

(uncorrected proofs – to appear in *Episteme* special volume on epistemic relativism)

Jonathan M. Weinberg

Moderate Epistemic Relativism and Our Epistemic Goals

Abstract

Although radical forms of relativism are perhaps beyond the epistemological pale, I argue here that a more moderate form may be plausible, and articulate the conditions under which moderate epistemic relativism could well serve our epistemic goals. In particular, as a result of our limitations as human cognizers, we find ourselves needing to investigate the dappled and difficult world by means of competing communities of highly specialized researchers. We would do well, I argue, to admit of the existence of unresolvable disputes between such communities, but only so long as there is a sufficient amount of fruitful exchange between them as well. I close with some speculation about when it is or is not legitimate to make an “appeal to discipline”: responding to another’s argument by saying something like, “we should do it this way, because we are philosophers (/linguists/psychologists/...), and that’s just what we do”.

Do I contradict myself?

Very well, then, I contradict myself;

(I am large—I contain multitudes.)

--Walt Whitman, *Leaves of Grass*

I. A *Moderate* Epistemic Relativism?

Epistemological discussions of relativism often take as a starting point that any such relativism will be a radical one, upturning much of our standard picture of inquiry and epistemic normativity, denuding the epistemic of its ambitions of objectivity. One can find this background assumption in play regardless of whether an author thinks that this makes relativism something to be feared (e.g., Boghossian) or celebrated (e.g., Rorty). There are those who would *épater le bourgeoisie*, and those among the *bourgeoisie* who are looking to avoid being *épatered*, but both sides agree that relativism is an *épatement*-producing thesis. In this paper, I will reject that starting point, and try to argue for a substantial form of epistemic relativism that, while avoiding radicalism, actually offers a better framework for inquiry than non-relativist approaches can.

We need to start by clarifying what work we want this term “relativism” to do for us. We do not typically extend the term “relativism” to each and every thesis that involves relativizing some set of propositions or other along some parameter or other. Time of day is relativized to time zone, yet this boringly quotidian fact hardly presents a form of

relativism worth getting worked up over. Contextualism, in its various forms, relativizes the truth of knowledge attributions according to various factors such as the goals of the subject or the recent moves in the attributor's conversation. But neither its proponents nor opponents generally consider contextualism to be a flavor of relativism.¹ My sense of what people generally want to mean by "relativism" requires that the relativizing parameter in question be somehow in terms of groups of people, such as culture or gender or educational background or religion or profession or socioeconomic status or language. Let us call this the *demographic* condition on relativisms worth discussing.

Moreover, we could, if we really wanted, describe any standard account of knowledge as relativistic, because the particular propositions that will count as knowledge will vary from person to person, depending on their different beliefs, evidence, etc. But that, too, just isn't an interesting notion of relativism – there's nothing to be gained by extending the term in that way. What such a "relativism" seems to lack is any interesting consequences for conflicts between members of different epistemic groups. In particular, on a relativist account, at least some such conflicts should turn out to be only apparent conflicts. Looking at a case of an epistemic disagreement between members of different groups, a non-relativist will consider that at most one of the parties to the disagreement can be correct, whereas the relativist will allow for the possibility that both are correct, and in some sense deny that the two parties are really disagreeing. The non-relativist finds only conflicts that she would like to settle by finding the correct

¹ Though see Neta (this issue). My claim here should not be seen as challenging Neta's arguments there, since he is largely trying to refute a certain strategy for arguing against relativism, and his contextualist example fits the form that that strategy takes as its target.

side; the relativist finds at least some that cannot be resolved rationally, but only defused or perhaps battled out politically. Let us call this the *unresolvability* condition on relativisms worth discussing.

The fear (or radical hope) is that it may be that *all* inter-group conflicts receive such deflationary treatment. Opponents of relativism have a well-justified fear that it is a short – perhaps immediate – slide from relativism to anything-goes-ism. Their fear might be summarized as: “If the absolute is dead, then anything is epistemically permitted.” Or, as Boghossian puts it (to specify his opponent), “[t]here are many fundamentally different, genuinely alternative epistemic systems, but no facts by virtue of which one of these systems is more correct than any of the others” (Boghossian 2006, 76). Their fears are surely at least somewhat justified, given that some proponents of relativism proclaim this consequence themselves, such as adherents to the “strong programme” in science studies; see also, e.g., Feyerabend (1975). Such an outcome could be called *radical unresolvability*, in that it seems to claim that *no* conflicts can be settled rationally. We might then ask whether there might be room for a relativism that, while meriting the name, does not have consequences quite so monotonously arational. Let us use *moderate unresolvability* for forms of relativism that allow for some sincere disagreements across the relativized parameter to be merely defused, but that also allows some, perhaps many, such disagreements to still persevere as legitimately substantial. We therefore can define as *moderate epistemic relativist* any view that asserts moderate unresolvability along some demographic lines.²

² Talk of “unresolvability” may connote that this relativism could be simply

epistemological epistemic relativism, in which at most one of a set of divergent groups

My goal in this essay is to defend the cogency, and indeed legitimacy, of moderate epistemic relativism (MER). In particular, members of different research communities (e.g., different academic disciplines, or movements, camps, even labs within disciplines) can legitimately operate with different epistemic norms, but the resulting unresolvability may be usefully moderate without succumbing to the dangers of radicalism.

II. Relativism and Metaepistemology

Given that MER has not generally been on the epistemological map³ – so much of the ink having been spilled about its more radical cousin – one might well wonder whether it is a real option in the first place. Are there perhaps good reasons to reject it outright?

One reason to do so would be if one thinks, as several epistemologists in good standing

may be strictly speaking correct, but the divergence is not rationally adjudicable. This is not my intention; rather, I have in mind a form of fully *semantic* epistemic relativism, which relativizes the truths themselves about which epistemic norms correctly apply to different groups. See White (this issue). It's not just that *we* can't resolve such disputes; the world contains no facts that could do it, even in principle.

³ See, e.g., Boghossian (2006, 94n5): “It is easy to see what might motivate someone to take seriously the idea that there are *no* absolute epistemic truths of any kind; it is much harder to see what would motivate the moderate view that, while there are some absolute epistemic truths, there are many fewer than we had been inclined to suppose...”

(emphasis original). One can find some exceptions, of course; perhaps Putnam's internal realism is one.

have thought, that principles of epistemic normativity have a kind of strong modal status that is inconsistent with any relativism at all. Necessary truths about the very nature of belief or rationality or knowledge, after all, are not the sort that it would make much sense to relativize according to group membership. And so it is likely that, if one thinks of the relevant epistemic norms as somehow given to us through direct insight into the structure of the universe, then any sort of relativism about those norms will be unsustainable.⁴ Nonetheless, there are reasons to doubt that anything like an apriorism *this* strong can be true. For starters, it seems to completely miss the ways in which our best methods have evolved slowly over a long, civilization-wide process guided in no small part by empirical discoveries and trial-and-error. But I will not push that line of argument here. Rather, I will exploit the fact that the kind of principles that are sometimes taken to be knowable through such intuition are palpably neither rich nor specific enough fully to guide our epistemic lives. It is, I grant, at least arguable that we have direct access to propositions like the general form of a principle of induction⁵, or

⁴Of course, it is conceivable that the norms thereby revealed could have a relativistic structure. But the norms typically cited have not been such.

⁵ E.g., Bonjour (1998, 208): “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).”

when some particular proposition can serve as a reason for another.⁶ But a great many of the epistemic decisions that matter to us are simply not addressed by anything we can so intuit. The propositions or principles that might be made available to us in this way are going to be fairly simple, general, and straightforward, if for no other reason than the limits of human attention and memory. As such, they are perhaps relevant to, but insufficient to answer, complex questions about what the vast and not-entirely-consistent evidence available says about the nature and extent of global warming, or the structure of human cognitive architecture, or whether gerrymandering is creating too many unbalanced partisan districts for the health of our democracy. And they are perhaps relevant to, but insufficient to answer, questions about how various of our investigations ought to proceed: under what circumstances and for what sort of inferences is fMRI evidence appropriate; or reasoning in evolutionarily teleological terms; or corpus data in linguistics; or the application of asymptotic theory in econometrics; or stationary replication in experimental economics; and, indeed, all use of laboratory experiments in economics⁷ more generally; and so on. So even if there is *something* to an apriorist metaepistemology, there's not *enough* to it. And that is all I need, for my purposes, to leave open the metaepistemological possibility of at least some sort of relativism about the rather large set of epistemic matters for which answers from armchair reflection will not be forthcoming.

⁶ E.g., Pollock (1987): "I am appeared to as if P" is a prima facie reason for me to believe P. See also Fumerton (1995) on our cognizing probability relations between propositions.

⁷ And, for that matter, in philosophy.

If I am not going to operate within a metaepistemological framework of entirely *a priori* access to necessary epistemic truths, then one may well ask just what kind of metaepistemological framework I am going to opt for. Since my goal here is just to make a plausibility argument for a form of MER, I will feel free to adopt without much argument my own preferred metaepistemological framework, which will render the relativism in question plausible. It will then be enough for my purposes if the relativism I argue for is not obviously inconsistent with the other main extant metaepistemologies aside from a strong form of apriorism that would insist that *all* principles of epistemic normativity must be discernible *a priori*. There are various extant options, including reflective equilibrium of intuitive epistemic principles and specific epistemic judgments; or a thoroughgoingly scientific naturalism of Kornblith and (some readings of) Quine; or an Austinian report from inside of our ordinary language at work.

The metaphilosophies of ordinary language and reflective equilibrium can both be handled easily. Ordinary language epistemology (or, it might perhaps be better called, ordinary *practice* epistemology) is clearly consistent with MER, since if there are different communities with different operative epistemologies, then the “what ought we say?” judgments of agents from those different communities will reflect appropriately different norms. Similarly, as Stephen Stich has argued,⁸ members of communities with different operative epistemologies may well find themselves converging on different reflective equilibria. Should such a diversity of communities obtain, however, there may well be significant areas of overlap, too, and so on these metaphilosophies, at least some

⁸ See, e.g., Stich (1988).

disagreements across communities may still be substantial, even if not all of them would be. I am not claiming that a community-relative version of MER is entailed by these metaepistemologies, since there may not really be communities with sufficiently divergent epistemic norms.⁹ My point here is just that MER is not obviously inconsistent with either set of methods.

It's a little harder to see what naturalism, as an epistemological methodology, can or should say about relativism about epistemic norms. For starters, it's just not clear what naturalism has to say about epistemic norms at all, even if we reject arguments like those of Kim (1988), Almeder (1998), or Bonjour (1994) to the effect that naturalism leads to a kind of norm-eliminativism. In his recent *Knowledge and Its Place in Nature*, Hilary Kornblith contends that knowledge is a natural kind whose essence – reliably-produced true belief – is revealed to us by scientists, in particular biologists studying animal behavior. Norms *per se* do not make an appearance, but he does consider the related question of the value of knowledge: such reliably-acquired true beliefs will be instrumentally valuable to any creature which can use such beliefs to pursue its goals, whatever those goals may be. The account of this value, though, is not itself a chapter in the naturalist analysis of knowledge itself, as it is not something revealed in virtue of uncovering the scientific essence of knowledge. It follows from the naturalist account, *plus* some other facts about what is valuable. So naturalism by itself may simply be silent on the question of norms.

⁹ Though see Weinberg, Nichols, and Stich (2001) and Maffie (2000) for evidence of cross-cultural epistemological divergence.

Nonetheless, I will offer a sketch of an argument for how a Kornblithian metaepistemology can find room for MER. How reliable a given mechanism is depends on what environment it is operating in.¹⁰ Epistemic norms can themselves be evaluated for their truth-conduciveness, and accordingly, Kornblithian naturalism might suggest that people operating in different environments should operate according to appropriately different norms. This is not yet a demographic form of relativism, until we recognize that “environments” are not just physical locations but also cultural and institutional – just think of the question of the reliability of testimony as applied to the non-geographically-localizable world of the internet. Then the possibility of a reliability framework for a sort of social relativism about epistemic norms can come into view. As with reflective equilibrium and ordinary language, any such relativism need not be radical and complete. For there may still be some epistemic strategies that have significant cross-niche reliability, and thus naturalism may involve a relativism that is moderate nonetheless. Now, one might suggest that this would all really be a form of non-relativism, because everyone is being assigned something like the norm, “believe in accord with whatever methods, mechanisms, rules, etc. are reliable in your circumstances”. But that instruction, like those of the apriorist, is too contentless to serve as the extent of our epistemic norms. Let me also note that, as above, I am not arguing that naturalism does anything like *entail* MER; I only need to show that it does not need to be unfriendly to it.

III. Reconstructive Neopragmatism

¹⁰ See, e.g., Sosa (1991).

So, a number¹¹ of major extant metaepistemologies are all consistent with MER, some rather directly (as with ordinary language philosophy) and some by having methods that fail to address all our epistemological normative needs (as with apriorism). However, being not-obviously-inconsistent with candidate methods is, of course, a pretty low standard for a claim to meet. So, for the next stage of the argument, I will appeal to my own preferred metaepistemology, and from within its framework, I will offer a more sustained and direct positive case for MER.

Edward Craig, in his *Knowledge and the State of Nature*, rejects the traditional conceptual analyst's question *under what circumstances do we attribute or withhold this concept?* and replaces it with the more fruitful question, *for what purposes do we have this concept in the first place?* In his words, “[i]nstead 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 1990, 2). Craig’s suggested method for understanding what a concept like KNOWLEDGE does for us is to contemplate, in the manner of a classical social contract theorist, what sort of norms might have been agreed upon by some hypothetical past pre-epistemic community. My preferred metaepistemology is heavily indebted to

¹¹ I do not pretend that this exhausts current methods in epistemology; at a minimum, virtue-theoretical and formal approaches have not been addressed here. But I do not see how either obviously has anything to say one way or the other about the likes of MER. (Though it may be helpful to show that relativisms can be formally coherent; see Hales 1997.)

Craig's redirection of epistemology to the uses of our concepts, but instead of a historically-oriented approach, my *reconstructive neopragmatism* is future-directed. If we wish to know what our norms governing a certain practice should be – such as our practices with using words like “knows” or “is justified”, or the norms guiding different forms of inquiry – we should not be hostage to even a hypothetical past version of ourselves. Rather, we should make the best use of what we take ourselves to know now about the sorts of creatures we are, about what sorts of goals we would like to have for the practice in question, and about the sorts of institutions and practices that have (and have not) worked for us in attaining such goals. The method I am suggesting would ask us to contemplate, not from where our epistemic norms might have originated, but to where we would like them to go; hence, “reconstructive”. And we answer that question by reference to what our goals are in having such epistemic norms; hence, “neopragmatism”. Although reconstructive neopragmatism may take our normative vocabulary as its direct target for analysis, it need not do so, as the norms governing such terms are only a part (albeit an important part) of our normative apparatus. Indeed, my project in this paper will not be particularly tied to any such piece of vocabulary.

The method has two distinct stages. First, whatever practice it is for which one is seeking to reconstruct the norms, one must begin by determining what the goals of that practice ought be. Second, one can thereupon determine what norms, for creatures like us in circumstances like ours, will best serve those goals. In an earlier (2006) paper, I considered norms of justification, for which it made sense to consider two main desiderata: *diachronic reliability* and *dialectical robustness*. The former is our goal of getting a high proportion of true beliefs, and keeping them once we have got them (hence

a *diachronic* reliability). The latter is our goal, as epistemically social creatures, of being able to have fruitful dialogue of epistemic import: to be able to exchange testimony with each other, to inform and be informed, to correct and be corrected. We want to minimize disagreements, and to be able to make good use of what disagreements we do have, by resolving them, or coming to an understanding, or at any rate not just having to ignore each other in silence.¹² In the second stage of the reconstruction, given the goals from the first stage, we ask: given the sorts of creatures we are, with the sorts of tools, practices, and social arrangements that are available to us, what norms would best serve those goals? If we could re-constitute our epistemic norms today knowing what we now know, how ought we choose to do it? (Of course, our current norms would remain in play during such deliberations. It would be folly to attempt such a reconstruction from some mythical frictionless position with no operative norms whatsoever. This method is, therefore, not meant to be a method for responding to a skeptic about epistemic normativity itself.) In practice, these two stages will not be totally distinct, and we may come to understand the full complexity of our epistemic goals only in the process of attempting to articulate some subset of them, and see what norms we would engineer to serve them. The order of RN's two stages is thus more logical than chronological.

¹² It might be argued that, for our norms of justification, we only have the latter goal as instrumental to the former. I don't think this is true, but even if it is, dialectical robustness is so completely essential to our lives as social epistemic agents that there is no harm in putting it basically in parallel with reliability.

There are other consequentialist approaches to epistemic normativity already on offer, so I might clarify my RN approach by contrasting it with two of the more famous ones.¹³ Larry Laudan's "normative naturalism" construes epistemic and methodological norms as hypothetical imperatives, such as "if one wants to develop theories which are very risky, then one ought to avoid ad hoc hypotheses" (Laudan 1987, 24). Laudan thus collapses epistemic rationality into a variety of instrumental rationality; e.g., "Good reasons are instrumental reasons; there is no other sort" (Laudan 1990, 320). But RN has no such commitment. To be guided by a norm, one is not required to make any sort of reference, explicit or implicit, to the goals that that norm is meant to serve (though it is in no way forbidden to do so, either). One may simply acquiesce in the norms of one's own community fairly unreflectively, for example. More satisfyingly, one may recognize with one's peers a mutual commitment to those norms, and thereby allow the collective commitments of the community to be its own source of normativity. It follows *a fortiori* that the norms need not be in the form of conditionals with various goals as the antecedents.

RN also resembles Alvin Goldman's social veritism, articulated in his (1999), in that we set a metric of epistemic goodness, and evaluate practices according to that metric. They also share the important emphasis on connecting epistemic norms with the real, empirically-ascertainable circumstances and capacities of human beings, with our particular competences and limitations. One key difference, though, is that Goldman's

¹³ Another philosopher working in a Craig-inspired framework is Ram Neta; see, e.g., Neta (2006). Our particular methods are sufficiently distinct that it requires no special disentangling here.

framework requires that we view practices in much the same way we would view mechanisms or instruments, as the source of cognitive outputs, with all normativity adhering to the quality of those outputs. Goldman's picture, like Laudan's, makes it hard to view oneself as fully participating in one's epistemic norms, as an epistemic agent. It would be like saying that one is participating in one's microscope or thermometer. RN allows for a more satisfying account of our relationship to our own norms – not as instruments external to our identities as epistemic agents, but as in large part constituting our roles as such agents. Who we are as agents cannot be fully disentangled from what practices we pledge ourselves to, and the norms we endorse to govern those practices. This is not a picture that is easily accommodated by social veritism.

Another important difference between my approach and Goldman's is that mine allows for a broader set of epistemic goals, not just pure veritism; accordingly, RN evaluates norms in large part by how well they manage the potential conflicts between those goals. I suspect that this will allow RN to make better sense than Goldman can of the role of only indirectly veritistic norms, such as publication credit (Goldman 1999, 260ff) or the special prestige of originating a novel theory. But that is not a matter I will pursue further here.

In my earlier paper, I postulated the two goals of diachronic reliability and dialectical robustness as governing the norms of justification, which regulate the legitimacy or illegitimacy of various epistemic acts or attitudes. Such norms will tend to be conservative, primarily telling us what we *cannot* do. They will thus need to be counterbalanced by other norms that encourage, demand, incentivize, or otherwise push

researchers into making new discoveries. We will thus need other norms than the norms of justification, norms that will be at least partially governed by a further desideratum, which I will term *discovery richness* (so that I can still call all of them the “DR desiderata”): their propensity to steer us towards new discoveries. In some ways this is similar to what Goldman (1987) has called the “power” of an epistemic practice, as its capacity to give correct answers to questions of interest, as distinct from its tendency to give correct answers about whatever questions it may simply happen to provide them for. But I would rather not build the interestedness into the goal itself.¹⁴ I am trying to retail a form of relativism to those who might not initially find it palatable or plausible, yet a danger of using a pragmatist framework is that it can result in too easy a relativism – a relativism of interests, such that groups with different interests end up with different norms. But then a more traditional epistemologist may object that we no longer have an *epistemic* relativism on our hands, as the norms in question are not really epistemic norms. This is a standard objection against forms of epistemic pragmatism that simply define, say, justified beliefs as those which serve the agent’s goals.¹⁵ My hope is that if

¹⁴ See Schmitt (2000) for arguments against defining our epistemic values in terms of such interests; and also for a detailed discussion of the various ways in which one might choose to define the truth-promoting values of practices. Although a fully-developed RN would need to tackle such questions, I do not believe that those details will bear on my arguments here.

¹⁵ E.g., Stich (1990); see Haack (1994) and Goldman (1999) for objections to such forms of pragmatism. I am also, for these reasons, setting aside in the current discussion arguments that various political interests should be a part of our epistemic norms, as in

the norm-setting goals are themselves impeccably epistemic, such concerns need not arise.

The particular norms I will focus on here, for purposes of setting a stalking horse in the hunt, are the norms of explanatory evaluation. How do we decide which, of a number of candidate theories, are worthy of the application of inference to the explanation?

There is reasonably broad agreement on a number of key components of a good abductive inference, such as “mechanism, precision, scope, simplicity, fertility or fruitfulness, and fit with background belief” (Lipton 2004, 122; he is building there on the likes of Hempel (1966) or Quine and Ullian (1978)). But agreeing on the basic list of good-making features of explanations does not help us very much in actually evaluating our theories. Sometimes one theory may dominate all its rivals on all or nearly all of these criteria. Much more often, though, the hard part is deciding what the right trade-offs are between rivals, in which one may (for example) seem to cover a larger set of phenomena, but in a drastically less simple way than another; or in which one theory wins in terms of ontological parsimony but its chief rival makes tidier predictions. For example, in the standing debate between classicalist and dynamicist forms of cognitive science, the former have the virtue of explaining some phenomena, such as the systematicity and productivity of linguistic behavior, that the latter have deep troubles

Longino (1990). Doing so still leaves open a large project of enumerating and evaluating other potential epistemic goals and their relations to each other, such as the inculcation of good epistemic virtues (Zagzebski 1996), or an increase in understanding (Kvanvig 2003), or the speed and efficiency with which results are provided (Gigerenzer *et al.* 1999), and surely others. But that is a project for another time.

with; but the latter avoid the kind of deep ontological tangles that cognitive representations seem to involve us in (such as the problem of grounding the intentionality of symbols, and indeed what can even *count* as a representation in a system as profoundly distributed as mammalian neurology). There is no reason to think that these schools disagree at the level of explanatory slogans or short, broad-stroke lists of criteria of good explanations. Rather, where they diverge is at the level of the particular explanatory norms that give substance to those slogans and lists. (See also §VI. below.)

Such norms will also determine just what phenomena a set of competing explanations needs to be responsible to. For example, can a theory of syntax legitimately bracket off some speaker judgments as “performance errors,” or do all judgments make equal demands on our linguistic explanations?¹⁶ Is it a flaw of some models of cognition that they do not even attempt to model the precise way that certain cognitive processes unfold in time?¹⁷ Different linguists or psychologists have answered these questions differently, and from these different answers have sprung the sparring camps of formal and functional linguistics, and classical and dynamical cognitive science. I would suggest that these debates are better viewed as at least in part operating more at a normative level than at a purely empirical one – there may be little disagreement between opposing camps as to the extant data, but deep and sometimes vehement disagreement about what sorts of explanations are viewed as comporting best with that data. Part of what I will ultimately be arguing is that relativism may be true of these debates, and view them as

¹⁶ See, e.g., Paradis (2003); Penke and Rosenbach (2004).

¹⁷ See, e.g., van Gelder and Port (1995); Eliasmith (1996); Thelen and Smith (1999); Chemero (2000).

containing at least some irresolvable disagreements. For now, though, I am just trying to locate what sorts of concerns our explanatory norms will need to address.

In addition to concerns about what the proper *explananda* are, we need norms to fix the admissible *explanantia* as well. One need not be a Wittgensteinian to agree that all explanations come to an end somewhere. The aforementioned debate between classicalists and dynamicists in psychology is not just about the relevance of temporal data, but also about the status of appeals to representations in psychological explanations. For a classicalist, it can be enough to explain a piece of an infant's behavior to propose that the infant has a relevant form of representations. For example, if the infant can distinguish displays in terms of the cardinality of the objects presented, then the infant has some capacity to represent number. Dynamicists, as contemporary scientific heirs to Quine, challenge the adequacy of such proposed explanations, considering representations to be creatures of darkness, yielding at best only a variant of "dormitive virtue" pseudo-explanations.¹⁸ (Again, see §VI. below.) One might note a similar disagreement within philosophy, as to whether anything is gained in one's philosophical explanations by declaring that concepts have various "constitutive" or "primitive" entailments, other than the most banally literal cases like RED being a constituent of RED STAR.

¹⁸ I am leaving aside here the interesting philosophy of cognitive science question of whether dynamicists *need* to be dynamicists, or if in fact they just tend to be so. A dynamicist version of the "representations without rules" theory of Horgan and Tienson (1989) might be a legitimate possibility. As things currently stand, though, the former position is rarely found without the latter.

IV. Relativism and Explanatory Norms

Our norms of explanation need, at a minimum, to give us not just a set of explanatory virtues but also some way of evaluating trade-offs between them; a guide to how to determine the phenomena that a set of competitor theories need to account for; and further direction as to what may be appealed to in such accountings. Such needs constitute the work that these norms need to do – yet we still must determine what counts as doing that work *well*. Applying the method of reconstructive neopragmatism, we therefore ask: in service of what goals might these norms be (re)constructed? Diachronic reliability seems an obvious contender: we want to give status to those explanations that are more likely to be true, and to hold on to such explanations once we have acquired them. Dialectical robustness is also in order, as we want to make sure that our discussions about how the various theories measure up can benefit from the sharing of information and the division of intellectual labor. Also, we want to minimize the amount of fruitless and unmanageable disagreement in our debates about the merits of the theories. Agreement need not be pursued for its own sake, but we do at least want to be able to learn from our disagreements, and see them resolved in such a way that all parties can judge that their concerns have been well addressed. Finally, the standard lists of explanatory virtues typically include the likelihood of the explanation's directing us towards new and interesting phenomena in the world ("fruitfulness" being a common denizen of lists of explanatory virtues). And so we will perhaps require our explanatory norms to serve the goal of discovery richness as well. There may well be others, and

these three themselves obviously would sustain a great deal more detailed articulation and development. But they will more than suffice to get my argument for MER off the ground. Also, performing a complete pragmatist evaluation of different norms of explanation would be a massive undertaking. For example, the entire literature on Ockham's Razor and the truth-conduciveness of the simplicity of theories would be relevant to such discussions. Fortunately, our purposes here allow us to focus narrowly on the question of the ways in which MER could fail – or foster! – these goals. We can do so by considering the following hypothetical. Suppose that we allowed the norms of explanation to vary in a manner both demographic and moderately unresolvable; that is, suppose we have a number of different groups, each of which operate under different explanatory norms, and at least some of the disagreements across these groups – let's call them “norm communities”¹⁹ – are not resolvable by the application of their different norms. Under what circumstances and in what ways will such a situation impact the goals against which any proposed explanatory norms must be evaluated?

If there is significant variation in how truth-conducive the various sets of norms are, such that some groups would be licensed to prefer explanations that are particularly less likely to be true than their competitor explanations, then we should expect the desideratum of diachronic reliability to be ill-served by such a form of relativism. We will not be helping ourselves, if we allow ourselves to operate with systems of norms that

¹⁹ In principle, nothing more need be required to constitute a norm community than that there be a group of persons who each individually are operating under the same norms as each other. But as noted earlier, a virtue of RN is that it can make sense of a more fully social and collective type of epistemic normativity.

do not even begin to steer us towards the truth. However, so long as all of the various sets of norms are in the same basic ballpark of truth-conduciveness, such variation will not be of direct concern for this goal. Now, it is at least conceptually possible that there be some system of norms that were perfectly geared towards our getting the truth, and if this possibility obtained, then concerns for diachronic reliability might disallow anything like MER. But such concerns will give us no quick reason to prefer a single set of imperfect norms over a family of equally-imperfect sets (as we are surely more likely to encounter in our lives as limited, human epistemic agents).

Dialectical robustness, however, immediately faces a substantial problem: if you and I are in different norm communities, then it might seem that you and I quite literally have nothing to say to each other. It would be like two teams taking the field, and then discovering that they disagree fundamentally about what even counts as scoring a point – it seems that there would be no way for them even to *begin* to play together. Relativism threatens to rob us of our ability to learn from each other, from both our agreements and our disagreements. Kai Nielsen states a version of this worry, about an inconsistency between relativism and wide reflective equilibrium; he is considering a conceptual version of relativism, but the concern applies here as well: “Considered judgments, principles, theories, descriptions of the world where they belong to different [groups] could not even in principle be compared ... and thus, because we cannot make such comparisons, we could not, even in principle, get our considered judgments into wide reflective equilibrium” (Nielsen 1993, 329). Such a result concerning dialectical robustness is unsurprising, as it is exactly this kind of breakdown of discourse that makes radical forms of relativism so vertiginously threatening.

And such breakdowns threaten the desideratum of discovery-richness as well. What we want is to keep acquiring ever more knowledge across an ever-expanding set of domains. And it will be impossible for individual creatures like us to become masters of more than a small handful of those domains. (In general, it's hard enough for someone to master even one sub-sub-sub-domain!) So the drive to make new discoveries will inexorably lead to greater and greater specialization. But this whole process may be short-circuited if the specialists in different areas are all sealed off epistemically from each other. If you and I are epistemically isolated from each other, then I cannot be credited epistemically with whatever you have learned, nor vice-versa. In particular, if you are in a different norm community, then it may not be possible for me to credit your discoveries as new pieces of *knowledge*, or whatever term of appraisal my community uses to tag propositions that may now be appealed to as evidence. We are often driven to new discoveries in the process of reconciling the received views of our own community with those of another; thus was the current field of behavioral economics born of trying to square economics' previously heroic assumptions about human rationality with some disquieting results originating from psychology. This engine of discovery would be put at risk in a relativistic system. For example, if geologists had not been able to grant status to the findings of paleontologists, then their own discussions would not have benefited from the fossil findings that helped settle the debate over continental drift. (In this way, diachronic reliability runs the risk of being *indirectly* impaired.)²⁰

²⁰ When discovery is driven by non-rational processes, then MER may not raise any difficulties. If all that one needs from another norm community is *inspiration*, then one need not grant the source of that inspiration any particular epistemic standing. (Consider

The reader has no doubt noticed that the previous paragraphs are seasoned heavily with words like “may” or “it can seem”. MER certainly *threatens* damage to the desiderata, but we still can ask the question, must those threats be realized? And, perhaps more usefully, under what circumstances will they *not* be realized? Indeed, under what circumstances might some forms of relativism *promote* these epistemic goals?

As our ever-increasing knowledge of the world collides with the limitations of the human mind, I noted above one epistemically centrifugal force that results: the drive towards specialization. An individual, even one as remarkable as a brilliant and accomplished scientist, can only master so much information in a lifetime – and indeed typically whatever mastery they will deploy must be acquired in merely the initial few decades of that lifetime. To stand on the shoulders of so many stacked-up giants requires a long and arduous ascent. Maybe some of us can still be Renaissance persons of a sort, but if we all were to try to do so, the DR desiderata would suffer. One can bemoan the fact that increasingly the typical inquirer can only make contributions to a particular sub-sub-field, yet one must also recognize just how hard it is to do even that, when that sub-sub-field may have its own handful of specially targeted journals, not to mention the journals dedicated to its sub-sub-sub-fields. (For example, a physician working in the emergency department of a children’s hospital who wants to stay current on her field’s research will need to keep up with not just general medical journals like *The New England Journal of Medicine* and *Journal of the American Medical Association*, and not even just the

the story of Kekulé’s dream of a serpent swallowing its tail, inspiring his postulation of the ring structure for benzene.) But often discovery needs not just inspiration but *evidence* from other realms of inquiry.

specialty journals like *Annals of Emergency Medicine* and *Pediatrics*, but even more targeted ones like *Pediatric Emergency Care* and *Journal of Pediatric Surgery*.) If we are to continue to enlarge our picture of the world on the whole, we must settle for seeing an ever smaller proportion of that picture as individuals.

Continuing with concerns about specialization, we must recognize that specialization leads not just to different individuals mastering different bodies of information, but very likely also to their seeking to operate in accord with different norms as well. The most obvious way in which this is so is that researchers in different specializations will be responsible to different data, and will be able to appeal to different *explananda*. A psychologist may take facts about our developmental biology as given facts that she may use in her explanations without being required to explore them further; obviously, the developmental biologist will not do the same. Furthermore, the norms governing the evaluations of explanations that serve one domain well may not serve another equally well, as physicists may seek out explanations in terms of exceptionless laws that economists would despair to pursue. A certain amount of relativism may be in order, then, as we tool our explanatory norms to the particular domains we are investigating.²¹

So we can see at least one factor in our cognitive situation that might require something like MER to accommodate: the apparently monotonic trend towards specialization that is the resultant vector of our human cognitive limitations on the one hand and our drive ever to push back the frontiers of human knowledge on the other. It is

²¹See also Dupré (1993), Cartwright (1999), and Longino (2002) for more elaborated versions of the argument that different scientific fields may require different forms of investigation and explanation.

not clear, though, that this relativism is really so unpalatably relativistic, for so long as each individual question or domain has its own unique optimal set of norms, then the absolutist can rig up a largely disjunctive and conditionalized, but entirely self-consistent, unified framework of norms. It is not clear that we yet have the unresolvability that we earlier decided should be part of any relativism worthy of being so called. We are only at pluralism, not yet relativism.

So we need multiple sets of norms to cover the multivarious, dappled world. But our human limitations are not just metaphorically spatial, but also literally temporal. We do not know all that we would like at any time, and as we learn more about the world, we also learn more about how it is we should go about learning about the world. At one time, we did not know about the need for double-blind designs, or the importance of operationalizing our claims, or that statistical methods are better used to reject a theory in the face of data than the other way around.²²

Such considerations might motivate a version of relativism, so that the flawed norms of our centuries past can still be seen as normatively appropriate for those periods. For we may want to distinguish between those in the past who were and were not doing as well as they legitimately could have been expected to do. And since we have every reason to expect our intellectual descendants to view us in much the same light, and we do not want our own investigations to count as normatively improper, *amour propre* could further motivate such a relativization. However, a non-relativist could well respond that

²²See, e.g., Gigerenzer (1991a) for an account of this development in the history of statistics, which Gigerenzer then argues allowed psychologists to take statistical inference to be a model for the operations of the mind itself.

all we need is the ability to distinguish *culpable* from *blameless* failure to comply with the correct norms. We perhaps violate fewer norms than those who went before us, but where we do violate them, we are just as wrong as they were. We do not feel an equivalent pressure to relativize truth because many of our beliefs will turn out false, after all. (Interesting questions do get raised here about the connections between complying with the relevant norms of inquiry on the one hand and counting as justified on the other; it is harder to sustain the claim that our predecessors and indeed, most likely, we ourselves are substantially unjustified in our beliefs.)

I would locate a different path for getting to relativism from our limited, corrigible, and ever-improving set of norms of inquiry. First we must recognize that inconsistencies between systems of norms are synchronic as much as diachronic. Some of the more interesting debates going on within various fields today concern normative questions about the what and how of inquiry in those fields. Just looking at the field I am best acquainted with, psychology, one can find those investigating human reasoning disagreeing about the relevant normative standards against which human performance is to be evaluated²³; and infant psychologists disagree sharply as to how readily we should attribute representations to the child, in explaining various of her abilities and behaviors; and animal researchers wrestle about the comparative merits of field work versus lab work²⁴; and perhaps most famously to philosophers today, the relevance of evolutionary approaches to human psychology is a matter of heated, sometimes even vituperative

²³ See, e.g., Gigerenzer (1991); Dawes (1990); Vranas (2000).

²⁴ See, e.g., Allen (2004).

dispute²⁵. These disputes are not settleable from the armchair, but must be addressed through (and as part of) the progress of the relevant science.

It follows that the norms of inquiry must aim to maximize the DR desiderata not just with regard to factually descriptive matters, but also about the norms themselves. And just as a free competition of hypotheses has proved to be a valuable framework for theory selection, so too will competition between norm communities serve our interests in developing and fine-tuning those norms. And just as we license the operation of divergent theories as normatively appropriate, so too will we need to license the operation of divergent systems of norms as appropriate as well. There are close analogies here with Philip Kitcher's (1993) contention that the diversity of different evidentiary temperaments was valuable to the progress of science, such that some scientists were quicker to adopt new theories than others, leaving their more conservative colleagues to preserve key insights of the older models. Similarly, Miriam Solomon has demonstrated in her (2001) the importance of dissent for preserving the empirical resources of our various theories, so that even the limited successes of disfavored theories are not lost to us, for as long as our most favored theory cannot achieve those particular successes. So, when we cannot determine which of a number of competing systems of norms is definitively to be preferred, we do well to keep them all in play. Just as our inquiries benefit from a diversity of temperaments and theories, so too do they benefit from a diversity of norms. This is not diversity for its own sake, but diversity in service of our epistemic goals. At any given time, our norms encourage the combative flourishing of a set of theories that are inconsistent with each other – and which we therefore think cannot

²⁵ See, e.g., Buller (2005); Machery and Barrett (2006).

all be true – because that facilitates our future determination of which among them is right. Because this proliferation of competitors is part and parcel of what makes our norms work as well as they do, it would be inappropriate to view all but one of the competitors as merely blamelessly wrong. Where our norms promote diversity, they surely ought also permit it. We commend, and do not merely excuse, a plethora of competing stories of our world’s truth, and my argument here is that we do well to adopt the exact same attitude towards our inquiries’ norms. This is no mere tolerance of alternatives, but a justification of relativism.²⁶

V. Relativism without Tears?

This is the situation before us: on the one hand, relativism about epistemic norms seems poised to issue a serious threat against the DR desiderata, primarily by weakening the possibilities of fruitful transactions across the borders of the different norm communities. Yet on the other hand, some form of relativism is supported by the way

²⁶ There are other benefits that can accrue within a relativist system of norms, if the dangers to our desiderata can be defanged. Although relativism threatens dialectical robustness at the *inter*-group level, it may well promote it on an *intra*-group basis, through group solidarity, and a sense of identity and community. Moreover, our desire as a group for ‘us’ to beat ‘them’ will motivate us to articulate as cogently and firmly as possible our challenges to ‘their’ way of doing things; see below. (My thanks to Fred Schmitt for this observation.)

our human limitations drive us towards specialization in our inquiries, and towards the recognition that the only way we can improve on our current norms is through allowing divergent sets of norms to battle it out on the fields of science. (These considerations are entwined, since the creation of new specialized fields opens up questions about the right way to go about investigating within those fields.) The challenge for the epistemologist is to devise a form of MER that will avoid the pitfalls while retaining the promise.

The perils of relativism arise, most fundamentally, from *the nontransferability of status*. Our ability to learn from and correct and be corrected by each other depends on our being able to grant status to others' claims as having a rational bearing on our own. In a relativistic system, when we are unable to recognize the evidence arising in other norm communities as such, we cannot use that evidence as a source of credit or discredit for the evaluation of our own theories or those of our rivals within the same norm community. A successful form of MER will thus require some ways of communicating status across norm borders, even when full recognition of status is only possible within a given community. The transmission of status within a community is part and parcel of what the norms governing that community are for. We must seek some special inter-community channels for the conveyance of status, operating outside of the standard intra-community ones, yet still recognizable as normatively appropriate within any community. Such cross-community channels of status exportation, if they existed, would defang relativism without denaturing it.

And fortunately, a number of such channels are available. Let me indicate some of them here, with 'H' indicating the *home* norm community whose perspective we are imagining as our own, and 'A' the alien norm community whose deliverances the good

citizens of H are considering importing. How might H come to make use of A's claimed evidence, even though A's norms diverge from H's?

Perhaps the easiest way for H to accept what A puts forward is to treat A as a testimonial black box. Within a given community, testimony is most often given and received in a robustly normative way, in which we do not merely treat others as "truth thermometers" but as agents offering the assurance of their own judgment and responsible reasoning.²⁷ But the conditions for such testimony do not, *ex hypothesi*, obtain between H and A. Nonetheless, H can retreat to the Lockean approach of evaluating A's track-record, if that record is at all evaluable from H's perspective. (A must have *some* results pertinent to H, for the question of import to have arisen in the first place.) Even while the full richness of A's reasoning will be lost in the process, their conclusions may still thereby be a potential source of valuable friction for H's deliberations.

It may also be possible for the H-researchers to recognize that the A-researchers are using *some* norms that they themselves endorse. This may allow a partial translation of status, depending on how molecular the application of those shared norms may be. For example – perhaps the easiest example – both communities may concur as to how new data is to be generated, even though they diverge as to the best ways to explain such data. Under such conditions, H can adopt A's experimental findings wholesale, even if it then uses them towards different explanatory ends. But other shared norms can play a role in enabling inter-community transactions. For example, Fagan (2007) closely documents an example from the recent history of biology in which a shared goal norm enabled fruitful

²⁷ See, e.g., Owens (2006).

competition between different immunological research groups. A particularly interesting variation on such scenarios of partial translation will develop when a researcher in A takes her findings to provide the raw material for an attack on one of H's norms.

Remember that part of the point of a relativistic framework here is to allow for the competition between sets of norms; we will therefore wish of H that it not just brush aside A's challenge, but consider whether it has the resources to re-interpret and then meet that challenge within its own system. (I will take up this question again at the end of the paper.)

It is important to note that norm communities do not form anything like a partition of inquirers, and most – maybe all – inquirers belong to a number of different communities. For example, a given person may be a philosopher, an epistemologist, a naturalized epistemologist, an experimentalist; and so on. This provides us with a possibility for interactions that are somewhat more nuanced than mere black-boxing. For if A and H are both sub-communities of a broader community B, then H may be able to recognize the normative force of some of A's moves inasmuch as they comply with the norms of B. An extended example in §VI. below is substantially of this variety.

A third, separate community C may provide a channel between A and H even if C is not a super-set community of them, if A and H both have some sort of channels into and out of C. For example, subatomic physics became relevant to linguistics only through both of those fields' connections to neuroanatomy. That happens to be a case in which the researchers in H can still recognize easily that the influence came from A, since words like “positron” still reside in the names of the technology. But we might expect that there are many cases in which the results from A are so fully absorbed by C, that

other researchers will not necessarily see A's fingerprints on them when they arrive in H.²⁸ Providing a very different sort of example, Michael Friedman has recently argued in his *Dynamics of Reason* that philosophy can play a key role as a mediating field during a period of paradigm transition in the sciences, which in our terms could be seen as a transaction from one norm community to another. Einstein, for example, was better able to advance and articulate his theories of relativity given the background of Poincaré and Helmholtz's philosophical discussions of geometry in the previous century.

Yet another important channeling norm community is the very broad community of our ordinary, non-specialist epistemic lives. Reading Chomsky's famous broadside against behaviorism in his review of Skinner's *Verbal Behavior*, one finds almost none of the argument operating within the particular framework of the kind of linguistics that Chomsky was looking to launch. Rather, his arguments turn largely on pointing out the apparent predictions of Skinner's system, and showing how poorly they fare against ordinary observation. For example:

...although the evidence for the contribution of each variable to response strength is based on observation of frequencies alone, it turns out that "we base the notion of strength upon several kinds of evidence"... [such as] emission of the response (particularly in unusual circumstances), energy level (stress), pitch level, speed

²⁸ I would note that this also raises at least the possibility of odd little circularities, since it would be possible for a result originally from A to be imported incognito into H, and then through some other channel wend its way get back around to A, and ultimately be mistaken there as independent confirmation of what, in fact, was just itself.

and delay of emission, size of letters etc. in writing, immediate repetition, and -- a final factor, relevant but misleading -- over-all frequency.

Of course, Skinner recognizes that these measures do not co-vary, because (among other reasons) pitch, stress, quantity, and reduplication may have internal linguistic functions. However, he does not hold these conflicts to be very important, since the proposed factors indicative of strength are “fully understood by everyone” in the culture... “For example, if we are shown a prized work of art and exclaim *Beautiful!*, the speed and energy of the response will not be lost on the owner.” It does not appear totally obvious that in this case the way to impress the owner is to shriek *Beautiful* in a loud, high-pitched voice, repeatedly, and with no delay (high response strength). It may be equally effective to look at the picture silently (long delay) and then to murmur *Beautiful* in a soft, low-pitched voice (by definition, very low response strength). (1959, 34 -35)

One needs no special grounding in either linguistics or psychology to see the problem that Chomsky’s last observation raises for Skinner’s apparent prediction.²⁹

Because MER can only succeed if there are substantial routes of status import and export, inter- and multi-disciplinary communities will take on a special importance within this framework, serving as a persistent and less ad-hoc version of what Peter Galison has

²⁹ I am putting aside here the question that has sometimes been raised by friends of behaviorism that Chomsky has misunderstood Skinner; I think that that is a mistaken charge, but anyway the important thing is that this *kind* of argumentative move is available.

termed a “trading zone” – a discourse created between groups operating within different operational vocabularies when they need to collaborate, such as the physicists and engineers involved in the development of particle detection technologies (Galison 1997). At their best, interdisciplinary communities can serve as a tremendous conduit across the boundaries of distinct norm communities, in large part because the people working within the interdisciplinary zone will be making a particular effort to see how to make such transactions work. Another role for philosophy, then, is as a contributor to such communities – sometimes by standing above them, in the manner suggested by Friedman, but also sometimes by participating in them. Philosophy’s analytic and conceptual tools will be especially relevant to the special task of aiding in the functioning of such interdisciplinary matrices.

There are thus a large number of ways in which two norm communities may have productive exchanges of the status of their separate results and arguments, even in the absence of the kind of singular overarching normative framework that the absolutist requires. And I have tried to suggest that these are not merely potential avenues of exchange; there are real examples of them in action that are easy to uncover, and I suspect many, many more where these came from. What we lack, however, is any sort of overarching master norms that will guarantee for us that such an exchange will always be possible any time a community A has a result it wishes to export to a community H, either to rule the exportation in as appropriately status-bearing or rule it as recognizably out-of-bounds. If we had such master norms – that is, if all disagreements across groups had some set of norms to appeal to in order to settle those disagreements – then we would not have a form of relativism at all. Rather, what we have are a set of status-

transfer channels that is motley, local, and frequently (though not always) improvised. The available channels may be merely listable without being systematizable. We can thus expect that many unresolvabilities will persist, without a good candidate way for H and A to come to any sort of agreement as to what A's offering really can do for the folks in H. We thus secure the possibility, indeed the plausibility, of a form of epistemic relativism, unresolvabilities and all, that is nonetheless moderate, and thereby promotes a set of our epistemic goals even while steering clear of the epistemic disaster that a more radical relativism would portend.

VI. A case study: the debate over infant arithmetic

Let me now illustrate some of these themes with a somewhat more extended example, in the recent and fascinating debate over infants' arithmetic capacities. (I should warn the reader that what follows is, in addition to of course being a very abbreviated and simplified account, also very much a *philosopher's* narrative of this research, and the issues of philosophical interest— the competition between different modes of explanation — are not as front-and-center in the psychological research itself as I will be depicting it here.)

One of the dominant research methods in developmental psychology of recent years is the looking-time paradigm. It is not easy to plumb the depths or shallows of infant cognition, since they differ even from toddlers in that one cannot give them linguistic directions, nor set them even slightly complicated tasks, nor simply ask them what they are thinking. But one can, it is thought, explore what aspects of infants' surroundings make at least some sort of difference to their minds, by seeing what sorts of stimuli they

will discriminate between. Infants' behavioral options are pretty limited, but what they can do (at least, if you keep their heads up) is direct their eyeballs. When they are confronted with a particular stimulus, they typically attend to its novelty, but after a little while they will stop paying it any particular attention. If one then presents the infant with another stimulus, one can measure to what extent the infant 'dishabituates', spending more time looking at it. No increase in looking time is frequently taken to mean that the infant is not discriminating the latter, 'test' stimulus from the earlier 'habituation' stimulus. An increase in looking time, however, indicates at a minimum that the infant discriminates between the two stimuli. Moreover, it is sometimes interpreted as meaning that the infant had been 'expecting' something different, and is perhaps even 'surprised' by the new stimulus. As we shall see, just how best to interpret these results is a matter of significant debate.

One classic version of such a study aimed to demonstrate that very young children have a grasp of elementary arithmetic. First, the infant subjects are shown one object on a stage, or two objects on a stage, to determine a baseline of interest in such displays. The subjects are then shown one of two trial conditions. In each version, the subjects see one object on a stage, and then a screen is raised blocking their view of what is on the stage. They then see another object placed behind the screen, and the screen is lowered. In one version, only one object is revealed, while in the other condition, two objects are revealed. Interestingly, the infants look much longer at the one-object version than at the two-object version. So the one-object version seems to strike the infants as something novel, while the two-object version is just more of the same. Psychologists of a classicalist (and typically nativist) bent have interpreted these results as indicating that

even very young infants have some rudimentary arithmetic competence: the one-object version only counts as novel if it can be viewed as a mathematical novelty. (“One plus one should equal two, but there was only one!”)

But psychologists of a more dynamicist (and empiricist) orientation have contested this interpretation. Mix and Clearfield (1999) argued that changes in numerosity are confounded with other changes, such as the total contour length of objects in the visual field; still other continuous aspects of the stimuli might also be driving infants’ reactions, such as total area (Clearfield and Mix 2001). Initial experiments which disentangled these confounds found that infants might be responding not so much to number as to amount. Classicalists have followed up with more elaborate experiments designed to reveal infant arithmetic competences, for example by exploring the auditory as well as visual modality (e.g., Kobayashi *et al.* 2005; vanMarle and Wynn 2003).

Unsurprisingly, neither camp has been persuaded by the other’s arguments – at least, not fully. One can see the variety of reactions on display in a series of peer commentaries on a paper in a 2002 volume of *Developmental Science*. Leslie Cohen and Kathryn Marks argue in their target article that Wynn’s original effects can be explained away in terms of infants’ preferences for larger stimuli and for more familiar stimuli. The dynamics-friendly Mix concurs classical interpretations of infant behavior are uncertain, “optimistic given the data, perhaps even recklessly so” (Mix 2002, 206). But where Mix sees irrational epistemic exuberance on the part of classicists, the more classically-oriented psychologist Susan Carey sees “abundant converging evidence” (Carey 2002, 203) for the operation of numerical representations. Furthermore, rival accounts like that offered by Cohen and Marks can only apply Ockham’s Razor against the classical ones if

they have on hand a fuller account of such variables as familiarity, and indeed the *representations* involved in infants' perceptions of familiarity (204).

We can understand these starkly divergent evaluations of the question as to who has the basic burden of proof here as evidence that there are distinct norm communities in play. As far as the classicalists are concerned, arguments like those of Clearfield and Mix or Cohen and Marks can simply be taken as interesting critiques of specific experiments. Yet dynamicists will construe these arguments as demonstrating evidential lacunae at the very root of classicalist approaches to infant cognition; indeed, Mix accuses her opponents of building on “a foundation of shifting sand” (206). For these different researchers are operating with a different set of explanatory norms. In their stock of candidate *explanantia*, classicalists see no significant cost in adverting to representations. While any particular representational story must be experimentally motivated, and representational stories must compete with non-representational ones, such competitions are on a basically even playing field. Yet dynamicists assign a much more severe cost to adverting to representations. Indeed, I do not think I am exaggerating to report their explanatory norms as committing them to the view that, so long as *any* non-representational explanation is available for some phenomenon is available, it is to be preferred over all representation-invoking explanations.³⁰ Each community thus sees the other as laboring under the burden of ad-hocery. To the classicalist, the dynamicists have a motley of perceptual cues from which they can construct rival explanatory hypotheses, but the set is so under-constrained and unsystematized that they do not

³⁰ See, e.g., Brooks (1991) and Thelen *et al.* (2001) for a fairly straightforward embodiment of this view on the part of some representative scientists.

present any privileged class of such explanations. To the dynamicists, the classicalists are reaching prematurely for the most prodigiously expensive explanatory tool in the box, while ignoring the wide and still under-explored class of purely perceptual-behavioral resources for constructing an explanation.

An anti-relativist could accommodate this normative divergence easily, if we could see the two communities as simply differing in the evidence they possess: if the classicalists had access to information that the dynamicists lacked, for example, where that information provided good evidence for the existence of cognitive representations. Also, an anti-relativist could accommodate this normative divergence if the two communities simply disagreed as to which phenomena they were trying to explain: if the dynamicists were only trying to account for certain low-level behavioral patterns, for example. But neither of these moves will succeed here. No such informational asymmetry obtains, and all the scientists in question have access to all the same data, and moreover they take their theories to be ultimately responsible to the same results as well.

However, if these two communities were thus doing nothing more than completely talking past each other, with an absolutely irresolvable difference in norms, then we would find ourselves with just the sort of dialectical breakdown that we wish to avoid. Fortunately, that is not the case. There are enough shared standards of experimental psychology that Mix and her colleagues take Wynn and her colleague's results seriously, and vice-versa. So *some* importation is going on at the level of data, at least. We want to see the possibility of an improvement in our explanatory norms, though, and the mere exchange of data does not seem to hold out much promise for helping with that project. Researchers like Mix mean to be challenging one of the fundamental explanatory tenets

of researchers like Wynn; can we see a way in which such challenges may move us forward, even if they are not unproblematically resolvable in a unified normative system?

In my lingo here: they have imported the dynamicists' arguments, but only as offering a very local, circumscribed challenge of a sort unproblematically met by doing further experiments. Doing so has proved a potent engine for the generation of novel results. Moreover, such exchanges have enabled the further refinement of the norms of each community, as they accommodate the experimental results produced by each side, even if they do not agree as to the larger implications of those results on the whole. It is possible that such parallel, interactive progress may continue indefinitely. But those interactions also make it possible that at some point, enough problems will be raised *inside* one or the other community, that it may fold its tent, its load-bearing normative principles fatally undermined from within. Or, much more likely, we may find one or both communities restricting its explanatory ambit so that it is no longer in real competition with the other. One can view Feigenson *et al.* (2002), Xu (2003), and Xu *et al.* (2005) as doing such, by positing distinct systems for large-number cognition in infants (which is still thoroughly numerical) and small-number cognition (which may be driven entirely by continuous factors). We can see the scientific endeavor of developmental psychology thus benefiting from competition between rival norm communities. The right kind of competition, that is – one that has significant opportunities for productive import and export.

VII. Conclusion: The Prospects for the Appeal to Discipline

We therefore find that under certain conditions, a moderate epistemic relativism can be not only harmless, but indeed requisite for meeting our epistemic goals. These conditions, for which there is considerable evidence, include (i) limitations on cognizers and their institutions that make specialization of and competition between sets of norms desirable, and (ii) the existence of a sufficient set of routes of the export of epistemic status across different communities. Within the methodological framework of reconstructive neopragmatism, at least, MER is not only defensible but indeed desirable. (Proponents of other methods will have to determine for themselves how best to accommodate this within their frameworks. For this result is one that any epistemology should take *some* sort of accounting of – that MER, under such circumstances, would serve our epistemic goals well – and it counts against any method that cannot make sense of it.)

Let me conclude by considering an application of this result to a question that I have been curious about for some time. It is not uncommon for partisans in various areas to deploy what might be called an *appeal to discipline*. In its positive form, one defends a given move in an argument by stating something like, “in field X, this is just how we do things”; more negatively, a putative result that would be trouble for one’s argument may be dismissed with a wave of a hand and an, “oh, but that’s not *really* field X, that’s field Y.” I have heard philosophers, linguistics, economists, and psychologists use this tactic; the psychologists perhaps most interestingly, in that I have heard it applied to inconsistent claims, at different times by different psychologists! The question, then, is what sort of account can be made of the cogency (if any) of such appeals?

If we really have here a form of relativism – and hence unresolvability – then at least *sometimes* such appeals must be appropriate. When our practitioners from H and A have accordingly different norms, and legitimately so, then sometimes the one must be able to demur when pressed by the other. However, this relativism will only serve our desiderata when it also promotes the competition between different systems of norms. And thus *sometimes* such demurrals must be disallowed, and the alien critic permitted to press her point home. The communities can be in general normative isolation from each other only so long as they are not universally and exceptionlessly so. Our epistemic relativism must be, in my terms, moderate. Which arguments can be deflected, and when not, will depend largely on whether or not the argument is offered by means of one of the channels of importation discussed above. When the researcher from A is arguing only from within her own normative framework, but without any means of exporting the argument to H, then no matter how well-constructed her argument is in A's terms, the good folks of H need pay it no heed. But on occasion there will be such a channel. A's argument may be cogent by the standards of a community that subsumes both A and H, for example. In such circumstances, the community of H cannot merely shrug off A's concerns, *even if those concerns threaten the very norms of H itself*.³¹ We can expect that H's own researchers will find it difficult to offer fundamental challenges to their own norms, and so the DR desiderata will only be well-served when legitimate challenges from outside can be acknowledged as such.

³¹ We can see how far this framework is from Carnap, then, in that it allows for *rational* challenges to a system that arise outside of it. This fact also indicates how un-Kuhnian this approach is as well (in addition to not being framed at all in term of semantics).

Now, I have been speaking in the maritime economic metaphor of “channels” and such, but an important disanalogy arises here in that when one attempts to ship a boatload of goods, at the journey’s end either the cargo arrives or it doesn’t, with typically no epistemic difficulties in determining whether the sortie was basically a success or failure. In contrast, after one has attempted an inter-community epistemic transaction, one may simply be unable to tell just how much has been successfully conveyed across the normative frontier. And one may moreover not be able to make such a determination in advance of launching it, either. After all, the success of such an exportation depends not just on one’s own norms (which one has presumably mastered) but on the norms of the other community (which one presumably has not). So we may need to err on the side of letting exporters make attempts that might, from the standpoint of the intended importer, stray from the mark; and, while the researchers in any given community will need to keep an honest look out for legitimate incoming imports, they will also need to fend off the illegitimate ones with appropriate appeals to their own discipline. But we must not hold it against exporters if they are not satisfied with such appeals, either, since it is hard for them to discern their failure from their own perspective – except, that is, when the target community might itself be able to counter-export an argument to the effect that the originally attempted import was ill-founded! What all this means, unfortunately, is that there may be no more general norms to distinguish fair from unfair applications of the appeal to discipline, and no expectation that all parties to a debate over such an appeal will agree as to what are or are not reasonable moves in such a debate. (This is a special case of the kind of lack of systematic master norms discussed at the end of the §V.) Sometimes there will be other locally applicable norms that allow such adjudication, as in

some cases where there is a shared norm community. But there is no over-arching system of norms that we can apply to make all such decisions. Such master norms would have to encompass all norm communities, and as such they would radically outstrip the capacity for any human to understand or act upon them.

Moderate epistemic relativism purchases for us for both specialization and the ongoing development of our norms; yet at the cost of any unitary normative standpoint from which all inter-community disagreements can be evaluated. For creatures like us, a degree of multiplicity is the price of progress. And RN can help us discover the circumstances under which that price is well paid.

References

- Allen, Colin. 2004. "Is Anyone a Cognitive Ethologist?" *Biology & Philosophy* 19: 589-607.
- Boghossian, Paul. 2006. *Fear of Knowledge*. Oxford: Oxford University Press.
- Brooks, Rodney. 1991. "Intelligence Without Representation." *Artificial Intelligence*, 47: 139 – 159.
- Buller, David. 2005. *Adapting Minds: Evolutionary Psychology and the Persistent Quest for Human Nature*. Cambridge, MA: MIT Press.
- Cartwright, Nancy. 1999. *The Dappled World*. Cambridge: Cambridge University Press.

- Chemero, Anthony. 2000. "Anti-representationalism and the Dynamical Stance." *Philosophy of Science* 67: 625 – 647.
- Chomsky, Noam. 1959. "Review of *Verbal Behavior* by B. F. Skinner." *Language* 35: 26-58.
- Cohen, Leslie and Kathryn Marks. 2002. "How Infants Process Addition and Subtraction Events." *Developmental Science* 5: 186 – 201.
- Craig, Edward. 1990. *Knowledge and the State of Nature*. Oxford: Oxford University Press.
- Dawes, Robyn. 1990. "The Potential Nonfalsity of the False Consensus Effect." Pp. 179-199 in R. M. Hogarth, ed., *Insights into decision making*. University of Chicago Press.
- Fagan, Melinda. 2007. "The Search for the Hematopoietic Stem Cell: Social Interaction and Epistemic Success in Immunology." *Studies in History and Philosophy of Biology and Biomedical Sciences* 38: 217 – 237.
- Feigenson, Lisa, Susan Carey, and Elizabeth Spelke. 2002. "Infants' Discrimination of Number vs. Continuous Extent." *Cognitive Psychology* 44: 33 – 66.
- Friedman, Michael. 2001. *Dynamics of Reason*. Stanford: CSLI Publications.
- Galison, Peter. 1997. *Image & logic: A material culture of microphysics*. Chicago: The University of Chicago Press.
- Gigerenzer, Gerd. 1991a. "From Tools to Theories: A Heuristic for Discovery in Cognitive Psychology." *Psychological Review* 98: 254 – 267.
- 1991b. "How to Make Cognitive Illusions Disappear: Beyond "Heuristics and Biases." *European Review of Social Psychology* 2: 83–115.

- Gigerenzer, Gerd, Peter Todd, and the ABC Group. 1999. *Simple Heuristics That Make Us Smart*. New York, Oxford University Press.
- Goldman, Alvin. 1987. "Foundations of Social Epistemics." *Synthese* 73: 109 – 44.
- 1999. *Knowledge in a Social World*. Oxford: OUP.
- Haack, Susan. 1994. *Evidence and Inquiry*. Oxford: Blackwell Publishing.
- Hales, Steven. 1997. "A Consistent Relativism." *Mind* 106: 33 – 52.
- Hempel, Carl. 1966. *The Philosophy of Natural Science*. Englewood Cliffs, N.J.: Prentice-Hall.
- Horgan, Terrence and John Tienson. 1989. "Representations without Rules." *Philosophical Topics* 17: 27 – 43.
- Kitcher, Philip. 1993. *The Advancement of Science*. Oxford: Oxford University Press.
- Kobayashi, Tessei, Kazuo Hiraki, Ryoko Mugitani, and Toshikazu Hasegawa. 2004. "Baby Arithmetic: One Object Plus One Tone." *Cognition* 91: B23 – B34.
- Kvanvig, Jonathan. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
- Laudan, Larry. 1987. "Progress or rationality? The Prospects for Normative Naturalism." *American Philosophical Quarterly* 24: 19 – 31.
- 1990. "Aim-less Epistemology." *Studies in the History and Philosophy of Science*, 21: 315 – 322.
- Lipton, Peter. 2004. *Inference to the Best Explanation*, 2nd ed. London: Routledge.
- Longino, Helen. 1990. *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*. Princeton, N.J.: Princeton University Press.
- 2002. *The Fate of Knowledge*. Princeton, NJ: Princeton University Press.

- Machery, Edouard and Clark Barrett. 2006. "Debunking *Adapting Minds*." *Philosophy of Science* 73: 232–246.
- Neta, Ram. 2006. "Epistemology Factualized: New Contractarian Foundations for Epistemology." *Synthese* 150: 247 – 280.
- Nielsen, Kai. 1993. "Relativism and Wide Reflective Equilibrium." *Monist* 76: 316-332.
- Owens, David. 2006. "Testimony and Assertion." *Philosophical Studies* 130: 105 – 129.
- Paradis, Cardita. 2003. "Is the Notion of Linguistic Competence Relevant in Cognitive Linguistics?" *Annual Review of Cognitive Linguistics* 1: 207-231.
- Penke, Martina and Rosenbach, Annette. 2004. "What Counts as Evidence in Linguistics?" *Studies in Language* 28: 480 – 526.
- Quine, W. V. O. and Joseph Ullian. 1978. *The Web of Belief*, 2nd ed. New York: Random House.
- Schmitt, Frederick. 2000. "Veritistic Value." *Social Epistemology* 14: 259 – 280.
- Solomon, Miriam. 2001. *Social Empiricism*. Cambridge, MA: MIT Press.
- Stich, Stephen. 1990. *The Fragmentation of Reason*. Cambridge, MA: MIT Press.
- vanMarle, Kristy and Karen Wynn. 2003. "Tone Discrimination by 6-month-old Infants." Poster presented at the annual meeting of the Society for Research in Child Development, Tampa, Florida.
- Vranas, Peter. 2000. "Gigerenzer's Normative Critique of Kahneman and Tversky." *Cognition* 76: 179 – 193.
- Weinberg, Jonathan. 2006. "What's Epistemology For? The case for Neopragmatism in Normative Metaepistemology." Pp. 26-47 in S. Hetherington, ed. *Epistemology Futures*. Oxford: Oxford University Press.

Xu, Fei. 2003. "Numerosity Discrimination in Infants: Evidence for Two Systems of Representation." *Cognition* 89: B15 – B25.

Xu, Fei, Elizabeth Spelke, and Sydney Goddard. 2005. "Number Sense in Human Infants." *Developmental Science* 8: 88 – 101.

Zagzebski, Linda. 1996. *Virtues of the Mind*. Cambridge: Cambridge University Press.

Jonathan M. Weinberg received his Ph.D. from Rutgers in 2002, and is currently an assistant professor in the Department of Philosophy and the Program in Cognitive Science at Indiana University, Bloomington. He works primarily in experimental philosophy and in various intersections of the philosophy of psychology, epistemology, and aesthetics, with articles in such journals as *Nous*, *Philosophy of Science*, *Philosophy and Phenomenological Research*, and *The Journal of Aesthetics and Art Criticism*.