(unproofed final draft for *Epistemology Futures*, 2006, Stephen Hetherington ed., OUP)

What's epistemology for? The case for neopragmatism in normative metaepistemology [*]

Jonathan M. Weinberg

How ought we to go about forming and revising our beliefs, arguing and debating our reasons, and investigating our world? If those questions constitute normative epistemology, then I am interested here in normative *meta*epistemology: the investigation into how we ought to go about forming and revising our beliefs about how we ought to go about forming and revising our beliefs -- how we ought to argue about how we ought to argue. Such investigations have become urgent of late, for the methodology of epistemology has reached something of a crisis. For analytic epistemology of the last half-century has relied overwhelmingly on intuitions,[1] and a growing set of arguments and data has begun to call this reliance on intuition seriously into question (e.g., Weinberg, Nichols, and Stich 2001; Nichols, Stich, and Weinberg 2003; Cummins 1998). Although that method has not been entirely without defenders (BonJour 1998; Bealer 1996; Jackson 1998; Sosa forthcoming; Weatherson 2003), these defenses have not generally risen to the specific challenges leveled by the anti-intuitionist critics. In particular, the critics have attacked specific ways of deploying intuitions, and the defenders have overwhelmingly responded with in-principle defenses of the cogency of appealing to intuition. An analogy here would be someone's responding to arguments alleging systematic

misuse of a particular scientific instrument, with accounts of how such an instrument could in principle be a reliable source of data.

But perhaps the best -- or most psychologically persuasive, anyway -- case for intuitions is a sort of 'what else?' argument. In the absence of a rival method to take its place, it is surely more rational to keep the problematic epistemology we know, than to abandon epistemology completely. My intention is not only to strengthen the case against traditional intuition-centered methods, but to articulate a better metaepistemology, too. First, I present a framework for debating these questions (a meta-metaepistemology?), based on articulating our desiderata for our normative epistemology. I then apply that framework to compare three basic methodological ideologies: intuitionism, naturalism, and pragmatism. I hope to suggest that pragmatism (or my version of it, anyhow) is an under-explored option on the table, and to that end, I further demonstrate applications of its method to some extant philosophical problems.

Metaepistemological desiderata

The mode of argumentation here is a normative parallel to inference to the best explanation. In inference to the best explanation, we consider some phenomena, and then evaluate competitor theories by how well they explain them. The winner need not perfectly explain all the phenomena, and indeed it may fail miserably on some, if only those failures are compensated with greater overall explanatory success. For our purposes here, we do not have phenomena so much as desiderata -- characteristics that we'd ideally like our methodology to

have. In comparing different methods, we see how well each would promote each desideratum. The winner will be the method that performs best overall, though it may not perform perfectly on all, and indeed may be outscored by competitors on some.

I shall put forward a list of seven desiderata. There is nothing sacrosanct about this list, each item is hopefully attractive as an empirical generalization as to what we really would want from a philosophical method here The list is also not necessarily a closed list, and I welcome the articulation of further desiderata that ought be addressed. But these should be plenty to get my argument rolling:

1. truth-conduciveness
2. normativity
3. dialectical robustness
4. progressivism without radicalism
5. interdisciplinary comportment
6. minimal naturalism
7. plausible relativism/universalism

I will briefly discuss each in turn.

The two most central desiderata are *truth-conduciveness* and *normativity*. We require that the method tend to produce true results, and the greater the ratio of true deliverances to false, the better. I won't try to use this desideratum as a consideration between intuitionism,

naturalism, and pragmatism, though I hope it justifies so abbreviated a list of candidates; i.e., astrological metaepistemology is right out. Moreover, we are looking for a method for normative epistemology, and as such, it had best give us results that are themselves normative in nature. Note that these desiderata receive a greater weighting than the others -- unlike the other five, failures here may not be compensatable with success elsewhere.

Perhaps in service of truth-conduciveness is the desideratum of *dialectical robustness*: we want our method to be one that supports, encourages, and enables successful conversation and debate between epistemologists. It should foment the discovery of both points of agreement and disagreement, and in cases of the latter, it should help us to resolve such disagreements fruitfully. (See section II.B below.) *Progressivism without radicalism* suggests that we want a method that can take us beyond mere common sense, and give us new norms as our overall epistemic and cognitive circumstances change. We want our normative epistemology to change its deliverances as our circumstances change -- the norms appropriate today may not be those that were appropriate in Plato's day or even Descartes's, and we want our methods to be able to register those changes. However, our proposed norms cannot fall so extremely remote from common sense that we cannot find them intelligible as epistemic norms for creatures like us to follow. Hence, we want progressivism without radicalism. Moreover, a method is better to the extent that it can learn from, or at least remain consistent with, the deliverances of other fields. An epistemology isolated from such sciences as

psychology and such humanities as history is, *ceteris paribus*, less desirable than an epistemology with rich *interdisciplinary comportment*.

The last two considerations are more metaphysical in nature. First, our methods should be consistent with the kind of naturalism that I take to be part of the contemporary philosophical *Zeitgeist*. I'm only insisting on a *minimal naturalism* here, however, which I take to be the requirement that all causally efficacious entities be materialistically respectable. It is less a strident reductivist sort of physicalism and more a loose anti-supernaturalism. There may be numbers, sets, or fictional objects in our ontology, but if there are any such, then we cannot have them causally interacting with the chairs and electrons and organisms of the world. Thus, e.g., various sorts of Platonic epistemologies are ruled out. Finally, we ask that our method's picture of the epistemic norms take a sensible stance on epistemic relativism. If it is a universalist method, we should be able to see why it licenses that universalism; if it allows a certain degree of epistemic relativity, then we should be able to see why that much relativism makes sense.

Let us see how the two main current metaepistemologies score.

I.A Intuition-driven romanticism

The main paradigm methodology in epistemology of the last few decades has been recently termed *intuition-driven romanticism* (Weinberg, Nichols, and Stich 2001). The proper epistemological norms are somehow already inside of us, and the job of the epistemologist is to get them out and set them out clearly; and the best way to do so is to pump our

spontaneous judgments about applying or withholding terms of epistemic praise or blame to various hypothetical cases. ‘Gettierological’ projects are paradigm instances of IDR methodology.

In terms of our desiderata, clearly the results of an IDR analysis are normative: they purport to tell us the structure of *the* concepts or terms that govern our epistemic lives. And the most basic Reidian self-trust in our own capacity for judgment requires that we take IDR to be at least moderately truth-tracking.

On other desiderata, however, IDR fares less well. For starters, IDR has not proven itself able to interact fruitfully with other disciplines. Perhaps it meshes well with some elements of logic and mathematics, though not all. (For example, withholding attributions of knowledge in standard lottery cases reveal a decided resistance on the part of our intuitions to be neatly mathematicized -- there is no probability p less than 1 that is sufficiently large such that a true belief with degree p of both objective and subjective probability automatically thereby counts as knowledge.) Despite the best efforts of the likes of Alvin Goldman, IDR has not learned anything from or taught anything to cognitive psychology.

Moreover, IDR runs the risk of being insufficiently progressive. One might worry that our folk epistemology reflects the last few centuries of development in the norms of reasoning and believing, but won't have had a chance to incorporate any lessons learned more recently. For example, it seems that the epistemic intuitions of educated Westerners tend to be very sensitive to even the mere possibility of a belief's turning out false. But, given that

modern science's results are so thoroughly lacking in claims to infallibility, , those more absolutist intuitions may simply not be up to the job of guiding us in today's epistemic world.

IDR doesn't obviously have to reject minimal naturalism, though at least some practitioners have felt compelled to do so. In BonJour's (1998) defense of rationalism, for example, he is clearly tempted towards a version of nonnaturalism in which our minds are somehow in direct contact with such abstract entities as triangularity itself. One can see why there might be a natural supernatural tendency here: it is natural to ask of IDR what explains the truth-conducivity of our intuitive judgments, but it may be hard to give a naturalistically acceptable answer without the consequence that our intuitions can tell us only about our own minds, and not about the norms themselves (Goldman and Pust 1998). So some IDR-practitioners may attempt to opt out of naturalism.

So far, IDR seems to score reasonably well on two desiderata, and somewhat poorly on three others. On the two remaining ones, however, it performs disastrously. Dialectical robustness requires that the evidence we cite to each other be evidence that we can each recognize the force of. But intuitions are damnably subjective. If I have a putative intuition that p, and you have a putative intuition that not-p, there's very little for us to appeal to other than mutual accusations of being captives of our respective theories. (I develop this worry about intuitions and dialectical robustness more thoroughly below, in II.C.)

The problem of varying intuitions also challenges IDR with regard to relativism. IDR practitioners typically invoke 'our' intuitions about a case, yet it is unclear just who 'we' are.

For Shaun Nichols, Stephen Stich, and I argue in our (2001) that the intuitions about various cases in the epistemological literature -- including Gettier cases and cases central to debates about reliabilism -- may vary significantly with ethnicity and socio-economic status. IDR practitioners would have to either find a way of contending that one such group's intuitions are to be privileged over the others, or surrender to some form of epistemic or linguistic relativism. The former of these options remains untaken, and the latter reduces IDR's score on this desideratum. At a minimum, we need an account of why it makes sense for epistemic norms to vary along *these* dimensions, since one would not have expected beforehand that, e.g., two native-English-speaking undergraduates from New Jersey *ought* to follow different norms, just because one student's grandparents are from Germany and another's from China. (We do not now have a practice of grading papers with different standards according to ethnicity, for example.)

In sum, IDR scores two hits, three so-so's, and two bad misses. We shall now see whether its chief contemporary rival -- metaepistemic naturalism -- fares any better.

I.B Metaepistemic naturalism

There are many positions that one could have in mind by the phrase 'naturalized epistemology', but only one of them has been put forward explicitly as a rival methodology to IDR: Hilary Kornblith's theory of knowledge as a natural kind. As he argues in his (2002), we should treat knowledge as a natural kind akin to how we treat water. It has a hidden essence, and we find that essence by seeing what it is that good science makes of it. The best

science involving water picks out a substance whose chemical structure is H_20. Similarly, he argues, the best science involving knowledge -- cognitive ethology -- picks out beliefs that have been produced by reliable processes. Now, one might accept Kornblith's framework but read the ethological literature differently; or perhaps plump for a different field to defer to, such as psychology or informatics, and thereby potentially get a different analysis of knowledge. Nonetheless, since our concerns here are at the meta-level, we should distinguish the methodological proposal from his intended results of that method.

It is of course unsurprising that *metaepistemic naturalism* (MN) scores highly on minimal naturalism and interdisciplinary comportment. We should grant it a default high rating on truth-conducivity as well, unless we are skeptics about science. MN is probably strong on dialectical robustness, though some recent worries about the objectivity of ethological observation have been raised.[2]

We might raise a concern about the implications MN has for relativism, in that perhaps it is *too* universalistic. Namely, we might worry that it has mistakenly conflated our proper epistemic norms with those that might govern chimps or plovers, when our greater cognitive and linguistic capacities, and more sophisticated social organization, might merit a distinct set of norms.

Finally, MN fails miserably on the last two remaining desiderata. MN has the potential for progressivism, but at the risk of extreme radicalism. Juan Comesaña recently proposed to me in conversation, only half in jest, that this kind of naturalism seems to allow

that knowledge could turn out to be fried potatoes -- but one need not allow for a hypothesis that radical to see that MN might permit our epistemic norms to fall unrecognizably far from the tree of common sense. Steve Crowley (manuscript) has suggested that the notion of knowledge operative in cognitive ethology is one that does not even require truth, in that the notion includes representational states that are too purely action-guiding or too widely distributed across a group of conspecifics to be propositional.

Worst of all, however, is that MN falls down on one of the two most key desiderata: normativity. (This is of course an old complaint against naturalism in epistemology, going back at least to Kim 1988.) Once we've learned what knowledge 'really is' according to science, the question still presents itself as to whether knowledge is something worthy of our pursuit. Kornblith argues for such a value, but his argument is, notably, not itself a matter of scientific inference. Rather, he tries to suggest that knowledge, as revealed to us by ethology, is something that in fact we might find of instrumental value. But there is no guarantee that MN's deliverances will have such a normative dimension. Indeed, MN might not apply at all to more explicitly normative terms in our epistemic repertoire, such as 'justified' or 'rational'.

So MN does not appear to score any better than IDR, and in fact does somewhat worse: three hits, two so-so's, and two bad misses -- one of them in the central desideratum of normativity. The poor performance of both of these methods should lead us to pursue other options.

I.C The case for reconstructive neopragmatism

The metaepistemology I am lobbying for is a variety of pragmatism, but it should not be confused with such brute versions of pragmatism that simply flat-out define the epistemic good in some other terms, such as the agreement of our peers (Rorty 1989) or the attainment of whatever we find intrinsically valuable (Stich 1990). I take such views to have been successfully harpooned by Haack (1993), and can be faulted as having given up on the epistemic altogether. My neopragmatism takes a subtler, two-stage approach.

Analytic philosophers typically focus on the conditions for the correct application of a concept, organizing their investigations along the axis of the question, 'What does it take for something to count as an X?' But of course we do not use concepts merely to categorize the world: we deploy those categories to help us make further judgments and generalizations. It is one thing to know that to be neon, it is necessary and sufficient to be a sample of an element with 10 protons. But then the cash value of knowing that something is a sample of neon comes from knowing what having 10 protons further entails about the substance, e.g., that it will be a noble gas. If there were no further entailments, we simply would not care to use the concept. (Philosophy is more than lepidoptery of the intellect.) Our interest in philosophical concepts like PERSON or VOLUNTARY is not just to parse the world in such-and-such a way. Rather, we think that persons should be treated differently than non-persons (only they get rights, perhaps), and voluntary actions should be treated differently than involuntary behaviors (only they are morally evaluable, perhaps). We carve up the world for

certain descriptive or explanatory or evaluative purposes, and if we wish to understand the role of a certain concept in our lives, we might well ask the question, to what end do we deploy this concept?

Thus, my neopragmatism's method is to ask: *for what purposes* might we reasonably prefer to have beliefs formed in accord with a given epistemic term? E.g., why might we choose to evaluate beliefs in terms of the presence or lack of justification, or knowledge, or certainty? Such questions are asked in a normative tone. We are not seeking simply an explanation for our happening to have such preferences -- there are probably many such explanations, of a psychological or evolutionary or cultural sort. We want to know why we should on reflection endorse these preferences (and not, say, decide to try to give them up as some distracting cognitive habit, a tic of the mind). Our question is: why ought we include such dimensions in our epistemic appraisals? To ask such questions is an attempt (in James's terms) to 'pump free air' around a given concept, and try to get a sense of how it fits into our epistemic lives on the whole. Should we be able to harmonize such a teleological view of a concept with the traditional attributional view, we would thereby attain a deeper understanding.

This teleological maneuver is also not unprecedented in epistemology itself. Here's a lovely and compact statement of it, with a bit of analysis:

> If the epistemic concepts can earn an honest living they must form a natural intellectual kind. Even if some multi-part analysis accurately matched our

> judgments in difficult cases, it would still need asking why we are interested in just *that* set of conditions …. But how can it be important to organize our lives around one complex of conditions rather than another? We need a role for the epistemic concepts, and the role which seems most natural is that of ranking and selecting titles to respect. We have to pick up our beliefs about the world from our senses and from each other. So we need a vocabulary to settle whether our sources are ones which themselves properly indicate the truth. This is a natural need, and it gives us the natural intellectual kind in which to place our epistemic verdicts. (Blackburn 1984: 169-70)

Blackburn is clearly dissatisfied with IDR here, when he questions our interest in having a successful descriptive account of our attributions. He wants to defend a concern for truth in our epistemic attributions, and he appeals to our desire to be able to categorize our doxastic sources with respect to their reliability. The passage also brings out the sense in which we can ask the question, 'Why do we care about knowledge?' in a manner that is not merely a psychological inquiry. The question is not, 'What underlies our mental pro-attitude towards knowledge?' but, more importantly, 'What role does the concept of knowledge play in organizing our lives?' (The question generalizes to other terms of epistemic appraisal.)

But we want to provide a more thorough analysis than can be supported in the one-paragraph form just quoted. We require not just a maneuver, but a method. We need a tool for organizing an investigation into the purposes of our epistemic norms -- something to play the

role for a teleological analysis that analysis-by-cases plays for attributional analysis. Fortunately, Edward Craig has recently given us a full-scale attempt at just such a methodology. In his *Knowledge and the State of Nature*, he expresses a similar dissatisfaction with the attribution-prediction type of project. Even were we to produce a successful version of such an analysis, he argues (1990: 2),

> I should like [that analysis] to be seen as a prolegomenon to a further inquiry: why has a concept demarcated by those conditions enjoyed such widespread use? There seems so be no known language in which sentences using 'know' do not find a comfortable and colloquial equivalent. The implication is that it answers to some very general needs of human life and thought, and it would surely be interesting to know which and how….
>
> Instead of beginning with ordinary usage, we begin with an ordinary situation. We take some prima facie plausible hypothesis about what the concept of knowledge does for us, what its role in our life might be, and then ask what a concept having that role would be like, what conditions would govern its application.

Craig invokes the 'state of nature' framework in social contractarian political philosophy (hence his title). Where Locke *et al.* concerned themselves with our basic social needs of cooperation and security, Craig concerns himself with the basic epistemic need of

telling whose testimony can be trusted and whose discounted. If we are to understand how our political and epistemic institutions ought to be structured, we should consider for what human purposes these institutions might have been founded in the first place.

I am extremely sympathetic to this framework, and am greatly indebted to Craig in what will follow.[3] But I have one significant difference in my choice of methodology here. For the 'state of nature' metaphor is fundamentally history-oriented. For example, it seems somewhat hostage to arguments about the actual past origins of the relevant institutions and practices. Moreover, the 'state of nature' approach assumes that our relevant human needs today are basically the same as they were back at the time of the mythological founding of our political or epistemic institutions and practices. Our more basic and biological needs probably are more or less unchanged; however, the epochal shifts in our social structures and base of knowledge over the last few centuries have changed what we want from the political and the epistemic. It could be argued that in the political realm, citizens have begun to require that the state become something stronger than just a guarantor of personal safety and property rights, and transform into a proactive agent for positive change and social justice. And clearly our demands in the epistemic realm have changed, becoming more tolerant of probability and less obsessed with certainty, less suspicious of testimony, and more suited towards socially cooperative investigations. Given such historical -- and ongoing -- developments, it could be unwise to attend too closely to the past conditions imagined in any analysis-by-state-of-nature.

In order to retain the teleological viewpoint forwarded by Craig, while avoiding the potentially historical biases of a state-of-nature approach, let me invoke instead a framework of analysis-by-imagined-reconstruction.[4] My operative question will be: were we to consider a radical re-constitution of our epistemic norms, what would we include, what might we strengthen, and what might we abandon as outmoded?

An analogy may help. Many modern societies' political norms are to an important extent encoded in their constitutions. And by and large questions of political acceptability can be simply referred to that document, with perhaps some consideration of the 'original intent' of its composers. But, significantly, a society also retains the capacity to amend that constitution -- or, if necessary, to call a new constitutional convention. Such foundational changes are sometimes necessary, because changing conditions render the original document less well-suited to performing its function. A society can consider altering existing structures (e.g. expanding the franchise to include African-Americans or women or 18-year-olds); introducing new norms (prohibiting the sale of alcohol; permitting a national income tax); or eliminating existing norms (repealing such a prohibition).

The epistemic world has changed continually as well. Early modern thinkers had scientific models and mathematical tools available to them unknown to the ancients. Since the early days of science, as Ian Hacking has documented in his (1990), it has been a slow path by which statistical reasoning came to be accepted as a source of knowledge. And our growing (if still nascent) science of the inner workings of human cognition, and the chunky

realism about mental mechanisms underlying our behavior that it requires, raises the question of the epistemic status of those faculties.

The objection might be offered at this point that, to perform an analysis-by-imagined-reconstruction on our epistemic norms, we would have to imagine ourselves in some sort of Archimedean point outside of the norms themselves. And how could a rational outcome arrive from beyond our norms? Bealer argues in his (1998) that we cannot construe as rational any radical rejection of our 'standard justificatory practices', since to reject those practices wholesale would be to reject the epistemic itself. One cannot step outside the realm of justification altogether and have a justified theory of that domain. I will grant the argument, because the sorts of critical revision of our epistemic institutions I am proposing do not require us to stand utterly outside of our standard justificatory practices. For those practices themselves include the means for reflection upon them. It is part of our standard practices to evaluate and re-evaluate our norms and procedures. (That is what philosophy is for, one might say.) It is for that reason that we no longer use trial by ordeal in our legal system, and we do now allow for the legitimacy of some probabilistic forms of inference while relying far less on appeal to Church authority in scientific matters.

So, how does this form of neopragmatism -- I will call it *reconstructive neopragmatism* (RN) -- score according to our desiderata? It is comparable to IDR and MN with regard to truth-conduciveness, in that only a skeptic would deny our ability to examine our goals and evaluate what rules would promote them. Normativity is also obviously

satisfied, and there is no obvious conflict with minimal naturalism. The reconstructive element accommodates progressivism, in that it allows for the possibility that the norms developed in the past may not be the best for our future; at the same time, radicalism is only possible if we are currently deeply wrong about our goals and what would promote them. And interdisciplinary comportment is guaranteed as well, since we will recruit psychological and social science to aid in our norm engineering.

RN's status on dialectical robustness and relativism is a bit trickier, because of the question of what our epistemic goals are. Once some set of epistemic goals is fixed, there is no worry about the robustness of our discussions about how best to satisfy those goals, for we have ample experience with such means-ends reasoning. And we may reasonably expect significant agreement about those goals -- the acquisition of true beliefs, for example. Yet there may be irresolvable disagreement about other of the basic goals themselves.[5] Under such circumstances, RN would admit of a degree of relativism as well, as different epistemic goals will generate different norms as well. But I hope that this limited relativism will seem plausible: if two groups truly wanted fundamentally distinct things for and from their beliefs, then perhaps it is appropriate for them to be governed by divergent rules. Moreover, that relativism about ultimate epistemic goals mitigates the worry about stubborn disputes concerning those goals. Where fruitful discourse cannot solve disputes about goals, relativization can dissolve them.

So, at worst, RN gets a solid rating on five of the desiderata, with a mixed rating on the last two. Compared to the poorer-scoring IDR and MN, RN promises hope for the future of epistemology. The rest of this chapter will aim at showing how we might begin to fulfill that promise.

II Reconstructive neopragmatism, internalism, and the *a priori*

The best argument for the viability of RN, and the only way to show that our hopes for it are not in vain, is to demonstrate it in action. I will do so here by applying this method to a connected pair of central epistemological topics: the internalism/externalism debate about epistemic justification, and the question of the existence and extent of *a priori* justification.

II.A What is justification for?

Since we are applying the RN methodology here, we shall address the issue of internalism/externalism not through intuition-mongering, but rather by asking: why should we care about our beliefs' having justification, instead of only caring about whether a belief is simply true or false? One can find in the literature at least two basic reasons we might want our beliefs to be justified in addition to being merely true. In a nutshell, they are *diachronic reliability* and *dialectical robustness*. (One may feel a bit of déjà vu from Part I; but it should not be surprising that two metaepistemological desiderata should also resemble two desiderata for justification itself.) I will call these the 'DR desiderata'. Let me articulate them here, and then discuss how they relate to this issue.

First, suppose that you were about to make a momentous, life-changing, potentially life-ending decision. You have no time for reflection or research, but must instead decide immediately. Which would you prefer to have guiding your actions at that moment: a belief-set that is mostly true but generally unjustified, or a belief-set that is mostly justified but generally false? Prudence prefers the former -- all the justification in the world comes to naught if you make the wrong decision, and without further reflection it is the current actual truth or falsity of your beliefs that will determine the outcome of your action.

Thankfully, we are rarely in such circumstances (except, perhaps, whenever we cross the street). Rather, when confronted by a decision of any importance, we can seek out further relevant evidence for our choice. But then we want to be sure that we make any and all appropriate revisions to our belief-set as a result of this information-harvesting. At this point justification becomes key. If we are able to trace the rational relations amongst our original beliefs, and between those beliefs and the new evidence, then we can adjust our doxastic condition, re-apportioning our epistemic resources as needed. This is not doxastic voluntarism, but just our ability to refocus our concentration and redirect our investigations. As our awareness of the facts changes -- or when the facts themselves change -- we want our beliefs to change with them, and the justificatory links between our beliefs are the channels along which such changes can rationally propagate. This desire for our beliefs to have this kind of across-time accuracy is one good reason to desire justification for our beliefs.

We also desire that we be able to integrate our beliefs with those of others. Polonius tells Laertes to neither a borrower nor a lender be; but *pace* such epistemic Polonii as Descartes, we borrow and lend each other's cognitive abilities, expertise, and information all the time. Our intellectual lives would be hopelessly impoverished otherwise, to our personal and collective detriment.[6] In seeking out others' testimony we require that we interweave their beliefs with our own. If two heads are truly better than one, then the outputs of each pair of eyes and ears had better be able to find their way into the other pair's cranium, and vice versa. And this process of informational exchange needs to be harmonious. By establishing such processes, you and I can form a community of cognizers, and can avail ourselves of an epistemic division of labor and investigate together, as a 'we', the nature of our world. To do so requires something more than just being willing to take others' testimony at its face, for we need also a way to resolve conflicting testimonies between co-investigators.

Here, I would argue, justification plays an important role. Each of the various contestants must put forward her justifications for her claims, which can then be interrogated by the other contestants. We can use the overall quality of their various justifications as a criterion for making our choice between contrary theories, each of us deciding which theory is the most justified so that each can believe it herself. To the extent that we all opt for the same theory, we can be said to have come to believe it, ideally. But even if none of us are able to change each other's minds, the next generation of graduate students will be able to make those judgments, and vote with their feet. Norms of justification are also required to

distinguish between legitimate and unwarranted challenges to a theory, and sufficient and insufficient responses to the legitimate challenges. Our justificatory practices should provide an infrastructure for our investigative communities. We thus further desire, in addition to general reliability, that our epistemic practices be *dialectically robust*.

So our justificatory norms ought to promote diachronic reliability and dialectical robustness: the DR desiderata. There may be other key desiderata for justification, and I would welcome anyone's making the case for such. But I trust that they are sufficiently central that we will not acquire too distorted a picture of our justificatory norms, if we view them through the lens of only those two values.

Having articulated two of our purposes in having norms of justification, we can next attend to what epistemic principles do -- or do not -- follow from them. It might very reasonably be asked first, though, whether these two values exhaust our epistemic desiderata. I grant that we will probably have other values that should be expressed in our justificatory norms. Indeed, I in no way intend these arguments to be considered the final stage of a pragmatist analysis of justification, but very much a first step. This method can work only by engaging in serious discussion about what our ultimate epistemic desiderata really are, and one can accept everything I have said up to this point and still get different results, if different epistemic goals are put forward. Nonetheless, diachronic reliability and dialectical robustness are clearly very central desiderata for our norms of justification, so I do not believe that we

will acquire too distorted a picture of our justificatory norms if we view them through the lens of only those two values.

With the DR desiderata for justification in sight, the appeal of a version of internalism becomes clear. It will generally serve the DR desiderata for agents to be able to tell, by means of reflection or introspection, what the sources of justification for their various beliefs are. If we wish to maximize diachronic reliability, it will often be useful for us to be aware of possible evidential grounds of our beliefs (e.g., BonJour 1985; Moser 1985). Holding beliefs on the basis of conscious reasons, or at least being able to provide such reasons upon reflection, allows an agent to direct her own investigations on relevant sources of evidence she does not yet have. She can take conscious control of her doxastic life, and apportion her resources as she sees fit.

Moreover, the agent will be able to take better advantage of information that crosses her path unexpectedly. Suppose that she believes on her brother's say-so that Microsoft stock is doomed to crash, but later learns to her surprise that her otherwise well-informed brother has been radically misled about matters of economics and computers. She can attempt to adjust her doxastic state accordingly, by reading *Fortune* and listening to 'Marketplace', and changing the topic whenever her brother brings up Bill Gates. Hopefully by doing so, she can ultimately replace any significantly false beliefs about that company with more accurate ones. Or, should she in the end concur with her Apple-loving sibling, she will do so in a way that keeps her in touch with the facts: if her justification is internally accessible, then, should she

learn that Microsoft's situation has altered in ways relevant to the justificatory anchors of her belief, she can commence a new investigatory cycle.

Furthermore, it is in general the case that only epistemically available reasons for beliefs can be put into public discourse. As Gilbert writes in her (1994), our ability to investigate and to know as a community depends on our capacity for a 'joint commitment' (246) to the terms of the investigation: each member must acknowledge an obligation to each other member to uphold those terms, and further expects that each other member is similarly obligated to her. The group can accept some proposition as a group only if each member can be explicitly committed to that acceptance, and recognize others' acceptance. So such commitments must be a public affair, that each member can see that they are being upheld by the others, and can indicate that she herself is upholding them. Of course, we can more easily make others aware of such acceptances and commitments when we ourselves are aware of them. Therefore, internalistically-accessible justifications will best subserve this role of maintaining joint investigative commitments.

Our success as an epistemic community will also depend on our mutually co-adjusting our beliefs. If I will not change my mind no matter what you say, and you are equally cognitively stubborn, then we can merely co-exist as cognizers, not truly cooperating. As individual organisms we each possess unconscious mechanisms for belief adjustment, but the deliberations of the investigative community require a more public medium. We must be able to make our justifications visible to each other, but if you are not aware of the grounds of

your belief, then you cannot cite them to me. And, when we can put our reasons out into public discourse, they can be confronted by the good reasons of others, and themselves confront the bad reasons of others -- thus increasing our ability to achieve a harmonious set of community beliefs.

So the DR desiderata motivate some internalism in our neopragmatist analysis of epistemic justification. We should place great value on an agent's capacity to have, within her reflective and introspective grasp, awareness of her bases of justification, and to be able to express that grasp to her fellow agents.[7] But much of the internalism/externalism debate in epistemology revolves around whether our norms require that an agent have this accessibility to what makes her beliefs justified. Paradigm 'externalist' authors have usually been willing to grant, at least, that there is some value to consciously-held justification,[8] but have insisted that we can also have justification beyond the scope of our introspective capacities. So we should next ask whether the DR desiderata can also motivate an *absolutely* internalist constraint on justification, or only something more limited. Ought we impose an exceptionless internal-accessibility clause as a necessary condition on justification?

Our neopragmatism suggests that we should ask whether instituting the norm of such a necessary condition would create favorable circumstances for achieving the DR desiderata. Any proposed strict constraint must satisfy two criteria: (i) it must not place so onerous a burden on our cognition that too little of our epistemic lives can be sustained, and (ii) whatever beliefs and inferences do pass through must be appropriately DR-promoting. If a

proposed constraint fails either condition, then it is inconsistent with our DR interests and should thus be weakened.

II.B Why strict internalism won't get us what we want

If a strict internalism is to be observed, then the source of epistemic justification must be available to introspection and/or reflection. We can divide the set of candidate internally-available sources in three: (a) the processes by which the belief was formed, or inference performed; (b) the belief or inference itself; and (c) some internally-available mental entity distinct from the target belief or inference. We must now consider whether our minds in fact have enough internally-available material, and indeed material of the DR-promoting sort, that would license a strict internalism. If not, then we must be willing to give up that kind of constraint as inconsistent with our epistemic values. If a norm rules most or all of our beliefs as unacceptable, then it will not help us achieve any lasting truth for our beliefs; if a norm silences all or most of our statements to each other, then it will not aid the smooth conduct of our conversations.

First, might our psychological processes be sufficiently open to the inner eye to be considered internally available? This question clearly has a large contingent component, concerning our actual psychological make-up. And indeed the transparency of the mental has seemed a tempting thesis at times in the history of the philosophy of mind. Actually looking at the empirical literature, however, leads one to suspect that there is no comfort for the internalist here. We have overwhelming evidence that many basic cognitive mechanisms are

predominantly unconscious. I will rely here on an *argumentum ex bibliographia*, listing a few paradigm references from the rather vast literatures here. Nisbett and Wilson (1977) launched the notion that humans are simply hopeless at telling what kinds of psychological processes underwrite their own reasoning. Going back at least to Helmholtz, psychologists have seen the conscious component of visual perception as representing a very slim portion of perceptual processing (e.g., Crick and Koch 2000). Researchers such as Evans and Over (1996) and Sloman (1996) investigating deductive reasoning have argued for a 'dual process' theory: an unconscious, preattentive filtering process selects only certain possibilities as salient to the problem at hand, which it feeds to our conscious, analytic processes. Those latter processes cannot function properly when the former do not. Inductive reasoning is even more unconscious; Reber contends in his (1993) that not only do we learn various sorts of contingent correlations better implicitly than consciously, but also subjects provided with explicit hints about the correlation to be uncovered perform worse than subjects relying on more purely unconscious forms of learning! Our epistemic capacities in various domains from basic physics (Spelke 1990; Leslie 1982) to psychology (Premack 1990; Scholl and Leslie 1999) to morality (Shweder and Haidt 1993; Cummins 1996; Darley and Shultz 1990) appear to depend crucially on unconscious processes, whose workings are closed to the very subjects who are deploying them. Empirical scientific investigation builds upon our ability to theorize about natural kinds, which is rooted in a brute capacity for biological essentialism (Atran 1990; Keil 1989).

What about the belief or inference itself, as an internalist source of justification? Do we sometimes have a self-recommending component in our phenomenology, in which certain beliefs or inferences just strike us as necessary, or at least appropriate and credible? A positive answer to that question is not enough -- such intuitive cognitions must also be of sufficient number and scope to provide the requisite internal buttressing to our otherwise unconscious cognitive infrastructure. I will grant here, for the sake of argument, that beliefs based directly on perception and memory have an appropriately phenomenological quality. But clearly those two sources of justification, as centrally important as they are, cannot pull all the weight of our cognition about the world. A vast proportion of our beliefs are about the unobserved (such as the future, or unsurveyable universal generalizations), or even unobservable (such as theoretic constructs or ethical principles). A strict internalism will require still further raw materials from within our introspective purview.

Intuitionists such as Bealer or BonJour have generally focused on the existence and cogency of rational intuition (or, in BonJour's terms, 'rational insight'). They have taken as their opponents the radical empiricist who would deny that we have any such completely non-perceptual forms of justification. I think their arguments are generally sound -- I too would endorse a substantive category of the *a priori* (see below). But the question at hand is whether a strict internalism makes sufficient room for justified cognition, or whether instead it presents in fact too small a lot upon which to build cognition's house. If intuitionist

internalism is to succeed, it must also address this further issue of not just the cogency but moreover the sufficiency of rational intuition.

And here I fear that they run aground. Perhaps we do have some such intuitions, in which the propositions have a phenomenology of rational compulsion -- one of Bealer's favorite examples is DeMorgan's Laws. But these cases only go so far, and do not extend much beyond certain simple forms of deductive and mathematical reasoning. Not even all of those forms are generally available: the great majority of subjects do not recognize the validity of instances of modus tollens, for example. And our intuitive powers are especially weak in the domain of non-deductive inferences. BonJour strives valiantly in his (1998) to formulate a principle of induction that elicits this phenomenology, but they lack (for me, at least) that sense of rational necessity.[9]

Moreover, many propositions that have this phenomenological glow about them are just not the right kind. Our actual inductive intuitions are generally a mess, for example. Indeed, instances of the gambler's fallacy seem more compelling on their face than any formulation of a proper principle of induction (Garnham and Oakhill 1994)! Again, I am not arguing from the existence of such undesirable intuitions that we can never rely on intuitions; rather, I am claiming that we might not want to rest the entire weight of *a priori* justification upon them.

I do not doubt that BonJour is sincere in his declaration that his formulations strike him as 'sufficiently obvious to require little discussion' (BonJour 1998: 208). I similarly do

not doubt Bealer, when he claims in his (1996) to have a rational intuition to the effect that intuition is a good source of evidence. But again I do not find myself sharing the intuition. I take this divergence to be emblematic of a more general lack of phenomenological univocity, which points to a further difficulty with appealing to intuitions. As discussed earlier in the context of the case against IDR, we lack good tools to resolve conflicts in intuition. So not only are our intuitions insufficient in number and scope, and frequently inconsistent with diachronic reliability -- but also a reliance on intuition would be deleterious to dialectical robustness as well. The internalist needs to argue that we can take the epistemic weight off of our tacit processes and place it on an introspectively-available structure. But given how massive that epistemic architecture is, intuition simply presents too thin a dialectical reed.

Another internalist strategy, which we might term inferentialism, seeks to find an argument to the conclusion that the outputs of our tacit psychological processes are rational and/or reliable. I have no doubt that some philosophers can make such a case -- though it will not be a simple matter, given the presence in the literature of many arguments to the contrary (e.g., Stein 1996), and it is indeed hard to see how it could be done noncircularly.[10] But put such concerns aside. What I question, however, is the relevance of such an argument to the issue at hand. For if such arguments are only available to the trained philosopher, then this strategy would entail that only we philosophers can have justified beliefs! I am here following Goldman (1999: 13), who has argued recently that '[i]t is very unlikely that someone who has never studied philosophy could produce a satisfactory justification for the

reliability of his inductive or deductive inference procedures. To conclude from this, however, that ordinary, philosophically untrained people have no inferential warrant would be a dramatic capitulation to skepticism.' In our terms here, it would be a complete surrender of the DR desiderata. Surely it would be nice for everyone to be able to consciously defend their unconscious cognition. But it cannot -- given the DR desiderata and the sorts of limited creatures we are -- be an epistemic necessity.

Note that I am not suggesting that it is a problem that ordinary folks in general cannot, as a matter of epistemological principle, form an argument to the effect that they have epistemically successful faculties of non-empirical cognition without depending on such faculties in the argument. It is not obvious that the defender of a strict internalist condition needs to take on such baggage, which seems to require something tendentious like a JJ principle. What I am suggesting simply is that ordinary folks in general cannot, as a matter of psychological fact, form such an argument at all.

The internalist might try to reply that, even if most ordinary folks are utterly unprepared to argue for the epistemic success of their non-empirical cognitive faculties in particular, nonetheless they are surely capable of launching a general defense of their own cognitive capability. They could observe that they are basically successful in navigating the world -- they do not (usually) walk into walls, or forget to come in out of the rain, and so on -- and then they could take this success as evidence that they are at least moderately decent

believers. *A fortiori* they could conclude that they are at least moderately successful in their non-empirically-based cognition. Call such an argument the 'Global Justifier' argument.[11]

But such an internalist reply would allow the ordinary folk to prove too much, and render the internalist constraint useless. I have been arguing that we would not want too strong a constraint, lest we place justification beyond the reach of most if not all humans. But we also have no use for too weak a constraint. Recall that the motivation for an internalist constraint was to require agents to remain in close cognitive contact with their specific sources of justification, such as our earlier example of the woman with the Microsoft-impaired brother. Internalism would not help her, if by relying on the Global Justifier she could satisfy the constraint without thereby improving her doxastic state, as measured by the DR desiderata. If all agents could make use of one big argument that justified all of their cognitive activity, then the original motivations for an internalist constraint would be vitiated. So we cannot allow merely possessing or deploying the Global Justifier argument to be sufficient for the justification of all of one's beliefs. To do so would be inconsistent with our goals in setting forth epistemic norms. Just as we must avoid setting the justificatory bar too high, lest no beliefs pass, we must also eschew placing it too low, lest all beliefs pass -- without regard for whether they are the sorts of beliefs that the DR desiderata would have us promote.

II.C Prospects for a less-strict internalism

So at the end of day we will not be able to endorse a strict internalist constraint on justification: none of the three possible avenues of justification are sufficient, or sufficiently internalistically-acceptable. Our justificatory norms must allow some significant loophole by which enough of the beliefs produced by our unconscious cognitive infrastructure can be allowed to pass, even when introspection and reflection prove unable to produce for us any reasons for holding them. At the same time, we cannot allow ourselves in general to hold beliefs on no reason at all, merely because no reasons are forthcoming. A lack of reasons can hardly itself be a reason for holding a belief! That would be a sort of epistemic suicide, and would fail the DR desiderata as badly as strict internalism does. Rather, we require that the DR motivations still be addressed. Under what conditions can X's true-but-brute belief that p still comport with our desired diachronic reliability and dialectical robustness?

First, if the fact of whether-or-not-p were itself unlikely to change, at least in any environment X has any chance of finding herself in, then we needn't worry about X maintaining true p-beliefs over time. If p's truth is not at all variable, then it won't be hard to track that truth. It must also be unlikely that X will ever come across evidence against p, either, since without further reasons for p, such evidence should ideally compel her away from her initial belief. So X has less of a need for consciously-available evidence for p, since her right belief will not be subject to the epistemic vagaries of more changeable facts.

Diachronic reliability would be guaranteed. When such conditions are met, let us call a belief that p *epistemically stable*.

Second, if most everyone else in X's epistemic community is also endowed with a similarly true-but-brute belief that p, then there will be no difficulty in integrating X's p-belief with her co-investigators'. X is unlikely to be challenged about her belief that p, since everyone else holds that belief as well, and so X has no need for any internally-accessible defense of p. Thus, under such conditions, dialectical robustness is not threatened by the absence of citable reasons. When this condition is met, we can consider a belief that p *epistemically universal*.

So epistemic stability and universality together provide a sketch of a good candidate for the principled loophole we need in any internalist constraints on justification. These are generally rare doxastic properties, but rather common in an epistemic domain of great epistemological interest: the *a priori*. Epistemic stability and universality are easily found in such paradigm domains of *a priori* justification as arithmetic. That domain is clearly quite stable, since its truths are necessary; and we have been fortunate in that arithmetical dissention occurs extremely rarely. Indeed, necessity is a traditional hallmark of the *a priori*, and necessity entails epistemic stability; and going as far back as Plato, and still more recently (e.g., Antony 2004), innateness has been cited as a sign of the *a priori*. So, even though rationalists might have thought themselves naturally allied with intuition-based methodologies, our reconstructive approach has revealed perhaps a closer relationship

between rationalism and neopragmatism. In the areas in which we most clearly wish to claim *a priori* justification, our neopragmatist considerations suggest that we least clearly need to insist on imposing internalist strictures on that justification. (One might wonder, though, whether the results cited above might indicate that philosophy itself demonstrates insufficient universality as a domain for such exemption.)

III Conclusion

We began by considering possible good-making features for a method in epistemology, and saw that reconstructive naturalism might be more desirable than either intuition-driven romanticism or metaepistemic naturalism. We applied RN to our norms of justification, and posited two general purposes for justificatory norms, diachronic reliability and dialectical robustness. These DR desiderata in turn motivated a general internalist constraint on justification. However, our best reasons for postulating the category of the *a priori* in the first place turned out to allow, perhaps even require, room for externalist sources of justification as well, in the special case of *a priori* knowledge. We can thus see how, on such age-old topics as internalism and rationalism, RN makes new positions available -- and, moreover, provides a new way to make a case for such positions. By asking what we want from our epistemology, reconstructive naturalism can give us more of what we want from our metaepistemology as well.

References

Allen, C. (forthcoming). 'Is Anyone a Cognitive Ethologist?' *Biology and Philosophy*.

Alston, W. (1989). *Epistemic Justification: Essays in the Theory of Knowledge*. Ithaca, NY: Cornell University Press.

Antony, L. (2004). 'A Naturalized Approach to the *A Priori*'. *Philosophical Issues*, 14: 1-17.

Atran, S. (1990). *Cognitive Foundations of Natural History: Trends in the Anthropology of Science*. Cambridge: Cambridge University Press.

Bealer, G. (1996). '*A Priori* Knowledge and the Scope of Philosophy'. *Philosophical Studies*, 81: 121-42.

----. (1998). 'Intuition and the Autonomy of Philosophy', in DePaul and Ramsey (1998), 201-39.

Blackburn, S. (1984). 'Knowledge, Truth, and Reliability'. *Proceedings of the British Academy*, 70: 167-87.

Boghossian, P. (2003). 'Blind Reasoning'. *Supplement to the Proceedings of The Aristotelian Society*, 77: 225-48.

BonJour, L. (1980). 'Externalist Theories of Empirical Knowledge'. *Midwest Studies in Philosophy*, 5: 53-73.

----. (1985). *The Structure of Empirical Knowledge*. Cambridge, Mass.: Harvard University Press.

----. (1998). *In Defense of Pure Reason: A Rationalist Account of* A Priori *Justification*. Cambridge: Cambridge University Press.

Craig, E. (1990). *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford: Clarendon Press.

Crick, F. and Koch, C. (2000). 'The Unconscious Homunculus', in T. Metzinger (ed.), *The Neuronal Correlates of Consciousness: Empirical and Conceptual Questions*. Cambridge, MA: The MIT Press, 103-10.

Crowley, S. (manuscript). 'An Inadequate Epistemology of Animal Cognition'.

Cummins, D. (1996). 'Evidence For the Innateness of Deontic Reasoning'. *Mind and Language*, 11: 160-90.

Cummins, R. (1998). 'Reflections on Reflective Equilibrium', in DePaul and Ramsey (1998), 113-27.

Darley, J. and Shultz, T. (1990). 'Moral Rules: Their Content and Acquisition'. *Annual Review of Psychology*, 4: 523-56.

DePaul, M. and Ramsey, W. (eds.) (1998). *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*. Lanham, MD: Rowman & Littlefield.

Evans, J. St.B. and Over, D. (1996). *Rationality and Reasoning*. Hove, UK: Erlbaum.

Fumerton, R. (1988). 'The Internalism/Externalism Controversy', in J. Tomberlin (ed.), *Philosophical Perspectives, 2*: *Epistemology*. Atascadero, Calif.: Ridgeview Publishing, 443-59.

Garnham, A. and Oakhill, J. (1994). *Thinking and Reasoning*. Oxford: Blackwell.

Gilbert, M. (1994). 'Remarks on Collective Belief', in Schmitt (1994*a*), 235-56.

Goldman, A. (1979). 'What is Justified Belief?', in G. S. Pappas (ed.), *Justification and Knowledge: New Studies in Epistemology*. Dordrecht: D. Reidel, 1-23.

----. (1999). 'A Priori Warrant and Naturalistic Epistemology', in J. Tomberlin (ed.), *Philosophical Perspectives, 13: Epistemology*. Oxford: Blackwell, 1-28.

----. and Pust, J. (1998). 'Philosophical Theory and Intuitional Evidence', in DePaul and Ramsey (1998), 179-97.

Greco, J. (1999). 'Agent Reliabilism', in J. Tomberlin (ed.), *Philosophical Perspectives, 13: Epistemology*. Oxford: Blackwell, 273-96.

Haack, S. (1993). *Evidence and Inquiry: Towards Reconstruction in Epistemology*. Oxford: Blackwell.

Hacking, I. (1990). *The Taming of Chance*. Cambridge: Cambridge University Press.

Hintikka, J. (1999). 'The Emperor's New Intuitions'. *Journal of Philosophy*, 96: 127-47.

Jackson, F. (1998). *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Clarendon Press.

Katz, J. (1998). *Realistic Rationalism*. Cambridge, Mass.: The MIT Press.

Lehrer, K. (1997) *Self-Trust: A Study of Reason, Knowledge, and Autonomy*. Oxford: Clarendon Press.

Keil, F. (1989). *Concepts, Kinds, and Cognitive Development*. Cambridge, Mass.: The MIT Press.

Kim, J. (1988). 'What is "Naturalized Epistemology?"', in J. Tomberlin (ed.), *Philosophical Perspectives, 2: Epistemology*. Atascadero, Calif.: Ridgeview Publishing, 381-406.

Kornblith, H. (2002). *Knowledge and Its Place in Nature*. Oxford: Clarendon Press.

Leslie, A. (1982). 'The Perception of Causality in Infants'. *Perception*, 11: 173-86.

Machery, E., Mallon, R., Nichols, S., and Stich, S. (2004). 'Semantics, Cross-Cultural Style'. *Cognition*, 92: B1-B12.

Moser, P. (1985). *Empirical Justification*. Dordrecht: D. Reidel.

Neta, R. (forthcoming). 'Epistemology Factualized: New Contractarian Foundations for Epistemology'. *Synthese*.

Nichols, S., Stich, S., and Weinberg, J. (2003). 'Metaskepticism: Meditations in Ethno-epistemology', in S. Luper (ed.), *The Skeptics: Contemporary Essays*. Burlington, VT: Ashgate, 227-47.

Nisbett, R. and Wilson, T. (1977). 'Telling More Than We Can Know: Verbal Reports on Mental Processes'. *Psychological Review*, 84: 231-95.

Plantinga, A. (1993). *Warrant and Proper Function*. New York: Oxford University Press.

Premack, D. (1990). 'The Infant's Theory of Self-Propelled Objects'. *Cognition*, 36: 1-16.

Pust, J. (2000). *Intuitions as Evidence*. New York: Garland Publishing.

Reber, A. (1993). *Implicit Learning and Tacit Knowledge: An Essay on the Cognitive Unconscious*. New York: Oxford University Press.

Rorty, R. (1989). *Contingency, Irony, and Solidarity*. Cambridge: Cambridge University Press.

Schmitt, F. (ed.) (1994*a*). *Socializing Epistemology: The Social Dimensions of Knowledge*. Lanham, MD: Rowman & Littlefield.

----. (1994*b*). 'The Justification of Group Beliefs', in Schmitt (1994*a*), 257-88.

Scholl, B. and Leslie, A. (1999). 'Modularity, Development and "Theory of Mind"'. *Mind and Language*, 14: 131-53.

Shweder, R. and Haidt, J. (1993). 'The Future of Moral Psychology: Truth, Intuition, and the Pluralist Way'. *Psychological Science*, 4: 360-5.

Sloman, S. (1996). 'The Empirical Case for Two Systems of Reasoning'. *Psychological Bulletin*, 119: 3-22.

Solomon, M. (1994). 'A More Social Epistemology', in Schmitt (1994*a*), 217-34.

Sosa, E. (1991). *Knowledge in Perspective: Selected Essays in Epistemology*. Cambridge: Cambridge University Press.

----. (forthcoming). 'A Defense of the Use of Intuitions in Philosophy', in M. Bishop and D. Murphy (eds.), *Stich and His Critics*. Oxford: Blackwell.

Spelke, E. (1990). 'Principles of Object Perception'. *Cognitive Science*, 14: 29-56.

Stein, E. (1996). *Without Good Reason: The Rationality Debate in Philosophy and Cognitive Science*. Oxford: Clarendon Press.

Stich, S. (1990). *The Fragmentation of Reason: Preface to a Pragmatic Theory of Cognitive Evaluation*. Cambridge, MA: The MIT Press.

Weatherson, B. (2003). 'What Good Are Counterexamples?' *Philosophical Studies*, 115: 1-31.

Weinberg, J., Nichols, S., and Stich, S. (2001). 'Normativity and Epistemic Intuitions'. *Philosophical Topics*, 29: 429-60.

Williamson, T. (2003). 'Blind Reasoning'. *Supplement to the Proceedings of The Aristotelian Society*, 77: 249-93.

Notes

[*] I would like to thank Steve Crowley, Stephen Hetherington, Henry Jackman, Mark Kaplan, Peter Klein, Adam Leite, Peter Markie, Mark McCullagh, Aaron Meskin, Ram Neta, Baron Reed, Ernie Sosa, and the faculty and students at Auburn University, the University of Memphis, and Texas Tech University for their many useful comments and suggestions.

[1] Joel Pust (2000: ch. 1) documents the central evidential role of intuitions in contemporary philosophy.

[2] See Allen (forthcoming), for a discussion both of those worries and some responses to them.

[3] Ram Neta (forthcoming) makes a similar metaepistemological appeal -- though towards different epistemological results -- to Craig's work.

[4] I intend this as more a friendly variant on, than a rival to, Craig's framework.

[5] There are other possible moves -- one could perhaps argue that only one set of goals is truly epistemic -- but I will not canvass them here.

[6] See various essays in Schmitt (1994*a*), including Solomon (1994), Gilbert (1994), and Schmitt (1994*b*); it is also a theme sounded in Alston (1989), Craig (1990), and Sosa (1991).

[7] The exact forms of acceptable reasons may be much broader than asserted propositions. Under many circumstances, simply pointing in the right direction may constitute a sufficient public reason.

[8] Cf. Goldman's willingness in his (1979) to accommodate the intuition that improper use of consciously-held evidence can disable the justification he thinks a reliable process can otherwise confer.

[9] Here's one such allegedly intuitive principle of induction in question: 'In a situation in which a standard inductive premise obtains, it is highly likely that there is some explanation (other than mere coincidence or chance) for the convergence and constancy of the observed proportion (and the more likely, the larger the number of cases in question)' (BonJour 1998: 208). Does this proposition really present us with the sort of luminous phenomenology that the intuitionist requires? It does strike me as a plausible sort of thing to believe, but of course most of the things we believe will strike us as such. And rational intuition ought, one would think, have higher standards than just that very ordinary and ubiquitous appearance of *prima facie* plausibility. Otherwise, relying on it would not help us promote our DR goals -- almost all our beliefs would be 'intuitive', so the internalist constraint would be trivially satisfiable. But see my discussion of the 'Global Justifier' argument below for reasons why such a trivial internalist constraint is worse than no internalist constraint at all.

[10] Cf. the exchange between Boghossian (2003) and Williamson (2003).

[11] Keith Lehrer in his (1997) seems to present something like this argument, when he invites us to accept that we are generally worthy of our own self-trust.