

Regress-Stopping for Neopragmatists  
Jonathan M. Weinberg  
Indiana University  
jmweinbe@indiana.edu

## I. A Brief Metaepistemological Preliminary

I want to begin by making a distinction between two families of epistemological projects. The first are concerned with the fundamental questions about how, or even whether, the mind can so much as gain a rational hold of the world around it. This family includes such debates as those concerning classical forms of skepticism; the grounding of chains of reasons and the existence of ‘the Given’; the possibility of *a priori* justification; and the factivity, and factorability, of knowledge. Such projects typically do not draw in any substantive way on our larger theory of the world, but can seem rather to be of a sort that, if they can be carried out at all, can utilize only the most absolutely general or purely structural considerations of the human condition. Let us call this family of projects *the epistemology of morning*, in its efforts to discern what happens at that very moment as light dawns upon the human mind. The epistemology of morning strives to take as little as possible for granted, and often its internal methodological debates concern just how dimly glinting a set of resources can be deployed while still yielding valuable insights.

We can contrast those projects with a different family of projects: to learn what particular norms should guide our inquiries in the world as we find ourselves in it; to

---

<sup>1</sup> (unless favorably)

glean what lessons can be cultivated from the specific successes and failures of the histories of our many cognitive endeavors; or to wrestle with the problems and puzzles that arise in the latest chapters of those histories. The pursuit of such goals requires one to be as deeply informed by our most accurate theories of human cognition, including its biological and social mechanisms, and the particular causal workings of its interplay with the environment. These projects take it as background that skepticism is refuted (or perhaps was a nonstarter in the first place), and that somehow-or-other our senses and intellects achieve significant purchase on the world. Let us call this family of projects *the epistemology of midday* – the dimness of the dawn is taken, for its purposes, to be past, and one strives to look at the epistemic world with as much illumination as can be mustered. Without putting too much weight on what is already a somewhat strained metaphor, I hope that the image of the brightness of early afternoon also suggests much of day's work yet to be done, and nowhere near yet that metaphorical period of dialectical resolution we often refer to as "at the end of the day".

One could argue that one or the other of these families is more important than the other, and I suppose one could read Quine's "Epistemology Naturalized" as urging us to abandon the former for the latter, and such authors as Stroud or Fumerton as contending that there's not much in the latter that *really* counts as philosophy. Although I will be pursuing in the rest of this paper a piece of the epistemology of midday, nonetheless I would stress that I'm not at all trying to contend that the more traditional projects should be displaced, but merely complemented. Indeed, I am offering this distinction in hopes

that it might suggest a way in which naturalistic and not-so-naturalistic approaches to epistemology can dwell together in a less combative relationship.

Now, even once one has decided to pursue an epistemology of midday, one still faces the question of how one is to proceed in doing so. One must figure out how to deploy the sorts of scientific, historical, and anthropological information that one is now planning on making use of. I suspect that we cannot just say vaguely that we are now doing *very* wide reflective equilibrium and leave our methodological considerations at that, since to do so would be to ignore all the meta-methodological content that would be a part of any such reflections. We will have only pushed the question back: upon reflection, however wide, how ought we to move forward?

To the extent that one can identify natural epistemic kinds, then perhaps one could follow Kornblith in looking to transform some epistemological questions directly into ones we could answer by consulting the relevant scientists, such as cognitive ethologists. But it is not at all clear that there are very many such natural kinds, maybe not even much more than knowledge, and it's not at all clear that knowledge itself counts. Philosophers could also perhaps hope to use the techniques of scientists themselves to investigate questions of epistemological interest, in the manner of some forms of experimental philosophy. But it is not at all clear that such methods can tell us much other than descriptive psychological facts about ourselves and other persons; such facts may be part of what we want to appeal to, but the question immediately at hand was *how* to go about using them. Being told that we'll get more of them, and gather them in the manner of the experimental sciences, does not address that question.

I hope that there will turn out to be a number of different, complementary methods that we may ultimately be able to deploy in pursuing these sorts of projects<sup>2</sup>, but I will just offer one such method here, which I have elsewhere<sup>3</sup> called the method of *reconstructive neopragmatism*. The game is to imagine that we were going to re-engineer some set of our epistemic norms, given what we now take ourselves to know about the kinds of creatures we are and the kind of world we operate in. We ask ourselves first: for the set of norms in question, what epistemic purposes are they meant to serve? (To focus on specifically *epistemic* purposes is part of what keeps this program from collapsing into a more amorphous brand of pragmatism, such as that which Susan Haack has attacked Steve Stich for advocating.) With such purposes in mind, we can then ask: what norms would best serve those purposes? Thus do we reconstruct the epistemic norms that it makes the most sense for us to be guided by.

## II. Getting the Regress Started

Here is a weak version of the kind of principle that many philosophers have found intuitively attractive, but not uncontroversially so:

(D) To be justified in believing *p*, an agent must be able to cite an appropriate reason in defense of *p*, and that reason must itself be justified.

(This formulation intentionally leaves a lot of key epistemological questions unaddressed, such as what more particularly constitutes an appropriate reason for *p*, and in what sort of

---

<sup>2</sup> Bishop and Trout's ameliorative approach might be one, which it would take me beyond the scope of this essay to consider.

<sup>3</sup> "What's Epistemology For?" (2006), in *Epistemology Futures*, Hetherington (ed.), OUP; also "Moderate Epistemic Relativism and Our Epistemic Goals" (2007), *Episteme*, 4, 66-92. This framework is heavily indebted to Edward Craig (1990) *Philosophy and The State of Nature*.

time-frame such a defense must be presented. Also, it may be that principles like (D) only come into play for claims, not beliefs themselves. I hope to be able to elide the discussion of such particulars in what follows.)

Under the neopragmatist rubric, then, we ask: does a principle like (D) well serve the epistemic goals that norms of justification are meant to promote? If so, can we determine how generally (D) should be applied? I will conjecture here, with only a brief defense, that the norms of justification are meant to promote jointly the goals of *diachronic reliability* and *dialectical robustness*. We wish to get a hold of truths, and to keep a hold of them once we've gotten them. Moreover, we wish to have fruitful discussions, collaborations, exchanges, and even more importantly, disagreements. I will not concern myself here with the question of whether the former of these two "DR desiderata" is in some sense prior to the other; creatures like us, in a world like ours, can best pursue the truth in promiscuously interactive ways, and it will be highly unlikely outside of pathological cases that a set of norms can do as good a job as it can of promoting either DR desideratum without also promoting the other. (It really does go both ways – a set of norms that help people get a hold of the truth will also tend to make it easier for them to have fruitful discussions and arguments, as it will be easier for participants to find premises that they can agree on.)

I should note that these two do not exhaust our epistemic goals on the whole. For example, we clearly place a high value on discovering new truths, and one could argue about whether securing truths of significant practical importance is itself an epistemic goal. But we are concerned here with justification, which is primarily a norm of

permission. It is not the job of justificatory norms to direct the targets of inquiry, but to ensure that the inquiry taking place is meeting basic standards of epistemic hygiene. So I take it to be appropriate to restrict our attention here to the two desiderata already identified.

Now, with these DR desiderata in mind, we can see why it makes sense *generally* to require that agents be able to muster good reasons on behalf of their beliefs, reasons that they not only can cite publicly but also which they can thereupon defend from further challenge as needed. On the front end, beliefs that are formed already able to pass such a requirement will be more likely to be true than ones that cannot. But just as importantly, after a belief has been formed, satisfying (D) will allow us to respond actively should any concerns be raised about it – either because of dissent, or the uncovering of new evidence, or perhaps simply because we now wish to put it to more demanding epistemic or practical work than we had asked of it until now. We can inspect its credentials, and if we find them wanting, we can apportion further epistemic resources either to shoring them up, or to rejecting the claim. Having citable, defensible reasons for one's beliefs enables a kind of epistemic bookkeeping for those beliefs, which is tremendously valuable when one finds oneself needing to pinpoint sources of error and dissent, and either eliminate them or at least quarantine them.

Another way of thinking about we might benefit from adhering to a principle like (D) is to consider what sorts of creatures would *not* particularly benefit from such a norm. Creatures with very limited inferential or reflective powers would have little to gain from following such a norm, because they would not have any resources to make use of the extra information that they were retaining. Whatever it is that they are getting in through

their fundamental sources of evidence, that's what they've got, and there may be little that can be done to make them better or worse. If they at least have some capacities for self-monitoring, they may be able to distinguish good belief-forming circumstances from bad, and thus guide their cognition somewhat, but further cognitive bookkeeping will not help them. (We have some reason to think that at least some other primates have this sort of cognition.) At the opposite end of the spectrum, creatures with spectacularly strong powers of coherence-based inference and a very broad sampling of the facts of the world might not be particularly aided by such bookkeeping, because they can simply survey their accumulated beliefs, which *ex hypothesi* are drawn from a sufficiently comprehensive stock of beliefs about the world, and can discern whatever outliers they may have among them.

Yet we lie in between these extremes. Our inferential ambitions vastly outstrip our sources of evidence, yet we have nontrivial, but nonetheless fairly limited, powers of holistic inference. And the sources of evidence that we have present us with only a very slender slice of the universe over the course of a lifetime or even many lifetimes. We are thus at risk of our belief-forming practices being dangerously *path-dependent*, subject to the luck of the draw concerning what initial inputs we manage to get a hold of to start our inquiries moving. What seems likely to us at one point in time will influence what we choose to investigate in the near future, and what look like successes early on will have a deep impact on what shape the rest of our inquiries take. This is a good strategy to maximizing the range of truths we might acquire using our limited resources, but it is not without risk; as Timothy Williamson has warned, in the context of urging philosophy to “do better”, “A few resultant errors easily multiply to send inquiry in completely the

wrong direction.” This is the human epistemic condition, and it makes it all the more important, therefore, that we be able to recover from such errors where they may occur, lest an initial wrong turn lead us into a permanent blind alley. Adhering to a principle like (D) or its ilk are among our primary tools for limiting path-dependence, because it requires us to carry with us whatever aspects of our cognitive past we are still looking to put to work. If we find ourselves in a dead-end, as it were, then we will know how to rewind the string and navigate back, and try a different path from there.

Nonetheless, if we try to make explicit defendability into an *absolute* requirement, we will end up with a norm that no one can actually manage to satisfy. This is not an in-principle argument, and I do not offer here any argument that, say, a infinite chain of reasons *would* automatically fail to justify, were an agent *impossibile dictu* to deploy such a chain. Our reasons are finite. Moreover (*pace*, e.g., Leite 2007), our reasons do run aground, and there are claims we would make for which most of us have no idea how to respond to a request for further reasons for them. Many of our beliefs are such that almost all agents end up “dumbfounded”, should they be asked to defend them (see, e.g., Haidt (2001), Hauser (2006)). Whether or not reasons *must* come to an end somewhere as a matter of conceptual or metaphysical necessity, as a matter of fact for creatures like us, they always *do* come to an end, and not simply because we reach a point at which no further reasons happen to be demanded. So we will need to discern where such a requirement can be relaxed, consistent with promoting the DR desiderata.

(Reconstructive neopragmatism, as an epistemology of midday, starts off by considering skepticism not as any sort of threat; nonetheless, skeptical concerns are not

irrelevant to our discussion, as they serve as a kind of quality-control – an unfulfillable norm will not help guide our actions, and should thus be rejected.)

There are a number of places in which it looks like our capacity for reason-giving screeches to a halt. Some of these can be located in terms of the content of the claims; basic moral claims, for example, or Wittgensteinian “hinge propositions”. Another locus of dumbfounding is at our appeals to our basic sources. As many have argued, we can’t defend them, certainly not without circularity. But we can’t have an exemption for *any and all* putative sources, lest we allow in far, far too much, and lose the advantages in terms of both reliability (by admitting in grossly unreliable sources) and dialectical robustness (by allowing different people or groups to effortlessly admit different, perhaps wildly divergent sources). We need a set of exemptions from the demands of (D), but they need to be tailored to the DR desiderata. And if we *can* find ways of restricting (D) consistent with promoting the DR desiderata, then it will be hard to see why we *should* adopt the absolute version over a pragmatically-preferable restricted version. And I think that such restrictions are available, under very particular sorts of circumstances.

### III. Regress-stopping for Neopragmatists

One such circumstance is when the appeal to the source in question is one that is *practically infallible*, and is recognized to be so. By “practically infallible” I mean that it is extraordinarily unlikely, in any sort of circumstances that the agent is at all likely to find herself, that the source could be mistaken. In using such a source, the agent is putting herself at very, very little risk of being wrong in the first place; moreover, by making an appeal that is recognized to be practically infallible, there will be no need to

settle cases of dissent, since we can require as part of these norms that any contending parties acknowledge that appeals to such a source are beyond challenge, except under extraordinary circumstances. (Should we have reason to think that such extraordinary circumstances actually obtain, then we will want a meta-exemption: in the presence of specific evidence that the appeal to the source is, in this instance, *not* practically infallible, (D) comes back online.)

This *practical infallibility exemption* yields at least two important appeals to sources as legitimately exempted from (D). First, it will apply to many simply cases of the *a priori*, such as basic arithmetic, the *modus ponens*, and the *cogito* inference (whether or not that really counts as *a priori*). Second, it will apply to the most central cases of introspection, such as a report on a current severe headache. Are our appeals to such sources perhaps more than practically infallible, but indeed infallible *simpliciter*? I don't know, but I'm inclined to doubt it. I suspect that no philosopher's liver is so strong that no amount of good scotch can possibly produce in them a state in which they could misdiagnose their own sensations and/or fumble their way into affirming the consequent. But one-hundred-percent, metaphysically- or conceptually-guaranteed infallibility will of course count *a fortiori* as practically infallible, so if such sources *are* completely infallible, that will certainly not pose a problem for this loophole. But I see no advantage to be gained making full infallibility the requirement for the loophole.

But the version of (D) we get even softened with this loophole is still almost certainly too restrictive. Surely we need something that actually gets us something like perception;

but perception isn't infallible, not even practically. We *might* could try to do so in the manner of truly old school foundationalists, and try to springboard into perceptual claims about the external world based on introspective claims about our inner happenings. But then we would have to require agents to be able to defend some *further* reason that could license their inference from a claim about the contents of their minds to one about the world outside of them. And I don't see how we can expect typical agents would have anything like the cognitive resources to defend such a reason; indeed, it's not even obvious to me that typical agents would have the resources to devise and articulate such a reason in the first place. This would be a locus of dumbfounding, and is thus precisely the sort of work we need to be performed by the norms themselves, and not require it of individuals.

An obvious candidate to consider here is simple reliability. This candidate is obvious both from its prominence in the literature, and also as a further variant of the conditions I have already considered: since I have already suggested that we don't need full-blown infallibility in our sources to merit an exemption, but merely practical infallibility, so why not weaken the condition further, to something less than even practical infallibility but still surely something well about a 0.5 truth ratio? Why, from a neopragmatist perspective, should we bother requiring people to defend their appeals to reliable sources? Or, if we want to include a nod towards dialectical robustness, why should we bother requiring people to defend their appeals to sources that are reliable and recognizes as such

But to do so would not, in fact, properly serve the DR desiderata. I suggested earlier that much of the value in having a principle like (D) is not just its promotion of reliability

on the front end of belief formation, but also on the back end of belief revision – the help it would give us in tracking down errors, localizing disagreements, reducing path-dependence, and so on. But mere reliability does not help much with these other components to the value of (D). Merely having used a reliable source to form a belief that  $p$  will not be of much help should one later have concerns about  $p$ . As noted above, much of the value of having reasons is to be able to backtrack and correct for errors that may arise in the course of inquiry, and to give us avenues for discussion should dissent arise. If one possessed a theory of how the source works, and in particular what errors it can fall prey to and under what circumstances; and if one furthermore had good records of the circumstances under which one received its deliverances, then one could perhaps go back and try to discern where any error might have crept in. But, obviously, if one has all *that*, then (D) will be more than adequately satisfied without the need of any sort of special exemption. The mere reliability of a source is not sufficient for us to exempt appeals to it from (D) – we still need some reassurance that our need to locate and correct or quarantine errors in the source will be addressed. If we only frame such conditions in terms of reliability, then practical infallibility seems to be the weakest such condition. (I believe it is a corollary to these considerations that even when one does have reasons to deploy concerning an appeal to a source, a premise of mere reliability will not necessarily be sufficient to secure that source's trustworthiness. But that is a point that lies beyond our current investigations.

What we need to do, it seems, is look beyond different degrees of reliability, and consider other ways that an appeal to a source of evidence can allow us to monitor and accommodate errors successfully, other than the possession of appropriately citable

defenses. If the advantage of practical infallibility was that little risk of error is incurred in the first place, we might consider whether there might be circumstances in which such a risk is present – but which is nonetheless a risk that can be managed in a fairly immediate way, in which we can easily investigate possible errors, and in which potential dissent could instantly be settled. Practical infallibility is like riding a train with very little chance of deviating from the desired path; this other sort of circumstance would be more like driving with a great GPS navigational system, able to respond flexibly in real time to any difficulties that may arise. In either sort of circumstance, (D) would have little work to do for us, and could safely be voided.

But this is the situation we are in with paradigm instances of perception, such as when we report on the layout of our immediate surroundings. We have excellent resources for the detection and correction of error with perception in such cases – a virtue I have elsewhere<sup>4</sup> termed “hopefulness”. Our percepts themselves carry substantial information as to whether they have been formed in more or less error-prone circumstances; when the lights are dim, we experience the world *as* poorly-illuminated, and not, or not just, as gray. Should a visual report seem amiss, it is very easy for us to look closer (or step back), to turn the lights up (or down), to look from different angles, to check with each others’ takes on the scene, even to apply various aides such as rulers or color-swatches. The hopeful nature of perception is thus partly a matter of our psychological hardware, and partly our folk practices that have been honed over millennia. Moreover, both the hardware and the practices are very broadly shared, and recognized as such. So when

---

<sup>4</sup> “How to Challenge Intuitions Empirically Without Risking Skepticism”, *Midwest Studies in Philosophy*, 2007.

someone appeals to perception, their interlocutors can generally acknowledge when such an appeal has been made correctly. Often they can even just check for themselves.

So we can allow regresses of justification to terminate at appeals to the deliverances of perception, without doing harm to the DR desiderata. We risk little harm to the goal of diachronic reliability, because perception is basically reliable, and when it makes errors we are generally very good at noting that we have done so, and in taking steps to discern the matter more clearly. And we risk little harm to the goal of dialectical robustness, because those we are in dialogue with can easily check, immediately and on the spot, whether we have made any such appeals in accord with our best perceptual practices. No further bookkeeping is needed, because the checking that such bookkeeping is meant to enable is something that we can already do without it. We thus have, in addition to the practical infallibility exemption, an *immediate hopefulness* exemption.

We might wonder whether it is to be the proposition so perceived, or the psychological state of the perceiving itself, that is the regress-stopper. The former does not, by itself, possess this range of hopeful resources; the latter seems, by itself, to lock us back behind a veil of appearances. Better, I think, to take the regress as stopped at the proposition so perceived, coupled to the fact that it is via perception that one is staking that claim. Our perceptual psychology and practices do not themselves serve reasons that need be cited – at least, not normally – but they play an ineliminable part of what entitles us to appeal to the propositions that do so serve. (One might compare them to the role of the right-hand column in a logical derivation; they make evident that we are complying properly with the rules of the game, but are not themselves premises or lemmas in the proof.) One interesting upshot of these considerations, though, is that it can make sense

of the tension that one can feel hard to resolve between allowing perceptions of the external world or one's own internal seemings to be the best candidates for an empirical regress stopper – namely, both are excellent candidates, but for very different reasons. There is no need to choose between them as *the* stopping point for regresses of reasons, because both are worthy regress-stoppers, as measured by the DR desiderata.

Two quick clarifications are in order. First, we need a meta-exemption here as well, for cases in which we have reason to think that our practices cannot operate in their usual way, for example, when one thinks that one has ingested hallucinogens. Second, although this exemption is of clearest relevance for our perceptual practices, it is at least possible that it may apply to other putative sources of evidence as well. At least some uses of memory, for the very-recent-and-not-yet-changed past, may be able to pass through this particular loophole.

#### IV. Questions & Upshots

We have thus seen that, although principles of epistemic responsibility like (D) have great promise for our epistemic goals of diachronic reliability and dialectical robustness, nonetheless we ought to allow some key exemptions – at a minimum, an exemption for the practically infallible (and recognized as such), and another for the immediately hopeful (and recognized as such). Let me close with two questions, and two upshots.

One question that this conclusion prompts is: are these the only exemptions? Are these sufficient for our purposes? Clearly we need to consider how, and in what way, appeals to testimony, to memory of the more distant past, or to moral intuition can be

admitted. Neither of these two exemptions seem a particularly good fit for any of those sorts of appeals, but maybe the grounds for making an appeal to them can be constructed from information about the world that is let in via those exemptions. I somewhat doubt it for testimony at least, on general anti-reductionist grounds. Nonetheless, there is important room for exploration here about the relation between “monitoring” conditions and the epistemology of testimony (e.g., Fricker 1994; Goldberg & Henderson 2006). Further investigation is warranted.

Also, in order to support the desideratum of dialectical robustness, we required that the practically infallible or immediately hopeful also be widely recognized as such; we will thus need ultimately to address what sort of “recognition” is to be required. Do agents need to have justified beliefs that the conditions for an exemption is met? Or is it enough that the agents in a given community all simply believe, perhaps without further justification, a given set of appeals to a source to have the requisite property? Indeed, is it even required that they believe that the appeals to have that property, or is it enough that they do a decent enough job of distinguishing apt from inapt appeals, without any particular explicit cognizing of the conditions themselves? I suspect that this very weak construal of “recognition” will be all that the DR desiderata require, but it is not in any way an obvious result.

One upshot of our considerations here is that it is the nature of our *practices* with our various faculties, even more than the faculties themselves, that license the regress-stopping. If someone were to make undefended appeals to introspection outside of the practically-infallible subset of its range, that person will be asking for epistemic trouble. If someone had a different set of practices with perception that lacked hopefulness –

perhaps relying on special, personal capacities like putative “aura seeing” – then the DR desiderata requires us to require them to defend such appeals. When it comes to stopping the regress of reasons, it’s not so much a matter of what you’ve got, but what you do with it.

Finally, I would like to offer a closing thought on the potentials for naturalism in epistemology, particularly in the context of pursuing what I called projects in the epistemology of midday. We are used to the idea of naturalism as a metaphysical constraint on the sorts of entities one can admit into one’s ontology; or as a methodological constraint as to how certain sorts of propositions can be known; or perhaps also as offering the promise of contributing additional information as premises in philosophical discussion. None of these has played much of a role in my arguments here. Rather, I think we have perhaps not yet tried fully to capitalize on naturalism’s promise as a source of *theoretical* resources for philosophizing as well. From purely armchair considerations, the kind of exemptions I have been arguing for here might seem strange and ad hoc. My hope is that the framework I am advertising for here is one in which we can see that they are natural and motivated. When we consider epistemic norms as abstract truths delivered to us by intuition, then perhaps a certain unity and simplicity will seem appropriate to them. But when we consider them as tools that are ours to configure as best we can to help promote our epistemic goals, then a certain amount of trickiness and complexity, as our norms warp a bit to match the contours of the crooked timber of humanity. We know to expect such complexity in, say, university degree requirements and state traffic codes – we should not necessarily expect any less in epistemology.