

BELIEVING ONE'S REASONS ARE GOODForthcoming, *Synthese**I. The Big Picture*

I take as my starting point the evident fact that people are capable of modifying their beliefs in response to reasons in the course of deliberation. This fact is sufficient to make notions such as responsibility, blameworthiness, and praiseworthiness applicable to people with regard to their beliefs. If a state is such, and one is such, that one is capable of determining it through one's best evaluations of reasons in the course of deliberation, then even if it isn't under one's voluntary control, it is attributable to one as something for which one is appropriately held accountable. There is thus conceptual space for the possibility of one's conducting oneself poorly or well with regard to it, and accordingly for the application of praise or criticism. And there is room for an evaluation of whether one has conducted oneself responsibly or irresponsibly – that is, whether one has proceeded in a way that takes proper account of the considerations which one reasonably could have been expected to take account of, or not.

What makes these terms of evaluation properly *epistemic* in their application to states of belief is their connection, through the notion of reasons, with the notion of truth. A consideration constitutes a good reason for believing a particular proposition only if it supports or indicates the truth of that proposition. In mulling over the question of what to believe on a certain subject, we consider (if we are proceeding properly) what we take to be the reasons that there are for and against the alternative views. We evaluate them, to the best of our ability, with regard to the extent to which they support or tell against the various options. To have arrived at one's belief *responsibly* through deliberation is therefore to have attempted to arrive at the truth

by doing all that can reasonably be expected to arrive at a correct appraisal of the reasons that there are for and against. It is perfectly possible for a belief arrived at in this way to be false, and circumstances are possible in which almost every belief arrived at in this way will be false. Still, what marks this notion as *epistemic* is its conceptual connection with truth: we can't explain the relevant notion of *proceeding responsibly* without talking about epistemic reasons, and we can't adequately characterize the notion of epistemic reasons without saying something about truth.

In what follows, I will use the phrases “explicit process of deliberation” and “explicit process of inference” to characterize what takes place when one asks oneself what is the case regarding a certain issue, evaluates the available reasons, and arrives at an answer – thereby arriving at a belief.¹ Since this is not the only way in which one can arrive at a responsible belief, the notion of responsible belief must be generalized. As a starting point, we might say this: a belief is responsibly held if one has done as much as can reasonably be expected to attain a true belief (e.g., one has been appropriately assiduous in gathering relevant evidence, has not culpably neglected any relevant information, has done one's best to avoid wishful thinking, rationalization, self-deception, and other forms of irrationality, etc.) and one holds the belief in a way that is appropriately sensitive to the relevant aspects of one's background conception. The latter requirement does not entail that one must hold the belief on the basis of an argument from relevant aspects of one's background conception; it is sufficient that one's belief-forming dispositions were (and continue to be) appropriately responsive to the presence or absence of relevant reasons in one's background conception (McDowell 1998a, 1998b).

The proposed characterization of responsible belief does not explicitly preclude the possibility of responsibly holding a belief despite lacking any reason whatsoever in its favor. However, it is a common thought that at least in the case of mature adults, responsible belief also

requires the possession of reasons — that is, recognition of something as telling in favor of the belief, a recognition that makes the reason available to one's explicit deliberations about what to believe. This requirement won't be met if one merely believes something (or enjoys some non-doxastic state) that is in fact a good reason for the belief, since the reason in question won't do one any good in deliberation unless one is in a position to deploy and respond to it as a reason for the belief. So the suggested requirement is that one must recognize something which one takes or treats as a reason or is prepared to take or treat as a reason. It is plausible that the target belief won't be responsible unless one is responsible in this regard as well. If this is right, then responsible belief always requires that one responsibly hold an attitude toward something else — toward some other proposition, or some perceptual or other non-doxastic state — of *taking or treating as a reason*, or being *prepared to take or treat as a reason*. My guiding question in this paper is whether any such requirement is so much as coherent.

To take something as a reason, or to treat it as a reason, is to utilize it or be prepared to utilize it in certain ways in one's thought and conduct. So, for instance, one might appeal to it when attempting to justify the belief, or rely upon it in a course of reasoning, or recognize that it is called into question by evidence against the truth of one's belief. Unfortunately, the notion of taking or treating something as a reason is nonetheless murky; in particular, the relation between *utilizing* something as a reason and *believing* it to be a good reason is quite unclear. One issue here concerns the relation between beliefs and tendencies to accept certain inferential or justificatory transitions: do such tendencies require — or perhaps even amount to — beliefs? However this may be, believing something to be a good reason has certain *ceteris paribus* dispositional consequences for one's thought and other conduct. So while utilizing something as a reason may not require believing it to be a good reason, believing it to be a good reason will

paradigmatically involve – cases of irrationality aside – being prepared to utilize it as a reason. Accordingly, for my purposes here we can harmlessly simplify matters by focusing on the requirement that one must *responsibly believe* that certain considerations are good reasons for one's belief.ⁱⁱ If such a requirement does not cause trouble, then neither will a requirement couched in terms of the notion of taking or treating as a reason.

The issue can be refined even further. Something can't be a good reason for believing p unless it indicates, supports, or tells in favor of the truth of p, or unless it makes p probable or "likely." These terms are all an attempt to capture what I call the *evidential support relation*; for my purposes here, the only assumptions I will make about this relation are these: 1) It is somehow related to truth, and 2) It does not concern the person's epistemic status with regard to any particular proposition, belief, or claim. In what follows, I will assume that when one believes something to be a good reason for a particular belief, one at least believes it to *evidentially support* the truth of the proposition in question. Given this assumption, we can focus upon one aspect of my guiding question: beliefs to the effect that something evidentially supports a given proposition. One benefit of this limitation is that the beliefs in question will generally not be higher-order, since they do not themselves need to make reference to belief states. This is not to deny that concerns about higher-order beliefs may be raised by my guiding question, but treating distinguishable issues separately is the key to progress.

Further clarity is gained by initially focusing upon cases in which one *bases a belief upon a particular reason*. To base a belief upon a particular reason is to hold the belief for that reason or on the basis of that reason — that is, to hold the belief in such a way that the justificatory status of the belief is directly tied to that reason's adequacy as a reason for the belief. This is the relation that is at issue when we ask "On what do you base that belief?" or "What are your

reasons for holding that belief?” One paradigmatic way in which one can base a belief upon a particular reason or reasons is to arrive at the belief through an explicit process of deliberation and inference. But even when one did not arrive at a given belief in this way, one can establish a particular consideration as one’s current reason for holding the belief, for instance by appealing to that consideration in the course of a sincere attempt to justify it.

Limiting our attention to cases in which a person bases a belief upon a reason, my guiding question can be put like this. Can a requirement with the following structure possibly be correct?

In order for one to have positive epistemic status \emptyset in virtue of believing P on the basis of R, one must believe that R evidentially supports P, and one must have positive epistemic status \emptyset in relation to that latter belief as well.

Following common usage, I will call principles with this structure *Inferential Internalism*. Such principles and their close relatives figure centrally in a surprisingly wide variety of epistemological discussions.ⁱⁱⁱ Still, there is confusion in the literature about what, precisely, Inferential Internalism commits one to and whether it is so much as coherent. My aim in what follows is to clear up some of the confusion by defending Inferential Internalism against charges of incoherence leveled by James Van Cleve and Paul Boghossian and then clarifying its relation to foundationalism. I will begin by detailing some preliminary motivations for accepting Inferential Internalism, and I will close by returning to my guiding question, whether it is coherent to suppose that in order to hold any given belief responsibly, one must recognize something else as a reason for it.

Before moving on, however, I want to set aside one concern. It is sometimes complained that the Inferential Internalist requirement cannot be met by very young children and most, if not all, non-human animals. This is true. However, I see no reason to think that every important and

viable justificatory status that can be possessed by mature human beings can also be possessed by dogs or two year old children. Most mature human beings have the abilities outlined above in virtue of which appraisals of responsibility are appropriate with regard to their beliefs; very young children and (most) other animals don't. Different abilities allow for different statuses. Don't sell yourself short! If Inferential Internalism is problematic, it will have to be for reasons other than this (cf. Wright 2001, 70).

I turn, then, to some preliminary motivations for accepting Inferential Internalism.

II. Motivating Inferential Internalism

First, a bit of *prima facie* evidence from our justificatory practice.

Suppose that someone claims that his lawn has moles. Being ignorant of gardening and curious by nature, I ask why he believes this. He says, "Because it is riddled with holes, hillocks, and collapsed tunnels." I ask whether these things are good reasons for thinking that one's lawn has moles. He replies, "Oh, I haven't the faintest idea. No views on that at all."

This response is bizarre. In fact, I can't fully imagine this conversation if I assume that the speaker is a mature, competent adult in full possession of his faculties and not joking or playing around. This strongly suggests that our justificatory practice is structured by a rule to this effect: If one sincerely offers a consideration *p* in response to a request for reasons for believing *q*, then one must be (or perhaps is thereby) committed to the claim that *p* is a good reason for believing *q*. If one doesn't accept this claim, then one's attempted defense fails even if *p* is a good reason for believing *q*; it fails because it doesn't even get so far as constituting an attempted defense. I take this to be *prima facie* evidence that a belief about an appropriate evidential support relation is required for the epistemic status that we could then mark the absence of by saying, "Then you aren't justified in believing that!"

Three related sets of considerations can help explain why such a belief would be required. They are related through the broad idea that in doxastic deliberation we exercise a form of agency in which we establish our commitments through the consideration of reasons — commitments which can be expressed, defended, and modified in private reflection and in the public activity of making and justifying claims.

1. Basing Relations. The basing relation is the relation that holds between a belief and the reasons upon which the person bases it. When one explicitly infers a conclusion from particular premises or sincerely defends it via an argument from those premises, one establishes those premises as the reasons upon which one bases one's belief (cases of irrationality aside).^{iv} Whatever else is involved in a person's basing a belief upon particular reasons, it plausibly requires regarding those reasons as telling in favor of one's belief. I don't base a belief upon particular considerations, in the sense of making them *my* reasons for holding the belief, if I don't so much as regard them as reasons for the belief: how can they be *my* reasons, if I don't even regard them as reasons?^v Such considerations strongly suggest that at least part of Inferential Internalism is something like a conceptual truth: In order for one to have any positive epistemic status in virtue of believing *P on the basis of R*, one must believe that R supports P — because otherwise, one wouldn't count as basing one's belief that P upon R.^{vi}

2. Responsible belief. I have already gestured at a connection between Inferential Internalism and the notion of responsible belief. Judith Thomson (1965) develops a further point this way:

... it might be asked why a man who says, 'p, so q' must believe that p is a reason for q. ... 'Surely he must believe that p is a reason for q or he can't mean his "so".' 'So' (and its cognates) rules out a guess. But if he does not believe this then he is at best guessing. For, for all he knows, it would be an accident if q, and a stroke of luck for him if he were right in saying that q. His 'conclusion' is not a conclusion at all (296).

Thomson makes two distinct suggestions here. First, that without a belief to the effect that the premise supports the conclusion, the “so” is merely a guess: the person is not engaging in reasoning (inference) or argument. Second, that because of this fact, his position is not worthy of positive epistemic appraisal. This second claim is explicitly made by Barry Stroud, who writes that if a person did not take his premises as grounds for believing his conclusion, then the person “would be no better off, his believing what he does would be no more worthy of positive rational appraisal, than if he had simply made a lucky guess” (1977, 60 – 61).

It might be thought that both points are incorrect. A cognitive transition from one belief to another, arising from a disposition to make transitions of that type, might be claimed to count as reasoning or inference. And if the transition is one that reliably yields true beliefs given proper inputs, then it is not merely coincidence or accident (relative to proper inputs) that the person arrived at a true belief; so in that case, it might be said, the person *is* better off than if he had simply made a lucky guess.

The last point is true, but it doesn’t touch issues about responsible belief. Part of what matters for appraisals of whether someone proceeded responsibly is how matters look *from the point of view of the deliberating agent*. From that point of view, to move from premise to conclusion without taking one’s premise to support the conclusion is simply to guess. Even if the truth of the resultant belief isn’t – relative to certain facts – merely a lucky coincidence, still from one’s own point of view it would look at best like a lucky coincidence, and that’s why the transition is irresponsible.

Might the transition nonetheless be an instance of *reasoning*, or of *inference*? One’s answer will surely depend in part upon the purposes of the classification. For instance, from the standpoint of certain projects in cognitive science, the transition will clearly count as reasoning

or inference. But Thompson is right that there is something this transition won't be: it won't be an instance of a person's drawing a conclusion from particular premises in the course of reflective deliberation or explicit argumentation. It *can't* be that, if the person does not regard the premise as a reason for the conclusion.

3. *Believing as an Activity of the Agent: Parallels with Action.* Cases of irrationality aside, one's evaluations of reasons in the course of deliberation determine one's beliefs. Likewise, cases of irrationality aside, one's evaluations of reasons in the course of practical deliberation determine one's actions. This parallel is sufficient to enable us to understand ourselves as *agents* with regard to our beliefs, and it suggests that we should expect structurally parallel requirements in the two cases.

Consider, then, the case of action. The mere fact that you are entitled to perform a certain action does not guarantee that your doing so will be beyond criticism. Suppose, for instance, that you and I have gone out to dinner at an excellent restaurant. During dessert, you get up to use the restroom, and as you pass me, you say, "Go ahead and eat my chocolate terrine. I know you'll love it. It's just too rich for me tonight." Unfortunately, I don't hear a word you say because of the din in the restaurant. While you are in the restroom, I eat your dessert.^{vii} Clearly, my action is open to criticism. But why? I am, after all, entitled to eat your dessert. As various modifications of the example reveal, the problem is that I don't responsibly believe that what you said is a reason for eating your dessert.

Such examples strongly suggest the following. If I recognize that I am in a situation in which performing a certain action is unacceptable unless certain entitling conditions are met, then my performance of the action will be open to criticism unless I at least responsibly believe that those conditions are met. Applying this framework to belief, we get the following. Suppose

that one of the entitling conditions in a case in which one bases a belief upon reasons is that the reasons have to be good ones. Suppose, moreover, that one recognizes this. Then in order for one to be proceeding acceptably in basing one's belief upon those reasons, one must believe that those reasons are good, and this latter belief must be responsibly held. Consideration of the parallel with action can therefore help us see not only why the additional belief would be required, but also why it would have to be responsibly held.

So far, I have discussed Inferential Internalism in relation to the notion of responsible belief. However, Inferential Internalism is often formulated in terms of the notion of justification, and in what follows I will do so as well for reasons of convenience. I will assume that the notion of being justified, like the notion of believing responsibly, is a notion which applies in the first instance to *persons* with regard to their beliefs, and that being justified requires responsible belief, though it may require more as well. Nothing whatsoever in my discussion turns on the shift in terminology.^{viii}

III. Challenges to Inferential Internalism

Despite its evident attractions, Inferential Internalism is widely regarded as unsatisfiable. Here is one argument to that effect.

Suppose that premise P is justified for subject S, that P entails Q, and that S infers Q from P. Shall we say that Q is not justified for S unless he is also justified in believing that P does entail Q? But if so, shall we not also have to add the requirement that S be justified in believing that if P is true and P entails Q, Q is true, too? A regress impends, and to avoid it we must say that in some cases the mere *existence* of an appropriate relation between premise and conclusion, whether the subject has a justified belief about it or not, enables justification to be transmitted from one to the other (Van Cleve (1984, 560).

The key step in the argument is the claim that the Inferential Internalist must require not just the justified belief that P entails Q, but also a justified belief that if P is true and P entails Q, then Q

is true. Taken at face value, however, Inferential Internalism just says that the person won't be justified in believing that Q on the basis of P unless she is also justified in believing that there is an appropriate relation between P and Q. That requirement doesn't even entail that she must believe (let alone, have a justified belief) that there is an appropriate relation between the propositions *P is true* and *P supports Q*, on the one hand, and *Q is true*, on the other. So what is supposed to start the regress?

Here is one suggestion. Perhaps the belief, that P supports Q, could play a role in justifying one in believing Q only by figuring as a premise in an inference, in particular in the inference to Q. If that thought is right, then (by Inferential Internalism) one will need a further belief to the effect that the conclusion of this latter inference is supported by its premises, and so on. Thus when supplemented by this thought, Inferential Internalism entails that *whatever your premises are*, they never suffice to justify you in believing the conclusion unless supplemented with a further belief that they support the conclusion. That's a recipe for an unsatisfiable regress. Conclusion: there must be cases in which one can be justified in believing a proposition on the basis of an inference, even though one does not have a further belief that the premises of the inference support the conclusion.

The conclusion is overkill, because the supplementary thought should be resisted. In fact, the need to resist it is one lesson to be learned from Lewis Carroll's (1895) parable of Achilles and the Tortoise. Achilles' problem didn't arise because he wrote down his inference principle, but rather because of *where* the Tortoise had him write it down: as a premise in his inference. There is a difference between inferring q from p because one recognizes that p supports q, and inferring q from "p and p supports q". There is likewise a difference between believing q on the basis of p (because one recognizes that p supports q) and believing q on the basis of "p and p

supports q.” If one ignores this difference, then, among other problems, one simply will not be in a position to see what the Inferential Internalist has in mind.

It might be said that to be entitled to this distinction, the Inferential Internalist must provide an account of the underlying facts in which the distinction consists — facts in virtue of which a proposition or belief functions on a particular occasion as inference principle rather than premise. For instance, we might try to spell out the difference in causal-explanatory terms, suggesting that if one bases Q upon P, then the belief that P *causally sustains* one’s belief that Q, and that this distinguishes it from the belief that P supports Q. However, in terms of standard counterfactual tests for causal-explanatory relations, the two beliefs are on a par. If your rationality is not impaired and you believe that Q solely on the basis of P, then if you give up the belief that P (without replacing it with anything you regard as appropriate), you will also give up the belief that Q or at least try to shore things up. The same holds if you give up your belief that P supports Q.^{ix}

This isn’t a fatal blow to the Inferential Internalist, however. The Inferential Internalist doesn’t have to shoulder the burden of explaining the underlying facts in which the distinction consists, because the distinction in role is evident in the surface descriptions of particular cases. Consider, for instance, a student learning to evaluate arguments using a natural deduction system. His homework asks him to determine whether a certain complex argument is valid. Here there is a clear distinction of role between the premises of the argument and the inference principles deployed in the argument: the premises are propositions from which further propositions are derived, while the inference principles license or forbid the transitions through which those further propositions are derived.

This distinction is needed in the non-deductive case as well. Suppose that I adopt the following policy regarding a possible non-deductive inference to some particular proposition Q: “I will not infer Q from my premises unless and until I determine that my premises support Q.” Even an opponent of Inferential Internalism should admit that it can be perfectly coherent to adopt this policy in a particular instance. However, this policy results in paralysis unless we draw the Inferential Internalist’s distinction. So even the opponent of Inferential Internalism needs this distinction in particular non-deductive cases.

It might be thought that the only distinction needed here is the one drawn in the deductive case, on the grounds that adding “P is a reason for Q” to the premises would yield a deductively valid argument. However, “P is a reason for Q” is not equivalent to “If P, then Q;” P can support Q – even support Q very strongly – when P is true and Q is false. Consequently, “P & P supports Q; so, Q” is not deductively valid, and further supplementations of the same sort won’t get you any closer to deductive validity. This is one reason to think that the distinction is often best placed right at the beginning. In the paradigmatic case of inferring Q from P, the consideration that P supports Q is not a premise. One infers that the solution is acidic from the fact that the litmus paper turned red. One wouldn’t do that unless one thought that the litmus paper’s turning red is good evidence that the solution is acidic. This latter consideration isn’t a premise, but rather what one goes on in taking the premise to bear on the conclusion.

There is a second, related way in which Inferential Internalism can seem unsatisfiable. Consider the notion of “having a justification for believing” a proposition Q. To “have a justification” for believing Q is for something to be the case in virtue of which you possess the status of being justified in believing Q. If Inferential Internalism is correct, then one does not have a justification for believing Q if one merely justifiably believes some P which supports Q;

one must also have a justified belief that P supports Q, and so having this latter belief is *part of what constitutes one's justification for believing that Q*. When formulated in this framework,

Inferential Internalism can seem to be committed to the following two claims:

1. If one must believe X in order to be justified in believing that Q, then one's belief that X is part of one's justification for believing that Q.
2. In order to be justified in believing that Q, one must hold a further belief about all the propositions which one must believe as part of one's justification for believing that Q.

So understood, Inferential Internalism would generate precisely the problem we've been looking at: *any* set of beliefs that are said to constitute the justification one has for believing Q would be incomplete; it would have to be supplemented with one additional belief. A regress is thus generated, but not from the mistaken assumption that the belief that P supports Q must play the role of a premise in an inference.^x

The objection arises when one uses phrases such as “the justification which one has,” “having a justification,” and “one's justification” in a way that brushes over the distinction between (A) the beliefs that are referred to in a complete statement of what it is in virtue of which one is justified in believing a particular proposition, and (B) the beliefs that are referred to in a specification of the reasons upon which one bases that belief. These two sets can differ, since it is possible for one belief to be relevant to the justificatory status of another even if the latter isn't based on the former. With the distinction clearly in mind, we can see that Inferential Internalism is not committed to claim (2) above. Inferential Internalism does not say that in order to be justified in believing Q, you have to have a further belief about all the propositions that you have to believe in order to be justified in believing Q. It just says that you have to believe something about *some of them*, namely, the ones upon which you base the belief that Q.

As long as that additional belief isn't lumped in as one of the things upon which your belief that Q is based, Inferential Internalism gets a clean bill of health.^{xi}

At the end of the day, this objection shares a fundamental mistake with the first objection considered above: both fail to distinguish between believing *q* on the basis of *p* and believing *q* on the basis of the complex reason, *p and p supports q*. The Inferential Internalist insists that these are to be distinguished. One indication that the Inferential Internalist is right is the structure of our ordinary justificatory practice. If someone offers *p* as a reason for believing *q*, we can correctly ask (given appropriate circumstances): "Do you really believe that *p* supports *q*?" Given an affirmative answer, we can (given appropriate circumstances) correctly pursue further questions, such as, "But does *p* really support *q*?" and "Why do you think that *p* supports *q*?" But there are no circumstances in which it would be correct to probe further by asking, "But do you really think that *p and p supports q* supports *q*?" Asking that question would indicate a failure to appreciate what the person was offering as a basis for her belief. The matter is completely different, however, if the person initially offers *p and p supports q* as her reason for believing *q*. In that case, one can (given appropriate circumstances) correctly ask precisely this question.

If the Inferential Internalist is right that the connecting belief (that *p* supports *q*) does not serve as a premise in the inference and is not part of what one bases one's belief upon, what role does it play in the creation of a justified belief? What is its contribution? This question generates what is perhaps the most important argument for the unsatisfiability of Inferential Internalism, an argument due to Paul Boghossian (2003).

Boghossian is concerned with a case in which one explicitly infers a conclusion via Modus Ponens. Suppose that one has knowledge or justified belief that a Modus Ponens

argument comprising particular premises (call them “(1)” and “(2)”) and a particular conclusion (call it “(3)”) is deductively valid. How, Boghossian asks, does this belief contribute to the transmission of warrant from the premises to the conclusion through the process of inference? How does it “bear on my warrant to infer [the conclusion]?” It can do so, he claims, only via inference, as follows.

- i. The inference from (1) & (2) to (3) is valid.
- ii. If an inference is valid, then anyone who is justified in believing its premises and knows its validity is justified in inferring its conclusion.
- iii. So, anyone who is justified in believing (1) & (2) is justified in inferring its conclusion (3).
- iv. I am justified in believing the premises (1) & (2).
- v. So, I am justified in inferring (3).

Obviously, this reasoning involves Modus Ponens inferences – inferences of the very sort in question. Boghossian therefore concludes that a “fatal circularity” arises from Inferential Internalism: it cannot explain warrant-transmission across a Modus Ponens inference without already assuming that warrant does transmit across such inferences; “the very ability we are trying to explicate is presupposed by the internalist account on offer” (233).

To begin to unravel Boghossian’s challenge, notice first that the charge of circularity is a red herring, since it is irrelevant that the inferences in question are themselves instances of Modus Ponens. To raise trouble, the opponent of Inferential Internalism only needs to claim that the belief that the inference is valid can “bear on one’s entitlement to draw the conclusion” only via an inference (of whatever form). For then, regardless of the form the additional inference takes, it can’t do its work unless one has a justified belief that *it* is valid. But then that further belief can only bear inferentially on one’s entitlement to draw *that* conclusion, and so on. A vicious infinite regress is thus launched by the combination of Inferential Internalism and the claim that the belief required by Inferential Internalism bears only inferentially on one’s

entitlement to draw the conclusion. In order to arrive at any justified belief through a process of inference, one would first have to perform an infinite number of inferences. This problem would arise for both the deductive and non-deductive cases.^{xii}

Boghossian's argument is fueled by the assumption that the connecting belief's role is to serve as the premise for an inference to the conclusion "I am justified in inferring [or believing] *p*."^{xiii} The underlying idea here is that the Inferential Internalist thinks the additional belief is required because a responsible believer regulates her beliefs by determining whether they are (or would be) justified. On this interpretation, the Inferential Internalist is thinking like this: "I won't be justified in believing *p* on the basis of an inference from *q* unless I have a justified belief that I would be justified in believing *p* on the basis of an inference from *q*. And I won't be justified in believing *p* on the basis of the inference unless I believe it *because* I recognize that the inference entitles me to believe it." Warrant transmits to the conclusion of an inference, on this picture, just when (and because) one recognizes that warrant can transmit through the inference and so performs it. As Boghossian puts it, "The simple Inferential Internalist insisted that the inference's justifiedness be transparent to the thinker – the thinker has to be in a position reflectively to appreciate that his inferring this conclusion from these premises is justified" (236).

This is not the Inferential Internalist's story. For one thing, Boghossian's story commits the Inferential Internalist to holding that one can't be inferentially justified in holding a particular belief unless one also has (justified) beliefs that one is or would be justified or responsible in holding that belief (under certain conditions), that holding that belief is or would be warranted or permissible, or something to similar effect. This is what we might call a "cross-level requirement." In the most general terms, such a requirement would state that one's belief that *p* can't have status \emptyset unless one has a belief with status \emptyset about one's belief that *p* and its status.

But Inferential Internalism, as I have characterized it, only demands that one have a belief with a certain status about the evidential relation between the propositions which comprise the premises and conclusion of the inference. That demand does not generate a cross-level requirement. And that's all to the good, because the fully general cross-level requirement generates an obvious and straightforward infinite regress: if any given belief can't have status \emptyset unless one has a belief with status \emptyset about that belief and its status, then the same requirement will apply to that further belief, and so on *ad infinitum*. An Inferential Internalist therefore should reject Boghossian's suggestion that the crucial role played by the connecting belief is to serve as the premise for an inference to a conclusion about the status of one's belief or the permissibility of drawing a certain inference.

Still, Boghossian's overarching question deserves an answer. How, according to the Inferential Internalist, does my belief that the premises support the conclusion "bear on my entitlement to draw the conclusion"?

To see the Inferential Internalist's answer, we first have to distinguish two questions. First, what role does the connecting belief play in a fully explicit process of inference? Second, how does one's possession or lack of a connecting belief determine whether or not one is justified in believing the conclusion?

To approach an answer to the first question, consider what takes place in a fully explicit process of deliberation. You wonder whether or not q is true. You believe that p is true, and it occurs to you that p strongly supports q . You don't notice anything that would defeat p 's support for q . Here, then, is how your explicit reflection would go. "P supports Q. Nothing defeats its support for Q. P is true. So, Q is true." Being rational, you thereby come to believe Q. In this process of deliberation, the objects of your explicit reflection are P, Q, and their evidential

relations, not anything about yourself and your beliefs. It is true that you perform an inference to the conclusion that Q. But though your process of inference and belief-formation could be an explicit focus of your deliberation, it needn't be; to form the belief that Q through an explicit process of inference it is sufficient that you reflect on the content of the relevant reasons and what they support.

Once we see this point, we can also see that the Inferential Internalist can retain a regulatory role for the connecting belief in the process of deliberation. Explicit deliberation is a process through which we regulate our beliefs in the light of reasons. To cash out the metaphor, it is a process through which we decide what to believe by considering (to the best of our ability) the available reasons and what they tell for or against. Of course, we don't generally ask ourselves, "What should I believe?" but rather, "What is the case?" The process of considering the available evidence and drawing a conclusion *is* the process of deciding what to believe. Insofar as we are rational, this process determines what we believe. Indeed, that is its point, and that is why it is properly called a process of "regulating" our beliefs.

There are two crucial points to notice here.

First, the process of belief regulation that I have just described does not *require* one to believe that one is or would be justified in believing the conclusion of the inference, though it may on occasion involve such beliefs.

Second, the role played by the connecting belief in such cases need not itself be inferential. The transition from belief in the premises to belief in the conclusion is an inference. But one need not infer anything from the connecting belief in order to make the transition from premises to conclusion in a way that is informed by it. Instead, one takes account of its content in deciding what the premises indicate to be the case. Utilizing a consideration as a premise for

an inference is just *one* way of taking that consideration into account in the course of deliberation.^{xiv}

Let us turn now to the second question. What does the connecting belief do to make one justified in believing the conclusion of the inference? What is its contribution to determining justificatory status? To make headway here, it is first necessary to remove a distortion introduced by Boghossian's imagery. Boghossian asks how warrant "transmits" through or across the inference. The imagery suggests that warrant (or justification, or what have you) is something like a current of water that flows through the inference, from premise beliefs to conclusion — but only under certain circumstances. The Inferential Internalist's proposal, *on this picture*, is that the warranted (justified, etc.) belief that the premises support the conclusion is like a sluice gate, or perhaps the sluice itself, permitting the current to flow to the conclusion. Boghossian's question, then, is prompted quite naturally, for it looks as though there is something here that cries out for explanation. What does the belief *do* that permits warrant to flow to the conclusion, since it wouldn't flow without it?

The desire for an explanation is a product of the imagery. Here's an alternative. Think of being justified as a status possessed by a person with regard to a particular belief, a status which accrues to the person when, and because, particular conditions are met. In the case in which a belief is arrived at through an inferential process, one of those conditions is that the beliefs which served as premises must themselves have the status in question. The Inferential Internalist maintains that a further condition must be met: one must believe that the premises support the conclusion, and this belief too must have the status in question. The Inferential Internalist traces this condition in part to the demands of responsibility, maintaining that a person can't responsibly draw a particular conclusion through a process of inference unless she

responsibly believes that the premises support the conclusion. On this picture, there is no mystery about what the connecting belief “does” to enable warrant to transmit across the inference. It doesn’t do anything to enable warrant to transmit, because warrant doesn’t “transmit.” The connecting belief is simply a necessary condition on responsible inferential belief, and so too on justified inferential belief.

Of course, it’s not enough that one merely *have* a justified connecting belief; one could have the belief and yet it might not play any role at all in the deliberations through which one arrived at the belief. In that case, one’s deliberations won’t have been responsible. For this reason, the Inferential Internalist is likely to add a further requirement: if a belief is arrived at through a process of explicit deliberation and inference, it won’t have been arrived at in a responsible manner unless one took proper account of the content of a responsibly held connecting belief in the course of one’s deliberation. This requirement can be generalized to cover cases in which a belief is held on the basis of particular reasons even though it was arrived at without any explicit process of deliberation and inference. What is required in such cases is that the connecting belief play the right sort of role in a commonsense psychological explanation of why the person believes as she does. This explanatory requirement doesn’t amount to the connecting belief’s “doing something” to enable warrant to transfer; it can instead be viewed as an additional condition a person must meet in order to be justified in holding a belief on the basis of particular reasons.

From this perspective, Boghossian’s explanatory question is removed; the explanatory question that remains is of a rather different sort: why is having and taking proper account of a connecting belief a necessary condition for responsibly holding a belief on the basis of particular

reasons? Here, the Inferential Internalist will appeal to considerations of the sort adduced in the previous section.

IV. Inferential Internalism, Foundationalism, and the A Priori

Suppose, then, that Inferential Internalism is structurally coherent. What are its broader implications, if any, for theorizing about justification?

It is widely held that if one accepts Inferential Internalism, then one must also accept a foundationalist account of justification. This is incorrect, however, and seeing why will clarify the prospects for the account of epistemic responsibility sketched in section 1.

The argument linking Inferential Internalism to foundationalism goes like this. According to Inferential Internalism, if one is to be justified in believing q on the basis of p , then one must also have a justified belief that p supports q . Call that latter belief S . How is S justified? There seem to be two options: inferentially or non-inferentially. If inferentially, then there must be some belief Z from which S is inferred. But then one must have a justified belief that Z supports S . Call that belief \emptyset . How is \emptyset justified? Inferentially or non-inferentially... and we're off to the races. If the justification is inferential in every case, then either an infinite chain or a circle results. It's doubtful, however, that any of our justified beliefs are part of an *infinite* chain of beliefs organized in such a fashion, and a circular chain of this sort seems unacceptable. Further problems arise for each option if we think of each of the beliefs as having been arrived at through a process of inference. Since inference is a temporal process, we cannot have performed an infinite number of inferences in order to arrive at the belief in question. It is equally impossible to arrive at a belief through a process of inference from that belief itself, which is what a circular chain of inferences would require. So it seems that we must be justified

in holding at least some beliefs about relations of evidential support without holding them upon the basis of any other considerations.^{xv}

The standard conclusion is that if Inferential Internalism is correct, so is a version of foundationalism: at least some of our beliefs about relations of evidential support must have a positive justificatory status that does not depend upon, require, or involve the positive justificatory status of any other beliefs. That is to say, at least some of these beliefs must be *immediately justified*.^{xvi} Moreover, since we don't have sensory experiences of one thing's being a reason for another, there is no plausible story about how beliefs about evidential support relations could be both *immediately justified* and also *empirical*. So there is general agreement: if one accepts Inferential Internalism, one is committed to a foundationalist story involving a priori insight into relations of evidential support (Fumerton (2004), Wright (2001, 72 ff.), Bonjour (1998)).

However, there are good reasons for wondering whether Inferential Internalism really requires such a view.

For one thing, the existence of immediately justified beliefs is arguably incompatible with the account of responsible belief mooted in section 1. That account held that one can't responsibly hold a given belief unless one responsibly takes or treats something else as a reason for it or is prepared to do so. If (as seems initially plausible) this requirement demands that one have a responsible belief about a relevant reasons-relations, then there can't be any immediately justified beliefs at all. So if Inferential Internalism requires immediately justified beliefs, then Inferential Internalism would appear to be *in tension with* the broader account of epistemic responsibility sketched in section 1. That would be a startling result. For it is hard to see how to accept that sort of account without accepting something like Inferential Internalism. So if

Inferential Internalism requires immediately justified beliefs, then that broader account seems not to be fully coherent. This argument needs to be probed at numerous points. One of them is precisely the claim that Inferential Internalism requires immediately justified beliefs.

Moreover, whatever one thinks about the possibility of a priori knowledge, it is extremely plausible that non-logical relations of evidential support depend upon how the world actually happens to work. But Hume was right about this much: we don't have purely rational, a priori insight into contingent nomological relations. So is it simply impossible to accept Inferential Internalism while also granting the contingency of non-logical evidential support relations?

In fact, the argument from Inferential Internalism to rationalist foundationalism is mistaken. There is another option. It's one thing to say that some beliefs are justified without being held on the basis of particular considerations; it's quite another thing to say that some beliefs are justified in a way that does not depend upon, involve, or require the positive justificatory status of any other beliefs. To break the threatened regress, you only have to say the former. Foundationalism is overkill.

Terminological confusion can obscure this point. The term "inferential justification" can be used in ways that skim over the differences between: (1) a justificatory status which requires the belief to have been arrived at via an explicit process of inference, (2) a justificatory status which requires that the person hold the belief on the basis of appropriate beliefs, (3) a justificatory status which simply requires that the person also hold some other justified beliefs. The regress argument shows at most that not every belief can have a justificatory status of sorts (1) or (2), given the temporal and normative constraints involved. But that's not enough to show that there must be immediately justified beliefs. One could still claim that no beliefs can be justified unless others are as well. That claim – the claim that all justified beliefs are

“inferentially justified” in sense (3) – is all that it takes to deny that there are immediately justified beliefs.^{xvii}

The distinction I have been stressing can be usefully illustrated in terms of a particular account of the basing relation. I have argued elsewhere (Leite 2004) that the mere fact that a person’s processes of belief formation and maintenance responded to certain considerations does not suffice to make it the case that the person based or bases the belief upon those considerations. This is because basing a belief upon particular reasons requires certain commitments on the part of the person, and the causal-explanatory considerations that figure in cognitive-psychological accounts of belief formation and maintenance do not suffice to establish a person’s commitments. If these claims are right, then the following sort of case will be perfectly possible: a belief might be formed and maintained in a way that is responsive to some of one’s other beliefs, even though one has not (yet) based the belief upon particular reasons since one has not taken on the relevant commitments. Seeing this possibility can help make sense of the thought that possession of (and appropriate sensitivity to) other responsibly held beliefs may be a necessary condition on responsible belief even in cases in which a belief is not held on the basis of particular reasons. As an illustration, consider my belief that I live in the United States. I believe many things that tell in favor of the truth of this belief. But it doesn’t seem quite right to say that I base this belief on any of them in particular, nor that I base it upon all of them. Still, this belief is both responsibly held and justified, and it seems that my background information will play a role in any adequate account of why that is so. The conceptual possibility of a view along these lines shows that even if responsible belief requires reasons, it needn’t therefore require that one currently base one’s belief upon them.

The following two claims are consequently perfectly consistent:

1. Because of the demands of epistemic responsibility, no belief can be justified unless one has some other justified beliefs.
2. Some beliefs can be both responsibly held and justified without the person's basing them on particular considerations.

The possibility of combining these two claims shows that Inferential Internalism doesn't require immediately justified beliefs about relations of evidential support. It just requires that some beliefs about evidential support relations count as justified without being currently held on the basis of particular reasons. That isn't to allow that such beliefs can be justified even if one entirely lacks reasons in their favor; as in the account of responsible belief sketched in section 1, the demand for reasons in their favor might be precisely what motivates accepting claim (1) above with regard to them.

The distinction we have been looking at also helps clarify the extent to which Inferential Internalism forces an appeal to the *a priori*. It is perfectly compatible with (1) and (2) to claim that one can't have a justified belief about non-logical evidential support relations unless one has some other justified *empirical* beliefs. We can consequently accept Inferential Internalism and yet preserve the thought that the relevant beliefs about evidential support aren't justified through some sort of *a priori* insight into non-logical matters.

The resulting view has a great deal of plausibility. Consider your belief that the fact that a pencil was dropped is good evidence that it will fall, or that a pencil's piercing your hand is good reason to believe you are in pain. These are both beliefs about relations of evidential support. They are responsibly held and justified. They are not *a priori*. But neither was arrived at through a process of explicit inference. Nor, if you are like me, are these beliefs currently held on the basis of any particular reasons. (What are they?^{xviii}) This isn't to deny that you have a great deal of empirical evidence in favor of these beliefs and little or none against them, nor that

these beliefs were formed in a way that is appropriately sensitive to this background evidence. But these facts are perfectly compatible with the suggestion at hand. In fact, they can help us understand how these beliefs could be responsibly held and justified.^{xix}

The possibility of this sort of view shows how the Inferential Internalist can evade a second argument of Boghossian's (2003). Boghossian urges that in many cases, one's knowledge about the relation between premises and conclusions will be *general*: one will know, for instance, the general truth that all instances of Modus Ponens are valid. To deploy this knowledge in the particular case, however, one will have to reach the conclusion that the particular inference at hand is valid. To do that, one will have to perform an additional inference, reasoning as follows:

1. Any inference of the Modus Ponens form is valid.
2. The particular inference at issue is an instance of Modus Ponens.
3. So, this particular inference is valid.

By Inferential Internalism, however, one must know or justifiably believe that this additional inference is itself valid, which in turn requires yet another inference from the general principle that instances of Modus Ponens are valid, and so on *ad infinitum*. Boghossian consequently claims (2003, 233) that Inferential Internalism is hopeless unless we can have immediately justified a priori beliefs to the effect that *particular arguments* are valid (or, to generalize the claim, to the effect that particular premises support particular conclusions).

This argument misses precisely the option I have been exploring. Boghossian is right that the Inferential Internalist is in trouble if every justified belief about the validity of particular arguments must be derived via inference from general principles. In fact, the Inferential Internalist will be in trouble if no such belief can count as responsible or justified unless it is *based on* a general principle. The Inferential Internalist must therefore allow that some beliefs

about the validity of particular arguments can count as responsible and justified even if they were not arrived at through an explicit process of deliberation and inference and are not currently held on the basis of any particular reasons. But this does not entail that some such beliefs must be immediately justified, since it could be claimed that such beliefs won't be responsible or justified unless one's background conception includes considerations which tell in their favor.

To see how this might work out in detail, consider the following argument:

1. Xanthippe is a human being.
2. If Xanthippe is a human being, she was born.
3. So, Xanthippe was born.

Is this argument valid? The merest glance – without any explicit inference – suffices to recognize that it is; you simply look and form the belief. Given this fact, it is doubtful that you already base this belief upon some particular reason.

This suggestion can be reinforced in the following way. There are at least two possible ways in which one might attempt to justify this belief. First, one could appeal to an argument from the premises “If an argument is of the Modus Ponens form, then it is valid” and “This argument is of Modus Ponens form.” Second, one could appeal to the truth-table for the conditional, assigning truth-values to the particular sentences (1) – (3). These two justifications are independent, in so far as one could offer the former without being able to offer the latter (e.g., if one has been trained to recognize the Modus Ponens form and taught that all inferences of this form are valid, but one doesn't know how to determine validity oneself via the use of truth-tables), and one could offer the latter without being able to offer the former (e.g., because one has no belief at all as to whether all Modus Ponens inferences are valid). You are presumably someone who could offer *either* of these reasons for believing that the inference is valid. But you haven't yet based the belief on either one, nor have you based it on both. You simply

identified the argument as valid. Given the spontaneous way in which this belief was formed, it seems that you haven't yet based it on particular reasons at all.

This is enough to evade the regress that Boghossian fears. If your belief about the argument's validity was not arrived at through an explicit process of inference and is not held on the basis of particular reasons, then the Inferential Internalist demand hasn't kicked in, and so the regress hasn't started. But we haven't said that this belief is immediately justified: we haven't said that its positive justificatory status does not require or depend upon that of any other beliefs.

What's the alternative? Just what we would expect the Inferential Internalist to say, given the approach to epistemic responsibility proposed in section 1. Though your belief about this argument's validity isn't (yet) held on the basis of particular reasons, you possess reasons which you take to tell in its favor. The suggestion, then, would be that your possession of reasons is a necessary part of this belief's positive status; it is responsibly held only because you believe other things which you take to tell in its favor.

There is a tempting objection to the argument so far. In the above example, you spontaneously formed a belief that a particular argument was valid. It is plausible that this belief was formed through the deployment of an acquired recognitional capacity that enables you to recognize certain arguments as valid by responding to their structures. This capacity doesn't operate on *nothing*, as it were; there must have been features of the argument to which your belief-forming mechanisms responded. So wasn't this process of belief formation itself inferential? And so isn't the Inferential Internalist in trouble after all?

The short answer is, "No." Not all cases of concept application can be understood in inferential terms, for reasons that have nothing to do with Inferential Internalism.

If concept application always involved an inferential process, the inference would be of this type:

1. It is sufficient for being an F that a thing have features a, b, c,
2. This thing has a, b, c,
3. So, this thing is an F.

However, not every instance of concept application can be understood as resulting from this sort of inference. For on this model, concept application would be impossible: the second premise requires the application of further concepts to the object, and so an infinite number of inferences would be required before any concept could be applied at all. Some instances of concept application therefore must consist in deployments of recognitional capacities which are not processes of inference and do not involve basing one's belief upon particular reasons.

This last point does not commit us to the existence of immediately justified beliefs. It may be that the beliefs resulting from the operation of recognitional capacities won't be justified unless one has appropriate justified background beliefs. If that were so, it wouldn't make all cases of concept application inferential, but it would prevent the resultant beliefs from being immediately justified. Since (as I have argued) this proposal is perfectly compatible with Inferential Internalism, Inferential Internalism does not force us to accept that there are immediately justified beliefs, let alone immediately justified a priori beliefs about relations of evidential support.

V. The Big Picture Again: a preliminary sketch

I now return to the question with which I began. Is it coherent to suppose that no belief can be responsibly held unless one also responsibly takes or treats (or is prepared to take or treat)

something else as telling in its favor? As before, I will begin by interpreting this requirement as demanding a responsibly held belief about evidential support relations.

It should be clear that this question is distinct from the question of the coherence of Inferential Internalism. One can take p to evidentially support q , believe both p and q , and yet not base q upon p . (For instance, this is the position I am in with regard to my beliefs that my driver's license shows an address in the United States and that I live in the United States.) The question at hand, then, is whether it is coherent to suppose that even in cases where you don't (yet) base your belief upon particular consideration(s), responsible belief always requires that you responsibly believe of something else – some consideration, or some non-doxastic state – that it evidentially supports the truth of the proposition at issue.

The obvious worry here is the threat of regress. Take any particular responsibly held belief. By the above requirement, one must also have a responsible belief that a relevant evidential support relation obtains. This in turn will require at least one more responsibly held belief, to the effect that some particular reason provides support for the truth of the preceding belief. And so on. If we assume, moreover, that reasons for beliefs about evidential support relations will always include other beliefs, the number of beliefs required quickly becomes staggering. It seems that either a circle (or, more broadly, a web) will be involved, or the set will have to be infinite.

Taken by itself, however, this situation need not be problematic. This can be seen by imagining that the set of beliefs is simply given and that the person has not (yet) based the relevant beliefs upon any particular considerations. In that case, the existence of a complex web of beliefs will not itself give rise to an objectionable circularity. The relations of evidential support may run every which way amongst the believed propositions, but the person will not

have engaged in a circular course of reasoning, nor will he have based any of the beliefs upon circular arguments. Similarly, if no basing relations are yet established, then no difficulty is introduced if an infinite number of beliefs is involved – so long, that is, as there is no independent difficulty with the suggestion that we have an infinite number of beliefs.^{xx} It is only when beliefs are arrived at through inference or held on the basis of particular reasons that temporal or normative ordering relations are imposed, and without such relations, no problematic regress is generated. There is consequently no inherent incoherence in the idea of a set of beliefs which includes, for each belief in the set, a belief to the effect that something else tells in its favor. As should be plain, this point can be integrated with the facts that we sometimes arrive at beliefs inferentially and sometimes base our beliefs upon particular reasons: in these cases, the ordering which is imposed is at the periphery of an appropriately large mass of beliefs amongst which no such ordering has been imposed. The chains of reasons terminate with beliefs which one does not base on particular reasons, but which one responsibly believes nonetheless.

Two problems remain. This can be seen most easily by adopting (for the sake of argument) a conception of belief as a conscious or unconscious occurrent representational state which assertorically represents something as so.

Here's the first problem: on this conception of belief, it is doubtful that we always satisfy the demand for a further belief about evidential support relations. Consider the example regarding validity used toward the end of the preceding section. You spontaneously formed the belief that the argument is valid. You were also in a position to provide considerations which told in favor of the truth of that belief. But did you – at the moment when you came to believe the argument valid – *also* believe that those considerations indicated that the argument is valid? For instance, did you then also form the belief that the argument, “All Modus Ponens arguments

are valid; that argument is an MP argument; so it is valid,” is *itself* valid? Of course, if you are asked, “Do you believe that those considerations about Modus Ponens show that this particular argument is valid?” you will answer, “Yes.” But the question here is whether you formed that latter belief *at the very moment (or before) you formed the belief that the first argument was valid*. It seems doubtful that you did — if possessing a belief is a matter of actually having an occurrent representation to the effect that something is the case.

This is not to deny that there is some inclination to say that you did have that belief. However, what drives this inclination is the thought that you had certain *dispositions* arising from certain recognitional and inferential abilities. This indicates that it is the recognitional and inferential abilities that are primary here. Insofar as you count as taking or treating something as favoring the truth of your belief that the argument was valid (or as being prepared to do so), this is in virtue of your possession of an appropriate network of recognitional and inferential abilities.^{xxi}

A similar point can be brought out by considering a second problem. Suppose that someone responsibly believes that *p*, and also that *p* supports *q*. In order to responsibly form the belief that *q* through a process of inference from *p*, the person must do so “in the light of” her recognition that *p* supports *q*. That means that she must recognize that the claim that *p* supports *q* is relevant to the question of the truth of *q*, given that *p* is true (and she must form the belief that *q* in the light of this recognition). However, we cannot think of this recognition as consisting simply in the presence of an occurrent representation. For if we think of the recognition in this way, then it won’t do her any good; she will also need to have another occurrent representation to the effect that the previous one is relevant, and so on, *ad infinitum*. Instead, what is needed here are recognitional and inferential abilities, abilities which do not

consist simply in the possession of particular representational states, but which instead enable her to move amongst them.^{xxii}

Both problems point in the same direction. What is crucial for responsible belief is that one have the dispositions and abilities that are captured by talk of “taking” or “treating” something as a reason. The responsible believer is able to operate intelligently and thoughtfully with reasons in the course of deliberation and self-critical reflection upon her beliefs – to deploy considerations in defense of her beliefs, to consider challenges, including challenges to those considerations’ relevance or adequacy, and to modify her beliefs accordingly. This requires possession of a whole host of abilities: abilities to recognize the relevance of certain considerations, what they tell for or against, what might tell for or against them, etc., and to respond appropriately. These abilities cannot simply be seen as consisting in and arising from a set of occurrent representational states.

At the same time, however, one would not be a responsible believer if these dispositions and abilities always operated “blindly,” without generating conscious beliefs about the matters they concern. It is through conscious, occurrent beliefs about reasons-relations that we can engage in reflective self-criticism, subject our beliefs and our reasons to the scrutiny of others, and bring the rational modification of our beliefs under deliberative control. Without such beliefs and the further capacities that come with them, we could not appropriately be held accountable for our beliefs at all.

I conclude, then, that it is at least coherent to maintain that one can’t hold any given belief responsibly unless one responsibly takes or treats something else as telling in its favor. The key is to recognize that at the end of the day, it is one’s recognitional and inferential

capacities — and the belief ascriptions that they license — that do the heavy lifting in enabling satisfaction of this requirement.

ⁱ I assume that consciously arriving at a belief in this way is an activity engaged in by *the person*, and so not merely epiphenomenal, in this sense: it is not merely a shadow cast in awareness by a process that could just as well run its course without such awareness. It is therefore not merely a process of belief formation of which one is aware. This is not to deny that a person's engaging in this activity may depend upon, or even consist in, processes taking place within the person of which the person is not conscious.

ⁱⁱ Or at least that one must responsibly believe something to that effect. I prescind here from worries about the precise degree of conceptual sophistication required in order for one to have the relevant beliefs, as well as from the question of whether some form of *de re* belief is required or sufficient. Though important, such questions are irrelevant to the main issues I want to pursue.

ⁱⁱⁱ Here's a sampling: recent "bootstrapping" arguments against reliabilism, discussions of epistemic circularity and knowledge of the reliability of the senses, of foundationalism and the *a priori*, of our entitlement to accept fundamental logical principles, and of skepticism about induction and about knowledge of the external world.

^{iv} For discussion of this point, see Leite (2004, section III). (The discussion in this paragraph relies heavily upon sections II – IV of that paper.)

^v This isn't to say that I will readily acknowledge that they are reasons, or that they are my reasons; irrationality can wreak havoc in this territory, as when repression leads one to deny, falsely, that one takes certain considerations as reasons, or self-deception leads one to deny that one's belief is based on particular considerations upon which it is in fact based.

^{vi} For an arguments along similar lines, see Stroud (2000, 40). Robert Audi (1993) concurs, though he weakens the requirement in order to render it applicable to very young children.

^{vii} This wonderful example comes from Ginet (1985).

^{viii} The differences between the notions of responsible and justified belief allow for the possibility that a full-fledged Inferential Internalist principle for justification is incorrect; perhaps all that is required is *responsible* belief about the evidential relation. For my purposes here, we can prescind from such matters of detail, since I will mainly be focused on issues arising from the structure of Inferential Internalist requirements that would apply equally to a requirement couched only in terms of responsibility.

^{ix} Michael Huemer (2002, 333) considers and appropriately rejects such an account. He incorrectly takes its failure to be a fatal blow against the Inferential Internalist.

^x Huemer (2002) offers a version of this argument. He writes, "to say that I am justified in believing P on the basis of E is to imply, among other things, that E is an adequate basis for P. If E itself is not an adequate basis for P but rather F must be added in order to form an adequate basis for P, then E, strictly speaking, does *not* justify P; it is only (E & F) that justifies P." But Inferential Internalism tells us that to be justified in believing P on the basis of E, I must believe

two things: E, and that E makes P probable. “These two beliefs are not two separate sources of justification for P that I might have; they are both held to be components of the same justification for P. I must have both of them at the same time, or I have no justification for P... [So] E is decidedly not an adequate justificatory ground for P; it is only the conjunction of E with *E makes probable P* that can give me adequate justification for P.” So, he claims, Inferential Internalism leads to a reductio ad absurdum. “[F]or any P and E, if I am justified in believing P on the basis of E, then I am not justified in believing P on the basis of E after all ... [So] I am not justified in believing anything on the basis of anything” (332).

^{xi} Huemer’s argument makes this mistake. He uses the phrases “justification for P,” “adequate justificatory ground for P,” and “justified in believing P” in a way that systematically elides the distinctions that Inferential Internalism insists upon. One thing can be an *adequate basis* for believing another all on its own – in the sense that it can be a basis which provides the sort or degree of evidential support which is requisite for having an inferentially justified belief; its adequacy consists in the fact that it is *good enough evidence*, or provides *good enough* support. It doesn’t follow, however, that in such a case one will be justified in believing that latter thing merely in virtue of believing it on the basis of the former, even if one is justified in believing the former. To say that there are other requirements which one must meet – even requirements which demand that one believe *other* things – is not to deny the claim that the one is an *adequate basis*. So if we clearly distinguish the notions of *having evidential support for P*, the belief that P’s *being based upon a good reason*, and a person’s *being justified* in believing P, Huemer’s argument fails.

^{xii} The charge of circularity gets its bite from the worry that the infinite regress is unsatisfiable. In general, it can be informative to learn that no fully general account of some phenomenon X is possible because explanations of what it is for any particular thing to be X must assume – as part of the explanation – that other things are X. (This would just show that one can’t provide a non-circular conceptual analysis; one might perfectly well be able to explain why, and in virtue of what, any given instance of X is such.) The situation becomes problematic only if it appears that the demands involved in the explanation of the particular cases can’t possibly be met. That’s what happens with the infinite regress Boghossian shows us how to generate.

^{xiii} In an earlier formulation of the argument, Boghossian (2001) suggested that this inference has as its conclusion the very proposition which was the conclusion of the initial inference. Taken in this way, however, Boghossian’s argument assumes that the connecting belief must be imported into the premises of the argument. It is consequently a version of the mistake brought to light by Lewis Carroll’s parable. This is pointed out by Wright (2001).

^{xiv} It’s plausible that *some* underlying process takes place here, and it may be useful for certain theoretical purposes to model this process as a series of transitions between representational contents analogous to a chain of arguments. However, even if that is so, it is not relevant in the present context, where the issue is requirements pertaining to the notion of inference as it figures in commonsense psychological explanation. It is part of commonsense psychological explanation that in deliberation we can take account of a consideration without utilizing it as the premise for an inference.

^{xv} For a representative statement of the argument, see Fumerton (2004, 162).

^{xvi} Fumerton, for instance, clearly indicates that he is using the term “inferential justification” in such a way that he takes the foundationalist conclusion to follow: “a belief P is inferentially justified if its justification is constituted by the having of at least one belief other than P. A

belief is noninferentially justified if its justification does not consist in the having of any other beliefs” (1995, 56).

^{xvii} This option can also be obscured by the conception of justification as a quality which “flows” from belief to belief. On that conception, it will appear that the “flow” must start somewhere: there must be springs which give rise to the torrent, and those springs are the immediately justified beliefs. But this conception is not obligatory. Foundationalists themselves hold that immediately justified beliefs possess positive justificatory status in virtue of the obtaining of certain facts; justification isn’t a quality which “flows” to the immediately justified belief from something else which has it, but rather a status which accrues when certain conditions are met. That conception can easily be extended to include the case of inferentially justified beliefs; rather than thinking of justification as “flowing” through the inferential link, we can instead say that when a belief is arrived at through inference from another, one of the conditions that must be met in order for the former belief to be justified is that the latter belief must also be justified. This approach can be generalized to suit the anti-foundationalist’s purposes as well. The anti-foundationalist just has to claim that one of the conditions that must be met in every case is that one have some other appropriate justified beliefs as well. There will then be no sense to the question, “How does the *first* belief get justified?” Positive justificatory status will accrue only to appropriate clusters of beliefs.

^{xviii} Cf., Wittgenstein, *Philosophical Investigations*, 478.

^{xix} It may seem that a dialectical regress threatens. Couldn’t someone persistently demand that one first justify the belief that R supports P, then the belief that what one has said supports the claim that R supports P, and so on? And wouldn’t one ultimately be forced into an infinite regress or a circular defense?

How the Inferential Internalist should respond at this point will depend upon broader, logically independent epistemological commitments. For instance, if one follows William Alston (1989, p. 82, p. 83 fn. 3) in sharply distinguishing being justified from the ability to succeed in the activity of justifying, then one could hold that the possibility of this dialectical regress is simply epistemologically irrelevant (cf. Pryor 2000). However, an Inferential Internalist can plausibly argue that no problematic dialectical regress arises *even if* one holds that being justified requires the ability to justify one’s beliefs. To cut off the impending dialectical regress, it will be sufficient if it is internal to the structure of the activity of justifying that not just *any* demand for a justification is in order on any given occasion. For if that is so, then the requirement that one be able to provide good reasons for one’s belief does not dictate that one must be able to satisfy the persistent interlocutor’s demands all at one sitting; there will be at least some beliefs for which one has the requisite ability though it is illegitimate for the interlocutor to demand its exercise. This basic point can be elaborated in many different ways, and how it is best developed is not my topic here (see, e.g., Leite (2005), Williams (1999), Brandom (1994: chapt. 4 sect. 1 – 4)).

^{xx} For an important discussion of this issue, see Klein (1999).

^{xxi} A possible parallel here is provided by cases of so-called “dispositional beliefs”, such as your current belief that the earth is more than 572,142 miles from the sun – a belief which it is plausible that you in some sense had before reading this sentence.

^{xxiii} If this appears mysterious, consider that it is also part of commonsense psychological explanation that belief involves appropriate dispositions, including dispositions of both thought and action. Believing that P supports Q involves defeasible dispositions to make certain

inferential moves in certain circumstances, for instance. In this way, it is part of commonsense psychological explanation that belief is not merely an occurrent representational state.

BIBLIOGRAPHY

- Alston, William (1989). *Epistemic Justification: essays in the theory of knowledge*, Ithaca, NY: Cornell University Press.
- Audi, Robert (1993). "Belief, Reason, and Inference," in *The Structure of Justification*, New York: Cambridge University Press, 233 – 273.
- Boghossian, Paul (2001). "How Are Objective Epistemic Reasons Possible?" *Philosophical Studies* 106, 1 – 40)
- (2003). "Blind Reasoning," *Proceedings of the Aristotelian Society, supplementary volume LXXVII*, 225 – 248.
- BonJour, Laurence (1998). *In Defense of Pure Reason*, New York: Cambridge University Press.
- Carroll, Lewis (1895). "What the Tortoise Said to Achilles," *Mind*, 4, 278 – 80.
- Fumerton, Richard (2004). "Epistemic Probability," *Philosophical Issues* 14, Epistemology, 149 – 164.
- (1995). *Metaepistemology and Skepticism*, Totowa, NJ: Rowman & Littlefield.
- Ginet, Carl (1985). "Contra Reliabilism," *Monist* 68, 175-187.
- Huemer, Michael (2002). "Fumerton's Principle of Inferential Justification," *Journal of Philosophical Research*, vol. 27, 329 – 340.
- Leite, Adam (2005). "A Localist Solution to the Regress of Epistemic Justification," *Australasian Journal of Philosophy* 83, 395 – 421.
- (2004). "On Justifying and Being Justified," *Philosophical Issues* 14, Epistemology, 219 – 253

-
- Klein, Peter (1999). "Human Knowledge and the Infinite Regress of Reasons," *Philosophical Perspectives* 13, 295 – 325.
- McDowell, John (1998a). "Knowledge and the Internal," in *Meaning, Knowledge, and Reality*, Cambridge, MA: Harvard University Press, 395 – 413.
- (1998b). Knowledge by Hearsay, in *ibid.*, 414 – 443.
- Pryor, James (2000). "The Skeptic and the Dogmatist," *Nous* 34, 517-549.
- Stroud, Barry (2000). "Belief, Reason and Inference," in *Meaning, Understanding, and Practice*, New York: Oxford University Press, 33 – 51.
- (1977). *Hume*, New York: Routledge.
- Thomson, Judith Jarvis (1965). "Reasons and Reasoning," in *Philosophy in America*, Max Black ed., Ithaca, NY: Cornell University Press, 282 – 303.
- Van Cleve, James (1984). "Reliability, Justification, and the Problem of Induction," *Midwest Studies in Philosophy IX*, 555 – 567.
- Williams, Michael (1999). "Scepticism," in *The Blackwell Guide to Epistemology*, ed. J. Greco and E. Sosa, Malden, MA: Blackwell Publishers
- Wittgenstein, Ludwig (1953). *Philosophical Investigations*, Oxford: Blackwell Publishers.
- Wright, Crispin (2001). "On Basic Logical Knowledge," *Philosophical Studies* 106, 41 – 85.