Epistemic Virtues and Cognitive Dispositions

DAVID HENDERSON AND TERRY HORGAN

I. Orientation:

The preponderance of the discussions in this volume make a case for positing dispositions, either by attention to the needs of some contemporary discipline, or by way of reflection on some historical thinker's insight, or by unfiltered metaphysical argument. The present piece does not argue for positing dispositions, and rather flatly presupposes that there are dispositions. Our concern is with an epistemological point–one concerning the range of dispositions (or cognitive processes understood dispositionally) that are significant for epistemology.

Most contemporary epistemologists are committed to the epistemic significance of cognitive dispositions. They may write instead of beliefs being the result of appropriate processes–but the processes in question are understood functionally, as dispositions to certain kinds of transitions. They are concerned with what are appropriate processes–functionally understood–by which to fix beliefs, and are thereby committed to dispositions. However, they tend to be concerned with a narrowly circumscribed set of dispositions–too narrowly circumscribed, we argue.

To appreciate this very standard epistemological commitment to dispositions, reflect on the distinction between an agent's *being objectively justified* in a belief and that agent's merely *having justification* for that belief. For an agent to be objectively justified in a belief requires that the processes by which that belief was formed be suitably responsive to the reasons (or information) possessed–it requires that, prompted by possessed information, the belief in question be brought about "in the right way." One can then *have justification* for a belief (having adequate reasons or supporting information) and yet *not thereby be justified* in holding that belief. For example, it is possible that an agent *have* adequate justification (possess information with a content significantly supportive of the content of the agent's belief) and yet would have held the belief even were that information not to have been possessed. To be justified, a belief must be dependent on the contentfully supporting information that the agent possessed–so that, were the agent not to have possessed that infor-

mation, the agent would not have formed or retained the belief in question. Yet, even this does not quite ensure that the agent is objectively epistemically justified in the belief. For example, an agent might have significant justification (supporting information) for a belief that *P*, and this information may serve to raise the question of whether of not *P*. This might then trigger wishful thinking yielding the belief that *P*–wishful thinking that would (once the question was somehow raised) yield the belief that *P* even were the agent not to possess the justification in question. The agent's belief is then counterfactually dependent on this "trigger," but is not thereby objectively justified. Being objectively justified requires a richer array of counterfactual dependencies. The agent's cognitive processes must exhibit patterns of counterfactual dependencies to ranges of relevant information. They must be describable in terms of "invariant generalizations", so that the resultant belief is caused in the right way – in a way that would be sensitive to various information that the agent has or might have.[1] In sum, being justified in some belief turns not just on the information possessed, but also on the cognitive processes in play. The cognitive processes in question are understood largely in terms of stable dispositions to transitions between representations or informational states. Thus, drawing the important distinction between being justified versus merely having justification leads epistemologists to focus on stable cognitive dispositions and to their manifestation in episodes of belief formation and maintenance.

Now, our central concern in this paper is not to examine the metaphysical status of cognitive dispositions, but rather to make a point about the range of dispositions that are epistemically important. Our suggestion is that much epistemology, particularly since the modern period, has treated only a narrowly restricted range of cognitive dispositions as epistemically relevant to being objectively epistemically justified. Epistemologists have been fixated on what we will call *classically inferential processes* (or dispositions to classical inference). We argue that a wider range of processes and dispositions is epistemically crucial.

The core idea of a *virtue* is that of a stable, good or excellent, disposition. This much has not changed all that much since Aristotle. Intellectual virtues were stable dispositions to reason well, when input and time were afforded the agent. *Epistemic virtues* are stable cognitive dispositions that are good (even excellent) from an epistemic point of view. *Virtue epistemology* is an approach to epistemology in which epistemic virtues are accorded an important and central place.[2] This is correct, so far as it goes. But, as it stands, it does not

1 Woodward's (2003) discussion of causes and invariance provides a useful framework for approaching the question of what it is to be caused in the right way.

2 This somewhat generic understanding of epistemic virtues and virtue epistemology aligns with Greco's (2000, chapter 7-8; 2002a) understanding. It is largely in keeping with Sosa's (1991a,

capture what is distinctive about virtue epistemology. For the traditional epistemology for which virtue epistemology is to provide an alternative is itself committed to the epistemic importance of dispositions. We suggest that virtue epistemology provides an important alternative insofar as it is open to the epistemic significance of a wider range of dispositions than has been much epistemology since the modern period. (This contrast will be developed in the next section.)

Insofar as virtues are stable dispositions systematically contributing to epistemically desirable or productive results, and insofar as agents value true systems of belief, one does and should care about being virtuous. Evaluative concepts such as *objective epistemic justification* have coalesced around concerns arising within the individual and joint project of producing true systems of belief. They were developed because, both in one's ongoing self-regulation of one's cognitive life, and in people's joint regulation of their joint or community's epistemic life, one must evaluate beliefs and the processes by which they arose and were maintained. Thus, a core idea informing virtue epistemology is that fruitful and revealing epistemic norms, standard, or models will commonly make reference to stable cognitive dispositions of agents. The idea is that such stable-disposition-featuring standards (a) reflect central aspects of our epistemically central evaluative concepts, and (b) describe (and facilitate becoming) the sort of epistemic agent who, individually or in community, tends to have satisfactory or optimum success at fulfilling that central epistemic goal of possessing systems of true beliefs.

What is optimal or satisfactory must be understood as relative to the kinds of creatures whose belief fixation is being evaluated. Typically, epistemologists have had in view adult humans with normal ranges of cognitive capacities and possibilities. (Such, after all, are the agents who developed the evaluative concepts in question, and who regulate their individual and joint project thereby.) What dispositions do human agent's tend to develop? What plasticity is found here? What stable dispositions can they develop with training? The virtue epistemologist seeks to understand which constellation of stable dispositions among those that humans can with training come to possess, would have the optimum or at least satisfactory tendency to produce and maintain systems of true belief. They insist that objectively justified beliefs are those that are produced by such constellations of stable dispositions–such virtues. Now, in determining what stable dispositions humans can come to possess, with what training, epistemologists must draw on empirical results. It is thus fitting that most virtue epistemologists have been naturalized episte-

1991b) influential epistemology. More restricted understandings of the epistemic virtues (such as Zagzebski's, 1996) are to be found, but we do not believe they provide the best understanding–although they certainly can be instructive at points. The central points made here would seem to carry over to several of the more restrictive virtue epistemologies.

mologists. (Much of what we say in what follows will be informed by contemporary cognitive science.)

A rather a diverse set of dispositions would qualify as virtues, given the core notion just set out. Very general traits of cognitive character having to do with belief fixation are sometimes mentioned as epistemic virtues. Open-mindedness and intellectual courage, for example, are sometimes mentioned–by Zagzebski (1996), for example. Also commonly mentioned are dispositions to employ general cognitive processes characterized in terms of classes of inference: the disposition to reason in ways that are deductively sound, for example, or the sensitivity of inductive strength of evidence might be cited as epistemic virtues. These *general inferential competences* are yet very general dispositions.

Stable dispositions of a more fine grained sort contribute to general inferential competences. One who is sensitive to the inductive force of evidence must be sensitive to certain sample characteristics–possible biasing factors, for example, or size with respect to the population at issue, for another. General inductive inferential competence requires such more specific stable dispositions or capacities–without which inductive inference from samples (in particular) would be subject to unacceptable pitfalls. (One might say the same thing for the simpler matter of deductive inferential competence.) Here one is reminded of Robert Cummins' (1975, 1983) treatment of componential and functional analyses as an explanatory strategy. Sophisticated capacities or dispositions of some system are typically explained in terms of the organized working of simpler capacities or dispositions possessed by that system or its components–the latter being required for the realization of the former. Thus, the virtue of inductive inferential competence turns in part on various simpler capacities–virtues–such as sensitivity to sample biases, and sensitivity to counter-evidence, to competing explanation, to possible common causes, and the like.

Not all cognitive competences are so clearly inferential in character. Consider a kind of *general perceptual competence*. To begin with, people develop *domain specific perceptual competences*: one might have developed a competence in identifying paintings and painters of the neo-romantic period, for example, and not in identifying beetle species. Agents can be more or less sensitive to when they have attained such a trained domain specific competence and when they have not. One who has such *general* stable perceptual sensitivities is perceptually competent in a particularly general way: such an agent readily arrives at perceptual beliefs only on matters on which they have reasonable domain specific competence–inhibiting judgments both where conditions do not afford a basis for confident judgment, and where training has not afforded a domain specific competence. This amounts to a *general perceptual competence*–and it seems on a par (with respect to both generality and epistemic

usefulness) with the general inferential competences mentioned already. Both general and domain specific perceptual competences seem to qualify as epistemic virtues. [3]

The virtue of general perceptual competence also depends on a range of sensitivities–including (often inarticulate) sensitivities to one's training or qualifications or track record in the relevant domain, and sensitivities to the quality of the particular presentation (light levels or background noise levels, degrees of obstruction, and the like). One's perceptual processes should be conditioned by relevant background information–which must be recognized and brought to bear. With sensitivities to such matters in place, one should freely form perceptual beliefs on matters of refined competence, unless inhibited by background, or sensitivity to degraded presentations, and one should inhibit judgments when training is sparse.

Perception and deductive, inductive, and abductive inference are not the only sources of belief, nor the only domains of general competence. Memory and testimony are also epistemically important. With respect to each, there would seem to be more specific competences or virtues.[4]

II. Why Virtue Epistemology

A. The Pivotal Contrast and Claim

Virtue epistemology is contrasted with what has become the traditional epistemological approach. Much epistemology from the last four hundred years has supposed that a rather limited range of cognitive processes is epistemically significant and crucial: (what we will here term) *classically inferential processes*. Intuitively, the idea has been that the epistemically relevant processes by which beliefs are formed and maintained are processes in the neighborhood of articulate (or at least articulatable) argument–processes in which informa-

3 It does not seem useful to extend the concept of epistemic virtue (and perceptual virtue or competence in particular) to apply to capacities within the retina. Visual acuity is not usefully thought of as a virtue. Rather, as a general rule, epistemic virtues will have to do with cognitive processes (although not necessarily articulate processes, as we will soon see)–processes determining what is done with the information one does receive. A perceptually competent agent learns to deal with, be sensitive to, the acuity of eye sight, hearing, and so on. A virtuous perceptual agent can have poor eye sight, provided that agent has developed dispositions that compensate for the degraded character of the visual input received.

4 One's competence with testimony would also seem to be dependent on a range of more particular sensitivities or dispositions (Henderson, forthcoming). Even one's general competence with memory might also yield to useful analyses (Goldberg and Henderson, forthcoming).

tion is represented (the premises) and in which these representations, and just these, figure causally and decisively in the transition to some new belief (or in the retentive re-evaluation and vindication of some antecedent belief). The intuitive idea can be unpacked by distinguishing three notions of inferential processes. In each, agents are taken to possess "information," where information is understood so as not to suppose truth (apparent information, but we will drop the qualifier 'apparent'). The information possessed provides some contentful (more or less strong) support for various claims or beliefs that the agent might form, abandon, retain, or revise.

Inferential processes, broadly understood, are simply those cognitive processes in which beliefs are formed or maintained on the basis of the information possessed. Being based on information is a causal notion, pointing to arrays of counterfactual dependencies and to dispositions. This is the broadest and most tolerant notion of an inferential process.[5] For an inferential process, broadly understood, to be epistemically appropriate, it must yield beliefs that are necessarily or probably true, given the correctness of the information on which it relies.

To qualify as a *classically inferential process,* a much more restricted class of processes, two additional things must be true of a cognitive process. First, in a classical inferential process, the information figuring in the inference is *explicitly represented* in the cognitive system that is the agent. This is not to say that the information need be consciously represented–just that, in taking its causal role in the cognitive system, that information is occurrently represented.[6] Second, the causal processes whereby beliefs are fixed (formed, revised, or retained) must be *occurrently isomorphic* with the deductive and inductive support relations obtaining between the information that the agent possesses. This is to say that the contentful support relations between the information possessed is mirrored by, or isomorphic with, causal dependencies between the representations in the agent's cognitive system. This obtains when the relevant psychological states (representations) are *occurrent*, and as such they causally conspire to generate or sustain the belief in question in a way that mirrors the way in which the pieces of information represented contentfully conspire to "support" the belief in question. A careful reading of a diversity of modern and contemporary texts suggests a wide partisanship for the pre-

5 Indeed, as noted later, this notion probably strains the notion beyond common usage, insofar as common usage associated inference with argument, and insofar as argument is understood as involving transitions between representations (premises and conclusions). We employ this broad usage here because, once one recognizes the possibility of the sorts of informed transitions we envision here, one will find it reasonably natural to refer to them as "inferential after a fashion."

6 This much does not require that the inference rules to which the system conforms are themselves represented.

supposition that empirically relevant processes are classical inferential processes.[7]

Cognitive processes might be broadly inferential without being classically inferential. For example, the agent might possess information that is not represented (as representations are standardly understood) and yet causally conditions the agent's belief fixing processes. (If this is difficult to envision, bear with us, and we will explain more fully.) If this obtains, then the agent's broadly inferential processes will not be occurrently isomorphic with the possessed information and contentfully appropriate transitions.

---

7 The epistemological project has commonly been conceived of as indicating by what arguments one can attain various forms of justified belief and knowledge regarding the wider world. This much is a fairly pervasive heritage of modern philosophy. Descartes' first philosophy sought a classically inferential path from self-evident truths to what is needed for modern science. Hume's more skeptical first philosophy sought to bring home the limited power of "reason"–by making evident the limits of classical argument to take us beyond certain starting points. Similarly, the justification of induction was thought to depend on deductive demonstrations of the likely reliability of inductive inference. Here, a priori argument, by tracking just the contentful relations between certain representational states, was to vindicate rules and arguments (or to fail to vindicate the inferences in question). Parallel demands were placed on the epistemic vindication of testimony. Kant reframes the epistemological project, not by repudiating the epistemic focus on classical inference, but by narrowing the ground that it is supposed to cover. The beginning of the epistemic chore–the starting position–is a set of perceptual judgments that are essentially more substantive than Humean appearances. The epistemic task is to do justice to the phenomenal world–the world as it appears to us. In effect, this takes us from the phenomena (the phenomenal objects, objects as they appear to us) to some suitable systemization of them (to accounts intended to be true of the phenomenal world). While much may have gone on in us to get us to the Kantian starting places–to the objects as they appear to us–the epistemological chore is the classically inferential chore of getting to a general account of the phenomenal world. More recent epistemology has retained this focus on classical inference–witness Carnap's attempt to develop an inductive logic in terms of observation reports and ranges of state descriptions. More recently, BonJour's (1985) coherentism, with its roots in Davidson (1983), provides a prominent example. Foundationalist approaches have generally faced an interesting problematic–that some finding regress stopping accessible states that can be understood as inferential bases (premises, reasons to think other claims true) and which yet do not themselves require inferential justification. Thus, the coherentist BonJour (1985), thought it unlikely that the foundationalist would succeed, while the foundationalist BonJour (1998) has decided that, on pain of skepticism, there must be the resources for such a structure of argument. The recurrent theme: epistemic agents must argue their way about on the basis of representations–that classical inference is the epistemic engine. This focus on inference from articulate reasons as premises has certainly been encouraged by access internalist epistemology. In Henderson and Horgan (2000) we develop a few examples of the contemporary focus on classical inference (BonJour 1985, Pollock 1986, Moser 1989, Audi 1993). It is clear that any epistemology that holds that epistemic agents must base their beliefs on representations, the content of which is internalistically accessible (any epistemology that conforms to access internalist constraints at least this far) will be an epistemology that is focused on classically inferential processes as the processes that make for objective epistemic justification (or its absence). It is worth noting that even epistemologists who incline towards a kind of reliabilism may think that the relevant (reliable or not) processes are those that make use of just such representations. Audi (1993) might serve as an example.

Our suggestion is straightforward: while some epistemically important cognitive dispositions may be dispositions to classical inference, not all are. Not all our epistemically important, productive, or needful chores are, or can be, well managed only at the level of classical inference. Insofar as some such chores must be managed in significant measure by way of dispositions that are broadly inferential but not classically inferential, a superior epistemological perspective will give significant attention to virtues–to epistemically good dispositions–beyond what has been the standard epistemological focus.

It will be helpful to have in play a third class of inferences–intermediate between broad and classical. Classical inferential processes are transitions between representations and representations, and are occurrently isomorphic with the contentful support relations obtaining between the featured representations. Broadly inferential processes move from "information," that may, but need not be represented, to some belief. The dynamics of the transitions may be conditioned by informational states that are not representations. Now, talk of "broadly inferential processes" is a little strained, for inference suggests argument, and argument suggests propositional representational states (premises and conclusions). Broadly inferential processes as we understand them need not begin with propositional representational states–and so need not begin with premises in the standard sense. Perceptual processes, for example may begin with much information that is not propositionally represented, and so are ill understood in terms of inference as a matter of argument.[8] Now, to get to our intermediate category of inference, one can envision cognitive processes that make for transitions between propositional representations while also being conditioned in part by some or much information that is possessed by the system and yet not represented in that processing. Such processes would *not be occurrently isomorphic* with the contentful support relations obtaining between the representational states–as they would be conditioned by information that is possessed but not represented. Such processes, involving transitions between propositional representational states, which do not satisfy the occurrent isomorphism requirement for classical inference, constitute a kind of inference intermediate to broad and classical. Because they are marked by having premises and conclusions (propositional representational beginnings and endings), we can write of *argumentative inference*.

8 These point are related to the power of virtue theory to move beyond an overly constraining concern with reasons in epistemology–a strength early appreciated by Sosa (1980).

We need not here commit to whether perceptual processes begin with nonpropositional representations of some sort (whether the initial states in the process–the initial total activation space in a given perceptual episode, perhaps–need be a representation)–and the matter will depend on just how representation is understood. In any case, the present point reflects at least one sense in which perceptual processes may be thought of as direct–as we suggest that they commonly are neither classically or argumentatively inferential.

Some of the stable dispositions that interest us here often mediate transitions between propositional representations–and thus are dispositions to argumentative inference. But, we will argue that the epistemically important dynamic of these transitions often cannot be understood wholly at the level of representations–articulate or not. When this is so, the processes are broadly and argumentatively inferential, but not classically inferential. (All classical inference is argumentative inference, but not all argumentative inference need be classical inference.) To the extent that processes crucial to human epistemic successes fail to qualify as classically inferential, and are rather either argumentatively or broadly inferential, epistemology must abandon the fixation on classical inference and must attend to a wider set of dispositions than is common. It must become a distinctively virtue epistemology.

Prominent in our argument for virtue epistemology will be our discussion of central processes of belief formation, processes of inductive reasoning and abductive reasoning. We argue that the epistemic chores an agent faces here cannot be managed wholly at the level of classical inference. Such argumentative inference draws on information that is essentially richer than what is or can be represented. The epistemically virtuous dispositions are dispositions to accommodate such information, and thus are not classically inferential.

We now flesh out the suggestion that certain subtle dispositions (commonly acquired or refined and developed by training), (a) are crucial to systematic human epistemic successes on a range of cognitive chores (chores that are commonly recognized by epistemologists to be important and necessary), and (b) are not readily understood fully in terms of classical inference. The virtue theorist sees the epistemological fixation on classical inference as analogous to the contemporary automotive industry's overwhelming reliance on the internal combustion engine–it is time we recognize the promise of hybrid technology.

B. The Argument

Human epistemic competence, that combination of processes by which humans do and should manage their epistemic chores, cannot be understood solely in terms of classical inference. Among those epistemic chores widely recognized as epistemically needful are ones managed, not by classical inference, but by dispositions keyed to essentially richer sets of information. By way of example, we now discuss two epistemic chores: (1) chores within central processes of belief fixation requiring the holistic accommodation of vast ranges of potentially relevant information, and (2) chores involved in fast, sensitive, formation of perceptual beliefs–an activity or process that is itself permeable to, or conditioned by, much background information. Both reflect

an essential place for dispositions essentially richer than what can be managed by classical inference. Our discussion of the first of these chores develops the case most fully, while our discussion of the second will need to be more suggestive.

The Chores, Part 1: Holistic Processes of Central Belief Fixation

a. Holism, the frame problem for computationalism, and the need for morphological content.

Certain problems and developments in cognitive science provide reason to believe that the human cognitive processes that typically produce and sustain beliefs are not of a sort that is well understood when preoccupied with classically inferential processes.[9] There is reason to believe that human belief generation *must*, and *epistemically ought*, proceed in an alternative manner. The alternative understanding emerges in the wake of a family of recalcitrant difficulties within the classical, computational, conception of mind in cognitive science–difficulties often classified under the rubric "the frame problem." In this section we first describe those difficulties, drawing upon an influential discussion by Fodor (1983), and recommend the alternative conception suggested by frame-type problems. We then sketch one way of elaborating this conception, inspired by connectionism and by the form of mathematics that goes naturally with it, dynamical systems theory. The frame problem is a direct challenge to computational cognitive science–but the alternative understanding of epistemically virtuous belief generation to which it points is incompatible with the standing epistemological fixation with classically inferential processes.

In the closing pages of *The Modularity of Mind,* Fodor (1983) argues that certain problems in classical cognitive science look to be in-principle problems, and hence that the prospects for understanding human belief fixation wholly within the framework of classical cognitive science are very bleak. These problems continue to plague the computational approach.

The main claim of Fodor's influential book is that the human cognitive system possesses a number of important subsystems that are *modular*: domain specific, mandatory, limited in their access to other parts of the larger cognitive system, fast, and informationally encapsulated. Where the classical computational approach has gotten somewhere, he says, is in understanding such modular subsystems, which by their nature delimit the class of relevant in-

9 Access internalist epistemology has encouraged this problematic preoccupation with classical inference, and it is also called into question by these results. The points made in the present section are closely related to discussions in Henderson (1994) and Henderson and Horgan (2000). We here draw on material in Horgan and Tienson (1994, 1996) and Horgan (1997a).

formation. However, classical computationalism has made very little progress in understanding what he terms *central* processes. Belief fixation on the basis of current input together with other beliefs–is a paradigmatic example. These processes are non-modular: they need to have access to a wide range of cognitive subsystems, and to information on an indefinitely wide range of topics. And the very considerations that point to non-modularity, he maintains, also constitute grounds for extreme pessimism about the prospects of explaining central processes within the framework of classical computational cognitive science. For our purposes here, this is reason for pessimism about the prospects for understanding central processes of belief fixation wholly in terms of classical inference.

Fodor articulates these considerations in terms of the analogy between belief fixation in human cognition and scientific confirmation. This is to garner lessons about central cognitive processes like belief fixation from what we know about "empirical inference in science" (Fodor, 1983, 104). Scientific confirmation, "the nondemonstrative fixation of belief in science," has two crucial features. It is (in Fodor's terminology) *isotropic* and *Quineian*:

> By saying that confirmation is isotropic, I mean that the facts relevant to the confirmation of a scientific hypothesis may be drawn from anywhere in the field of previously established empirical (or, of course, demonstrative) truths. Crudely: everything that the scientist knows is, in principle, relevant to determining what else he ought to believe.... (1983, p. 105)

> By saying that scientific confirmation is Quineian, I mean that the degree of confirmation assigned to any given hypothesis is sensitive to properties of the entire belief system; as it were, the shape of our whole science bears on the epistemic status of each scientific hypothesis (1983, p. 107).

Isotropy brings in the whole of current theory: any bit of apparent information from any portion of one's belief system might, in some circumstances, be evidentially relevant to any other. Being Quineian makes confirmation holistic in a deeper way: confirmation depends upon "such considerations as simplicity, plausibility, and conservatism" (Fodor, 1983, p. 108), which are determined by the global *structure* of the whole of the current belief system and of potential successor systems.

Since belief fixation in human cognition is commonly a matter of inductive inference from the information provided by input systems and the information antecedently possessed, evidently it too must be isotropic and Quineian. Fodor concludes that it must be non-modular. He also stresses that these global aspects of belief fixation look to be at the very heart of the problems that classicism has encountered in attempting to understand such central processes:

> The difficulties we encounter when we try to construct theories of central processes are just the sort we would expect to encounter if such processes are, in essential re-

> spects, Quineian/isotropic.... The crux in the construction of such theories is that there seems to be no way to delimit the sorts of informational resources which may affect, or be affected by, central processes of problem-solving. We can't, that is to say, plausibly view the fixation of belief as effected by computations over bounded, local information structures. A graphic example of this sort of difficulty arises in AI, where it has come to be known as the "frame problem" (i.e., the problem of putting a "frame" around the set of beliefs that may need to be revised in light of specified newly available information) (Fodor, 1983, pp. 112-3).

Let us take a closer look at the challenges posed by isotropy in particular. The distinction between being justified versus merely having justification turns on the idea that, when one is justified in one's belief, the justification that one has for that belief is not causally inert, but rather "comes into play" or is "appropriately operative" in generating or sustaining that belief. Whatever information one might possess, if one does not systematically use it, and use it in an appropriate way, then, even if one somehow arrives at true beliefs, one does so by accident, and not by justificatory processes. For such distinctions to make sense, it must be possible for one to have justifications for a belief which do not, and did not, "come into play" in the fixation of that belief. In such cases, something else must then be relevantly operative in belief-fixation. Perhaps one has "other justification" for the belief in question, and these might be there operative in the relevant belief-fixation (in which case one may yet be justified). Or perhaps what is in play in fixing the belief in question is not justificatory information–not something that one should talk of as justification that one has (in this case one has justification, but is not justified). Such possibilities are presupposed by the distinction between being justified in believing and merely having justification for a belief. The distinction requires the *selective efficacy* of beliefs and information states representing the justification that the agent has.

Holy Cross 9/20/07 1:54 PM

Comment: David, I think I know what you are getting at with your sentences up to the end of the paragraph. Yet I find them a bit hard to comprehend. Concrete examples might be helpful here. I am also wondering whether you really need this paragraph?

One might think that holism is incompatible with such distinctions. But, a closer look at the holistic dimension that Fodor terms "isotropy" should serve to dispel this suggestion..When Fodor writes of confirmation as isotropic he insists that "the facts relevant to the confirmation of a scientific hypothesis *may* be drawn from anywhere in the field of previously established empirical (or, of course, demonstrative) truths. Crudely: everything that the scientist knows is, in principle, relevant to determining what else he ought to believe" (p. 105, emphasis added). Fodor's first formulation reflects a point that is lost in the crude reformulation. Isotropism is the *potential relevance* of everything (all the beliefs and information one has) to everything else one believes–but this is *not* to say flatly that everything one believes is *actually relevant* to everything else one believes. One must be careful in conceiving of relevance here, and it will be helpful to distinguish a strict sense, *actual relevance*, and an extended sense that might best be thought of as *relevance to relevance*. In the extended and attenuated sense it is correct that everything that one believes (and more gen-

erally, all the information that one in some sense possesses) is relevant to any episode of belief fixation. Take any two beliefs that are mutually irrelevant on their face–that is, when considered of themselves, without taking into account what else one might believe. Then notice that, with various sets of possible auxiliary information or beliefs serving as intermediaries, the beliefs that were of themselves irrelevant would become mutually relevant. What is *actually relevant* to what depends then on what else one *actually* happens to believe–or on what other information one actually has (and how it "adds up"). To then be sensitive to what is actually relevant to what one does or might believe, one must be sensitive to how the set of what else one believes makes for various stripes of mediated actual relevance. Much actual relevance is mediated relevance, so this is important. To be sensitive to mediated relevance, and thus to much actual relevance within one's doxastic and informational set, one must be sensitive to the full set of one's information or beliefs. Thus, in view of the *potential* relevance of everything to everything, everything that one believes, *all one's information, becomes relevant in the extended or attenuated sense–it is relevant to relevance–and one must be appropriately sensitive to this in belief fixation*. Only then can the actually relevant beliefs and information "come into play," as required if the agent is to be justified in the resulting belief as opposed to merely having justification.

Holy Cross 9/20/07 2:03 PM
Comment: Is there something wrong with the sentence? Is an "is" missing? i.e…."be sensitive to what is relevant to what-actually is relevant- one must…"

The prospects for understanding central processing within the classical computational paradigm look very discouraging indeed. Not only do we have no computational formalisms that show us how to manage the epistemic chores associated with the Quinean and isotropic aspects of central processes of belief fixation; it's a highly credible hypothesis that a (tractable) computational system with these features is just impossible, for belief systems on the scale possessed by human beings. Human central processing evidently does not operate via any kinds of computation we currently know about or can even contemplate. Something else is needed. What might it be?

These frame-type problems arise largely because of the apparent computational intractability of managing all relevant information, insofar as that information gets *explicitly represented* in the course of cognitive processing. What this suggests is that belief fixation and related cognitive processes operate in a way that accommodates much relevant information *automatically* and *implicitly*. The suggestion is that the holistic aspects of belief fixation involve not the finding and fetching of relevant representations from memory-banks where they are stored in explicit form (to accommodate isotropy), and not the overt representation and comparative evaluation of large-scale alternative belief-systems (to accommodate the Quinean dimension). Rather, these holistic aspects are somehow implicit in the structure of the cognitive system, in such a way that temporal transitions from one occurrent cognitive state to another accommodate the holistic aspects automatically. In the terminology

of Horgan and Tienson (1995, 1996), important holistic informational content is *morphological*, rather than occurrent.

Morphological content is information that:

(i) is implicit in the standing structure of the cognitive system (rather than explicitly represented in the system's occurrent cognitive states or explicitly stored in memory), and

(ii) gets accommodated in cognitive processing without getting explicitly represented in occurrent cognitive states, either conscious or unconscious.

The apparent moral of the frame problem is that, in general, human belief fixation *must* operate in a way that draws heavily upon morphological content, in order to avoid computational intractability. As we will put it, these processes are *essentially* morphological. Belief fixation is not accomplished (simply) by computationally manipulating explicit, occurrent, representations of all relevant information. Nor are they accomplished by "proceduralized" computational processes that are mere shorthand algorithms for computations that could instead have been carried out in a way that renders all relevant information explicitly. Essentially morphological processing is a fundamentally different way of accommodating the holistic aspects of belief fixation.

Recognition of the holistic elements of belief-fixation, of its Quinean and isotropic character, is not, at least not primarily, a product of cognitive science. Rather, it arises out of fairly traditional epistemological reflection–for example, out of reflection on the role of ancillary hypotheses and webs of belief that can be found in writings of philosophers such as Hempel and Quine. Reflecting on how the implications of one sentence commonly turn on what other sentences one combines with it highlights the information to which a successful epistemic agent must be sensitive–revealing the informational tasks facing an epistemic agent. (One could point to a vast range of philosophical work–to name a few examples, to BonJour (1985), to Shapere's (1982) discussions of the range of information and understandings that inform "observations" in physics, to Goodman's (1973) work on induction and the projectibility of certain predicates–itself gauged in terms of much background information.)

Central processes of belief fixation then must be handled essentially morphologically–so that much possessed information is accommodated automatically, by way of morphological content. Not only does this doom classical computational cognitive science, *it also indicates that these epistemic chores cannot be managed by way of processes wholly at the level of classical inference. Significant ranges in information must condition belief fixation without being represented–and thus these processes cannot be occurrently isomorphic with contentful relevance relations obtaining between the information deployed here. Dispositions subtler and richer than those evinced at the level of classical inference are epistemically essential.*

b. A way to think about morphological content.

How might the daunting task of essentially morphological processing get accomplished in human cognition? To our knowledge, there certainly are no models in cognitive science that come close to achieving such processing for cognitive tasks even remotely comparable in complexity to those handled in real human thought. Nor is there any good reason, as far as we know, to think that any extant models are likely to "scale up" or be straightforwardly extended to provide an adequate account of real human cognition. This is no less true for connectionist models than it is for classical computational models.

Nonetheless, a general conception of human cognition has begun to emerge within cognitive science that is potentially more powerful than the classical computational conception of mind, and that provides the broad outlines of an answer to the question, "How is essentially morphological processing possible?" This alternative conception draws cautiously upon connectionist modeling, in a way that eschews unduly optimistic assumptions about the scale-up potential of extant models. It also draws upon a form of mathematics that is natural for describing connectionist models–dynamical systems theory. Here we offer a very brief sketch of connectionism, dynamical systems theory, and the nonclassical framework–with emphasis on features that are especially germane to morphological content.[10]

In a connectionist system, information is actively represented as a pattern of activation. When the information is not in use, that pattern is nowhere present in the system; it is not stored as a data structure. The only representations ever present are the active ones. On the other hand, information can be said to be *implicitly* present in a connectionist system–or "in the weights"–if the weighted connections subserve *representation-level dispositions* that are appropriate to that information. Such information constitutes morphological content in the system, rather than explicitly-represented content. Among the apparent advantages of connectionist systems, by contrast with classical computational systems, is that morphological information "in the weights" gets accommodated automatically during processing, without any need for a central processing unit to find and fetch task-relevant information from some separate memory banks where it gets stored in explicit form while not in use.
Learning is conceived quite differently within connectionism than it is within the classical approach, since connectionist systems do not store representations. Learning is the acquisition, "in the weights," of new morphological content.

10 This nonclassical framework for cognitive science is described at length in Horgan and Tienson (1996).

The branch of mathematics called dynamical systems theory is often applied to connectionist models. To describe some physical system (e.g., a planetary system, or a connectionist network) mathematically as a dynamical system is to specify in a certain way *its temporal evolution, both actual and potential.* The set of all possible states of the physical system--so characterized--is the mathematical system's abstract, high-dimensional *state space.* Each magnitude or parameter of the physical system is assigned a separate dimension of this mathematical space, and each possible state of the physical system, as determined by the values of these magnitudes, corresponds to a point in state space. A dynamical system, as such, is essentially a complete mathematical description of *how the physical system would evolve temporally from any possible initial state*; it is a collection of trajectories through state-space, with a trajectory emanating from each point in state space. In essence, it is a description of a complex disposition. A useful geometrical metaphor for dynamical systems is the notion of a *landscape.* A dynamical system describing a physical system involving n distinct magnitudes is the n-dimensional analog of a two dimensional, non-Euclidean, contoured surface: i.e., a topological molding of the n-dimensional state space such that, were this surface oriented "horizontally" in an (n+1)-dimensional space, a ball would "roll along the landscape," from any initial point *p*, in a way that corresponds to the way the physical system would evolve from the physical state corresponding to *p*.

Connectionist systems are naturally describable, mathematically, as dynamical systems. The state space of a network is its "activation space" (which has as many dimensions as the network has nodes), and the dynamical system associated with the network is its "activation landscape." In connectionist models, cognitive processing is typically construed as evolution along the activation landscape from one point in activation space to another--where at least the beginning and end points are interpreted as realizing intentional states.

So in terms of the mathematics of dynamics, occurrent cognitive states are realized mathematically as *points on the activation landscape*, which are then realized physically as distributed patterns of activation in the nodes of the network. Morphological content–the information implicit "in the weights"–is embodied in the *topographical contours* of the network's high-dimensional activation landscape. Thus, the various superimposed slopes on the activation landscape subserve trajectories from one occurrent cognitive state to another that automatically accommodate the morphological content.

From this mathematical perspective, training a network is a matter of (i) *molding* the activation landscape, thereby inducing new topological contours embodying new morphological content, while simultaneously (ii) *refining* the cognitive/mathematical realization relation whereby intentional states get realized mathematically as points on the landscape. (The weight-change train-

ing procedures employed in connectionist modeling bring about a co-evolution of these two factors.) Once "trained up," the system's temporal trajectories from one occurrent intentional state to another will automatically accommodate the relevant morphological content.

Horgan and Tienson (1994, 1996) describe a non-classical framework for cognitive science that they call the *dynamical cognition* framework (the DC framework). This alternative approach offers an answer, in principle, to the question, "How could the holistic, Quinean/isotropic, aspects of cognitive processes be accommodated automatically and morphologically? The answer is this: In principle, Quinean/isotropic information could be embodied morphologically in the complex and subtle topography of a high-dimensional activation landscape subserved by the human central nervous system. Given a sufficiently nuanced realization relation from cognitive states to points on this landscape, the landscape's multifarious, superimposed, topographical contours guarantee that the cognitive system's transitions from one occurrent cognitive state to another are automatically appropriate not only to the explicit content of these occurrent states themselves, but also to very large amounts of implicit Quinean/isotropic information.

We maintain that the frame-type problems encountered in classical cognitive science provide a strong reason to maintain that the holistic, Quinean/isotropic aspects of belief-fixation are accommodated in an essentially morphological way in human cognition. We also maintain that this conclusion is further reinforced by the in-principle account of such essentially morphological processing that is provided by the DC framework. But it should be noted that the appeal to the DC framework is not essential here. Even if one were thoroughly dubious about connectionism and about the usefulness of dynamical-systems ideas in cognitive science, one still could accept the argument from frame-type problems to the conclusion that processes like belief-fixation must *somehow* be essentially morphological.

c. Morphological content, dispositions, and virtue epistemology

The central claims in the above subsections are these: First, we have argued that there are holistic aspects of belief fixation. Isotropy means that the central processes must accommodate the potential relevance of each piece of possessed information to any given belief, and thus that all information possessed is relevant to what is relevant. The processes by which an individual manages to fix beliefs must be sensitive to what is relevant–and this cannot be managed computationally. Being Quinean means that global features of belief systems are themselves relevant, and it seems likely that this itself also cannot be gauged computationally. Second, we have suggested that, in managing the

indicated epistemic chores, the competent human agent must accommodate much possessed information automatically without representation, and thus without it featuring as a premise in a classical inference. The competent human cognitive agent must make use of morphological content. Third, while there are certainly broadly or argumentatively inferential moments in the processes here envisioned, the whole cannot be fully understood as classically inferential. Here, it is good to reflect on the dynamical system framework we have recommended. One might think of *some* of the transitions between representational states as inferences on the order of those that have been the common epistemological fare–as transitions between representations (premises and conclusions). But remaining only at that level, one cannot fully understand the dynamics of belief fixation by which crucial holistic chores are managed. Instead of remaining at the level of what is classically inferential, it is best not to loose sight of *the cognitive system's dynamics* and *the morphological content* on which it relies. *This is to focus on epistemically crucial dispositions* beyond those that have been the common epistemological fare. Such virtues are broadly inferential, and are sometimes argumentatively inferential, but are not restricted to the sort of classically inferential transitions that have too commonly marked the limits of epistemic concern. Only with this less constrained focus can one appreciate the full range of information in play within central processes of belief fixation, and how such information might have an impact on such processes.

Consider the dynamical system as a topography of the system's weight space. This is a matter of its tendencies to move from one total state (or pattern of activation) to others–to descend in state (or activation) space. Only some of the patterns of activation through which it would descend, from a given point in activation space, are themselves representations–the sole stuff of classical inferences. But the whole dynamic is a more subtle disposition, one that reflects the capacity of the human cognitive system to make sensitive use of vast ranges of information. The capacities or dispositions that the system acquires through courses of experience, through training, are highly desirable epistemically–are epistemic virtues. Some of the patterns of activation through which the system passes are representations, so some of the transitions between such representations will count as argumentative inferences. But, just as not all the information that the system possesses is represented, or even could all be represented, and just as not all the transitions are determined or conditioned only by representations, not all the epistemically desirable dispositions of the system are to be understood as at the level of classical inference. We thus arrive at the general theme or claim on which this paper pivots: some epistemically important cognitive dispositions are dispositions to inference broadly conceived that cannot be fully understood at the level of classical inference. Insofar as not all our epistemically important, productive,

or needful chores are or can be well managed only at the level of classical inference, insofar as they must be managed in some significant measure by way of dispositions that are not, strictly speaking, classically inferential, then a superior epistemological perspective will give significant attention to virtues–to epistemically good dispositions–beyond the classically inferential sorts that have been common epistemological fare.

In his virtue ethics, Aristotle was notably reticent regarding precise principles or rules by which the morally appropriate action could be determined. He seems to suggest that no set of precise exceptionless rules can ultimately capture what it is to live a morally good or correct life–stated somewhat anachronistically: one cannot write a program for the moral life. Rather than rules–he urges kinds of training, kinds of sensibilities or capacities to foster. Famously, he insists that the correct action is the action that the morally virtuous person would undertake–and that this is a matter of a rational judgment drawing on trained sensibilities.[11] The suspicion of precise rules (moral rules or rules of inference on the order of a program) has been a recurring feature of a number of virtue theories, moral or epistemic. In light of our discussion of central processes of belief fixation, and of the place for morphological content in particular, we can at least suggest why there might be something importantly right about this suspicion of rules as the end all and be all of epistemology–at least if one thinks of rules as something on the order of a computational program for belief-fixing processes. Belief fixation is probably too complex and subtle to conform to programmable rules at all.[12]

In this paper, the central processes of belief fixation serve as our central exhibit of the epistemological importance of dispositions to broadly inferential transitions that cannot be understood in terms of classical inference. We have developed the case at some length. The general picture of epistemically important processes and dispositions that has emerged can be applied to other epistemological topics–to the epistemic generation or use of perception, memory, or testimony, for example. In each case, one would find that belief formation can and should be conditioned by, or permeable to, indeterminate ranges of possessed information that must be accommodeated quickly and it seems largely automatically. The central processes of belief fixation we have

11 In the *Nichomachean Ethics,* one reads:

> I mean moral virtue; for it is this that is concerned with passions and actions, and in these there is excess, defect, and the intermediate. For instance, both fear and confidence and appetite and anger and pity and in general pleasure and pain may be felt both too much and too little, and in both cases not well; but to feel them at the right times, with reference to the right objects, towards the right people, with the right motive, and in the right way, is what is both intermediate and best, and this is characteristic of virtue (Aristotle, Nicomachean Ethics, II 6, 1106b, 15-20.).

12 See Horgan and Tienson (1996) for a fairly extended discussion of this point.

discussed presented us with transitions that are broadly and argumentatively inferential in character.

Before closing, we want to discuss (in a more abbreviated fashion) an epistemic context where the transitions are broadly inferential, less clearly argumentatively inferential, and clearly not classically inferential. Consider perceptual processes. Note that perceptual processes are akin to central processes of belief fixation in a very important way–they are permeable to, and necessarily sensitive to, wide ranges of background information. Again, this is managed very quickly–and automatically. Again, it is a highly plausible hypothesis that this is managed by reliance on morphological content within a kind of cognitive dynamical system. Again, virtue epistemology seems vindicated.

The Chores, Part 2: Sensitive Perceptual Processes

What is it to be perceptually competent with respect to some limited matter? What is it to be perceptually competent generally? When one is a competent perceptual judge of the wildlife in one's environment, for example, one can with reasonable sensitivity respond to common episodic encounters (and to more or less enduring traces) in one's environment. Due to one's sensitivity, one can produce reasonably accurate identifications of the species of wildlife involved in the episodes (or leaving the traces). One who is perceptually competent with respect to a delimited domain is able to make reliable judgments in a range of common environmental conditions. Thus, one who is perceptually competent with respect to local wildlife can make reliable identifications when various creatures are partially obscured, quickly glanced, or imperfectly illuminated (common conditions). Importantly, such capacities thenselves require reasonable sensitivity to when conditions do not allow the agent to make a reliable identification. A competent perceptual judge on a given subject matter is thus one who has been trained up to be reasonably reliable on the matter in question–one who can consequently render verdicts that are likely true–and this requires not rendering verdicts in certain ranges of difficult cases.[13]

13 For reasons having to do with the new evil demon problem, and pointing to the epistemic significance of certain nonstandard forms of reliability, these remarks need to be qualified. They hold true of perceptual competence for agents in reasonably hospitable epistemically possible global environments–such as that provided by the actual global environment. However, from the epistemic point of view, one can be perceptually competent in a demon-infested global environment, where all ones training has been deceptive, and apparent successes and failures have resulted in an exquisitely sensitive perceptual system that is nevertheless not globally reliable. For more on the relevant forms of reliability, see Henderson and Horgan (in press). Montmar-

It is important to pause and notice the range of information that the agent – as cognitive system – must possess and use, if the agent is to exercise a domain specific perceptual competence in a given episode. The agent will need to sensitively possess much information particular to the episode. This information need not, and commonly will not be represented in the perceptual processes. For example, there need be no representation of the degree of occlusion, nor of just what "parts" are shielded from view, in the course of the perceptual process by which the agent generates or refrains from generating a judgment in response to a fleeting glimpse of some member of the local fauna. Yet, a competent judge of local fauna will have acquired a sensitivity to such matters. Similarly for light levels, shadows, and the like. Auditory input may condition the visual reception–and again in ways that are not fully susceptible to argumentative reconstruction. In competent and epistemically highly laudable perceptual processing, much such information, as morphological content, is likely accommodated automatically and very quickly. The information accommodated in a competent perception is typically not limited to information about the particulars of the episode. The perception of the competent judge will be conditioned by much background information–in ways that seem not at all dissimilar to the way in which wide ranges of relevant information can be accommodated in central processes of belief fixation. Consider the hesitance that a competent ornithologist would experience when, in the United States, confronted with a passing glimpse of a large woodpecker having what appears to be a white bill and white on both leading and trailing edges of its underwings. Perhaps the location is one in which Pileated Woodpeckers are common. An ornithologist would have some reasonable feel for the relevant ranges and characteristic population densities, of Pileated Woodpeckers and other candidates–of base rates. But the Pileated Woodpecker has a relatively dark bill and a black trailing edge of the underwings. What of the Ivory Billed Woodpecker? Obviously very different base rates (even given some uncertainty regarding whether it is extinct or just exceedingly rare). The perceptual process of a competent agent is permeable to such information, and much more of relevance. One might think of the processes as giving rise first to a perception or perceptual seeming that serves at the basis for a perceptual belief. The substantive perceptual seeming might be something on the order of the appearance of a large bird with a light colored beak and wing patches on a black background, moving through the dappled sunlight of deep forest cover. Perceptual seemings, given or taken, involve what seem to be salient objects standing forth–for example, an animal of a more or less familiar kind situated with respect to the agent and certain famil-

quet (1987) rejects the association of virtue with reliable processes, but we believe that the better move is to refine the understanding of the relevant form of reliability.

iar things in the agent's immediate environment. Whatever goes on in the processes yielding such perceptions must be quite sophisticated and complex, being conditioned by the ranges of information already suggested. Typically, the competent agent proceeds without hesitation or reflection to a perceptual belief. Occasionally, however, the perception or perceptual seeming will arise with an uneasiness. It may arise with what we might call a "warning flag," commonly associated with an issue demanding further attention. Further observation may seem needed because of light levels or angles, or occlusion. Or there may be the thought that the seeming is very unlikely to be true. Here, background information will have occasioned restraining thought–much as, in central processes of belief fixation, background information may be accommodated so as to make certain information stand out as relevant. In such cases, further thought may have an argumentatively inferential character. But, in the more straightforward episodes, it is plausible that the agent's perceptual processes need not pass through propositional representations on the way to the forming of a perceptual belief–and thus may be broadly inferential without being argumentatively or classically inferential.

We believe that perceptual processes–with their conditioning by, or permeability to, much background information that is accommodated automatically–require dispositions that are much like those involved in central processes; they are dependent on complex and subtle dispositions that are rightly thought of as epistemic virtues. Since the starting places–the appearances–are not comfortably seen as premises, it seems best to categorize the rudimentary perceptual processes making for the transitions here as neither classically nor argumentatively inferential. These broadly inferential perceptual processes can then occasion argumentatively inferential processes in some cases.

What then of an agent who is *perceptually competent generally*? This would seem to be an agent who, by virtue of training, has acquired a reasonable range of domain specific perceptual competences, and who tends to refrain from forming perceptual judgments with respect to domains about which such competence is yet to be acquired. Such sensitivity to one's sensitivities is itself an epistemic virtue.

We have argued that human epistemic competence cannot be understood solely in terms of classical inference. Among the epistemic chores widely recognized as epistemically needful are ones managed, not by classical inference, but by dispositions keyed to essentially richer sets of information. Central processes of belief fixation turn on holistic sensibilities that require the use of morphological content. So also do competent human perceptual processes. These have served to illustrate the epistemological importance cognitive dispositions beyond those commonly of concern in epistemology.

Holy Cross 9/20/07 2:40 PM

Comment: A few concluding sentences seem to be in order that would provide some psychological "closure" to the reader. Otherwise the article just ends rather abruptly.

Literature

Aristotle. 1941. *Nicomachean Ethics.* W. D. Ross, trans. In *The Basic Works of Aristotle.* Ed., R. McKeon. New York: Random House..

Audi, Robert. 1993. *The Structure of Justification*. Cambridge: Cambridge University Press.

BonJour, Laurence. 1985. *The Structure of Empirical Knowledge*, Cambridge: MA: Harvard University Press.

BonJour, Laurence. 1998. *In Defense of Pure Reason.* Cambridge: Cambridge University Press.

BonJour, Laurence. 2002. *Epistemology: Classic Problems and Contemporary Responses.* Oxford: Rowman and Littlefield.

Burge, Tyler. 1993. "Content Preservation." *Philosophical Review* 102: 457-88.

Burge, Tyler. 1997. "Interlocution, Perception, and Memory." *Philosophical Studies* 86: 21-47.

Coady, C. 1992. *Testimony: A Philosophical Study.* Oxford: Clarendon Press.

Coady, C. 1994. "Testimony, Observation and Autonomous Knowledge." in *Knowing from Words.* Eds. B. K. Matilal and A. Chakrabarti, 225-250. Dordredcht: Kluwer Academic Publishers.

Cummins, Robert. 1975. "Functional Analysis." *Journal of Philosophy* 72: 741-60.

Cummins, Robert. 1983. *The Nature of Psychological Explanation.* Cambridge, MA: M.I.T. Press.

Davidson, Donald. 1983: "A Coherence Theory of Truth and Knowledge. In *Kant oder Hegel.* Ed. Dieter Henrich. Stuttgart: Klett-Cotta.

Fodor, Jerry. 1983. *The Modularity of Mind: An Essay in Faculty Psychology*, Cambridge, MA: M.I.T. Press.

Goldberg, Sanford and Henderson, David. Forthcoming. "Monitoring and Reductionism in the Epistemology of Testimony." *Philosophy and Phenomenological Review.*

Goodman, Nelson. 1973. *Fact, Fiction, and Forecast*, 3rd edition. Indianapolis: Bobbs-Merrill.

Greco, John. 2000. *Putting Skeptics in Their Place.* Cambridge: Cambridge University Press.

Henderson, David.1994. "Epistemic Competence and Contextualist Epistemology." *Journal of Philosophy* 91: 627-49.

Henderson, David. Forthcoming. "Testimonial Belief and Epistemic Competence." *Nous.*

Henderson, David and Horgan, Terrence. 2000. "Iceberg Epistemology." *Philosophy and Phenomenological Research* 61: 497-535.

Henderson, David and Horgan, Terrence. forthcoming."The Ins and Outs of Transglobal Reliabilism." In *Internalism and Externalism in Semantics and Epistemology.* Ed. Sanford Goldberg. Oxford: Oxford University Press.

Horgan Terrence and Tienson, John. 1994. "A Nonclassical Framework for Cognitive Science." *Synthese* 101: 305-45.

Horgan Terrence and Tienson, John. 1995. "Connectionism and the Commitments of Folk Psychology." *Philosophical Perspectives* 9: 127-52.

Horgan Terrence and Tienson, John. 1996. *Connectionism and the Philosophy of Psychology.* Cambridge, MA: M.I.T. Press.

Kitcher, Philip. 1992. "The Naturalist Returns." *Philosophical Review* 101: 53-114.

Montmarquet, James. 1987. "Epistemic Virtue," *Mind* 96: 482-97.

Moser, Paul. 1989. *Knowledge and Evidence.* Cambridge: Cambridge University Press.

Pollock, John. 1986. *Contemporary Theories of Knowledge.* Lanham, Md: Rowman and Littlefield.

Sosa, Ernest. 1980. "The Raft and the Pyramid: Coherence Versus Foundationalism in the Theory of Knowledge." *Midwest Studies in Philosophy* 5: 3-25. (Reprinted in Sosa 1991.)

Sosa, Ernest. 1991a. *Knowledge in Perspective.* Cambridge: Cambridge Univ. Press.

Sosa, Ernest. 1991b. "Reliabilism and Intellectual Virtue," In *Knowledge in Perspective.* E. Sosa, 131-45. Cambridge: Cambridge University Press.

Shapere, Dudley. 1982. "The Concept of Observation in Science and Philosophy." *Philosophy of Science* 49: 485-525.
Woodward, James. 2003. *Making Things Happen*. Oxford: Oxford University Press.
Zagsebski, Linda. 1996. *Virtues of the Mind*. Cambridge: Cambridge University Press.