition that is such that, when it's added to a causal transition between informational states, or even to a causal transition between judgements, what results is a clear-cut case of reasoning. Nor has he exhibited a clear-cut case of reasoning that doesn't involve a reflective element. For all that he has shown, something like Taking is required if a causal transition is to be reasoning.

Later on in the article, I will look at Kornblith's negative case for claiming that no correct account of reasoning need have reflective elements. Before doing so, though, I want to consider whether the issue that divides Kornblith and me is purely terminological.

Let's suppose I am right that the plover does not reason and that content-sensitive transitions between judgements are not sufficient for reasoning. To what does that matter?

Suppose we introduced a new word, 'schmeasuring', which is such that it applies both to reasoning properly so-called, and to merely content-sensitive transitions. What important differences between the plover's behaviour and adult human reasoning would such a manoeuvre obscure? Which word has a better chance of 'carving epistemology at its joints'?

The main features that would be obfuscated by such a conflation would be reasoning's *normative* features. You can be *held responsible* for the way you reason; and you can be blamed for having reasoned badly and praised for having reasoned well.

What do these observations imply about the nature of reasoning? What constraints do they place on our account of the nature of reasoning?

Obviously, we are here in quite treacherous territory, something analogous to the question of free will within the cognitive domain. We cannot hope to discuss the issue here in full generality, let alone settle it.

But unless you are a sceptic about responsibility, you will think that there are some conditions that distinguish between mere mechanical transitions and those cognitive transitions towards which you can adopt a participant reactive attitude, to use Strawson's famous expression.

And what we know from reflection on cases, such as those of breathing and the habitual depressive, is that mere transitions between mental states – no matter how 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 they are transitions that a thinker *performed* that he/she can be held responsible for them. That is the fundamental reason why such transitions cannot, in and of themselves, amount to reasoning.

If you insisted on using a single word both for reasoning properly so-called and for merely content-

sensitive transitions, we would have to point out that you would have to adopt entirely different sets of normative attitudes towards the two types of processes, despite calling them by the same name.

You might think that the case of the plover shows that this train of thought is completely wrong-headed. For isn't it clear that, although the plover doesn't control its 'inferences', and so can't be held responsible for them, those transitions are still subject to substantive normative assessment? <sup>4</sup>

After all, we can regard the plover as moving from given premise propositions to a given conclusion proposition. And for any such set of propositions, we can always assess whether the conclusion *follows* from the premises, in a suitably broad sense of 'follows' that includes inductive arguments. So, of course we can assess the plover's inference as good or bad, as done well or badly, despite the plover's not having control over it.

This argument is mistaken. To assess whether a conclusion follows from a set of premises is to assess only the validity of the *argument* represented by the reasoning; it is not an assessment of the *reasoning* itself. You don't count as having reasoned well just because your conclusion follows from your premises.

As we all now know, Fermat's Last Theorem (FLT) *does* follow from the Peano axioms. But if you inferred FLT *directly* from those axioms you wouldn't count as having reasoned well unless you were also in a position to provide a chain of deductions from the Peano axioms to FLT that proceeded in steps that are 'small enough' (what we call a 'proof'). (I will be suggesting in a moment that no normal human adult could so much as infer FLT from the Peano axioms, let alone do so well.)

What does 'small enough' mean? The exact theoretical explication of this notion is still unclear (see Schechter m.s.). But a reasonable hypothesis, one that confirms Taking, is that the steps have to be small enough for a thinker to be in a position to take them to be valid (or to have an impression of them as valid).

Indeed, this very constellation of facts explains not only why certain inferences would be bad, even if valid, but also why certain inferences seem *impossible*, even though valid.

Suppose Goldbach's Conjecture to be true. Suppose someone were to claim to infer Goldbach directly from the Peano axioms, without the benefit of any intervening 'small' deductions. We wouldn't believe him. It's not just that we think he couldn't be justified in making such a transition. More than that, we feel that no such transition could be an *inference* to begin with, at least for creatures like ourselves. No one like us is in a position to take in the relation between the premises and the conclusion (we make allowances for Ramanujan).

By contrast, *any transition whatsoever* could be hard-wired in. One could even imagine a creature in which the transition from the Peano axioms to the FLT or to Goldbach is hard-wired in as a basic transition.

What explains this discrepancy between a transition that could simply be wired-in versus one that could be the result of inference. A tempting answer is that merely wired-in transitions can be anything you like because there is no requirement that the thinker be in a position to approve the transition; they can just be programmed in; by contrast, inferential transitions must be driven by a recognition (or at least an impression) on the thinker's part that the premises support the conclusion.

Another aspect of reasoning that indicates the intimate connection between reasoning and reflective control is that, even with respect to inferences which we regard as most automatic and unreflective, we take ourselves to have the capacity to reflect on them, to evaluate the principle of reasoning that they presuppose and to change our allegiance to that principle if we see fit to do so. We have the freedom to change the principles of reasoning that we employ; that is why we can be held responsible for them.

But this capacity requires second-order abilities: to reflect on what one's rules of inference are; to evaluate whether they are correct; and to amend them if need be.

Kornblith doesn't want to deny that we have such a capacity. But he does want to deny that our capacity to modify our principles of inference requires second-order capacities. He seems to think that the plover is just as capable of changing its manner of reasoning as we are. And he believes that this is illustrated by the way the plover adjusts to the established harmlessness of a potential predator in modifying its broken-wing display.

As I said above, this claim seems to me completely unwarranted. I see no case for saying that the plover is *modifying* its 'reasoning', as opposed to saying that it already had in place a complex algorithm that called for it to do distinct things under distinct input conditions.

This contrasts starkly with the capacities that we credit adult humans with. A human being may reflect on the fact that all his life he has been relying on an unrestricted version of Modus Ponens (MP); he may believe that he has come up with some decisive counter-examples to this principle in certain sorts of case (e.g. when conditionals are embedded in larger conditionals, to pick a realistic case); resolve henceforth never to reason with unrestricted MP ever again; and stand a chance of fulfilling that resolution if he is careful enough.

Each of these normative differences between adult human reasoning and the plover's transitions is well explained by Taking-style accounts but seems inexplicable on purely first-order accounts.

As far as I'm concerned, these normative features are essential to reasoning. An account that cannot vindicate them cannot be correct.

However, even if you insisted on introducing a single word that covered both the plover's mechanical transitions and human reasoning, you would have to agree that it is only appropriate to adopt a participant reactive attitude towards human reasoning and not towards the plover's transitions.

## **7. Kornblith against accounts that incorporate reflection**

Kornblith doesn't worry about much of this, or with providing an adequate first-order account, because he believes he's got something on the order of a proof that no account incorporating a reflective element could work. The 'proof', of course, centres on the threat that a reflective requirement on reasoning will lead to a vicious regress. How does Kornblith argue for this? He first maintains that on the most straightforward way of appealing to higher-order resources the result will be nothing short of a vicious regress:

... one particularly simple version of the view that reasoning requires second-order beliefs about reasons leads straightforwardly to an infinite regress. If one holds that my belief that p is a reason for believing that q, then trouble is just around the corner. This belief about reasons must surely play a role in my reasoning, for the whole point of this account is to distinguish between cases in which one merely has a belief that plays no such role, and the cases in which beliefs are genuinely involved in reasoning. If the higher-order belief is merely a by-stander to reasoning it cannot do the work of distinguishing between genuine cases of reasoning and cases in which one has a reason but is not moved by it. But once the second-order belief about reasons must itself play an active role in reasoning, the requirement that reasoning involve a higher-order appreciation of one's reasons comes into play again. So now one needs a (third-order) belief ... and so on. (45)

Now, I know something or two about regress arguments: I've made a career out of them. And there is no doubt that a danger of this sort looms for second-order accounts of reasoning. I myself have pressed

various versions of such an objection to such accounts ( Boghossian 2014 ).

But the friend of reflection-based accounts of reasoning has many options for evading such a regress.

Let's grant for the moment Kornblith's assumption that the second-order element would have to be an *explicit belief*, as opposed, for example, to something more akin to an intuition or (as I am inclined to favour) a *tacit belief*; and let's agree that if this belief is to be of any use, it has to 'have a role' in the reasoning.

Still, it wouldn't follow that the role it's got to have has to be that of *another premise*. It would be fatal to the view, of course, if its role were that of another premise, because then the relation to the conclusion would have to be inferential and a regress would be launched.

However, the obvious alternative is that the belief in question has the role not of another premise but of a background enabling condition. There is nothing obviously incoherent in the idea that a set of propositions (or beliefs) count as the premise bases for a given conclusion only when that conclusion is partly explained by a certain sort of second-order belief. Crispin Wright (2001) put the point very well when he said in response to one of my attempts at a regress argument:

It is clear how the simple internalist must reply to Boghossian. To staunch her view against all threat of Carrollian regress, she must insist that recognition of the validity of a specific inference whose premises are known provides a warrant to accept a conclusion not by providing additional information from which the truth of the conclusion may be inferred, but in a direct manner ... (79–80)

Wright is talking here about justification by inference rather than inference itself; but the point can be adapted to the topic that we are concerned with.

The possibility that Wright is describing here is even more evident if we construe the second-order element as an intuition rather than a belief. Elijah Chudnoff (2013) has a particularly interesting recent treatment of this option, one in which he attempts to answer my Lewis Carroll-style objections.

In some recent work ( 2014 ), I've explored a different way of evading a regress objection to accounts of reasoning that incorporate reflective elements. In that work, I consider the possibility that a thinker's 'taking it that  $p$  supports  $q$ ' is a tacit belief that guides the thinker's inferential behaviour in the way in which a thinker's rule of inference guides his behaviour. Indeed, the view I explore is that the two states are at root the same.

Since I believe that there are independent reasons for us to take this sort of non-inferential rule-guidance as primitive ( Boghossian 2008 ), that gives me reason to think that we get the notion of taking  $p$  to support  $q$  for free.

This is obviously not the time at which to look into the details of such a view. I mention all this merely to point out that there is no easy regress argument with which to dispatch all accounts of inference that incorporate reflective elements. <sup>5</sup>

## 8. Conclusion

I understand the powerful motives that impel Kornblith to shy away from reflective accounts of central philosophical phenomena, like reasoning. Regress problems loom for such accounts and they also seem both phenomenologically oversophisticated and unfair to children and non-human animals.

And if one has been antecedently seduced by the charms of a broadly naturalistic worldview, the case for seeking purely causal first-order accounts can seem overwhelming.

But just as it won't do to deny the reality of consciousness, simply because it's difficult to understand, so it won't do to deny the reality of reasoning.

1 Talk given at an 'Author Meets Critics' session at the Eastern APA in December 2014. I am grateful to my co-symposiasts, Hilary Kornblith and Declan Smithies and to our chair Meg Schmidt, for useful discussion both during and after the session, and also to the members of the audience.

2 'Reflection' can be used in a variety of ways and I'm not sure Kornblith is always consistent. I will use it to stand for any form of higher-order thought, regardless of the process that leads up to it. In particular, I don't intend for it to designate a process of 'reflecting on' one's thoughts.

3 I have stated the inputs and outputs of reasoning in terms of *acceptances* rather than beliefs, to allow room for the fact that one can reason from suppositions, imperatives, and so forth. For the sake of concreteness, I will follow Kornblith in working with the case of beliefs or judgements; but everything I say should apply quite generally.

4 Smithies, in his response to Kornblith, also highlights our responsibility for reasoning when it is under our reflective control, although he sides with Kornblith in allowing for the possibility of unreflective reasoning in non-human animals, and thereby shares the burden of giving an adequate first-order account of reasoning.

5 My thinking on the content of the rules guiding reasoning and on the exact relation between taking state and rule-guidance has evolved since Boghossian (2014). I hope to explain some of this in greater detail elsewhere.