

Iceberg Epistemology

DAVID HENDERSON AND TERRENCE HORGAN

The University of Memphis

Accounts of what it is for an agent to be justified in holding a belief commonly carry commitments concerning what cognitive processes can and should be like. A concern for the plausibility of such commitments leads to a multi-faceted epistemology in which elements of traditionally conflicting epistemologies are vindicated within a single epistemological account. The accessible and articulable states that have been the exclusive focus of much epistemology must constitute only a proper subset of epistemologically relevant processing. The interaction of such states looks rather contextualist. It might also be called quasi-foundationalist. However, in attending to our epistemological tasks we must rely on processing that is sensitive to information that we could not articulate, that is not accessible in the standard internalist sense. When focusing on the full range of epistemologically important processes, the structure of what makes for justification is rather more like that envisioned by some coherentists.

Contemporary accounts of what it is for an agent to be justified in holding a given belief commonly carry substantive commitments concerning what cognitive processes can and should be like. In this paper, we argue that concern for the plausibility of such psychological commitments leads to significant epistemological results. In particular, it leads to a multi-faceted epistemology in which elements of traditionally conflicting epistemologies are vindicated within a single epistemological account. We suggest thinking of the epistemologically relevant cognitive processes in terms of the metaphor of an iceberg—the accessible and articulable states that have been the exclusive focus of much epistemology constitute, for reasons that we explain, only a proper subset of epistemologically relevant processing, even as only a part of an iceberg is exposed to view. When one focuses on the interaction of accessible states and articulable information, the structure of epistemic justification looks rather like what has been called structural contextualism (Timmons 1993, Henderson 1994b). It might also be called quasi-foundationalist. Yet, given the sort of creatures we are, in attending to our epistemological tasks we must rely on processing that is sensitive to information that we could not articulate, that is not accessible in the standard internalist sense. When one focuses on the full range of epistemologically important processes, the structure of what makes for justification may be rather more like that envisioned by some coherentists.

I. Epistemological justification and psychological processes.

Much epistemological discussion of justification has been concerned with detailing an abstract normative structure of beliefs—or, commonly, beliefs and other cognitive states such as perceptual or experiential states. The normative structure has been understood in terms of various sorts of evidential support relations involving a more or less extended notion of cognitive content (supposedly fitted to the range of relevant states). However, contemporary epistemologists typically require more for an agent's being justified in holding a particular belief than that the agent merely have states that are evidentially related as described. They commonly recognize that being justified cannot be understood simply in terms of an abstract structure of evidential relations between cognitive states that the agent happens to have. This is to insist that, in an instance of justified believing, the actual mechanisms of belief-formation and belief-maintenance in the agent must themselves evince dependencies that mirror those featured in the abstract evidential structure of contentfully related justificatory states. For an agent to justifiably hold a given belief is a matter of that belief being produced or sustained by the appropriate kind of causal psychological processes, ones that mirror the normative structure of support relations.

The recognition that epistemic justification turns on psychological processes is reflected in commonly acknowledged distinctions between (a) having justification, (b) being justified, and (c) giving one's justification. Each is a matter of an agent's psychological processes, but in very different ways.¹ Giving justification is a matter of articulating, or formulating and expressing, one's justification. It is thus focused exclusively on what can be articulated; and hence is concerned only with internally accessible states. Obviously, giving justification is a psychological matter insofar as articulating or listing such states is (or involves) a psychological process. Having justification for a belief is a matter of merely having psychological states that bear normatively appropriate "content-relations" to the belief in question. Being justified is the matter of central concern in this paper. It is a matter of the agent's belief states arising and being sustained by the right sorts of causal structure. When having justification is understood rather generically, as above, then we can say that having justification is a matter of having states that are appropriately related in "fitting" relations and "supporting" relations—but just having such states, just having justification, is not sufficient for being justified. The more that is needed is a matter of the justified belief being caused in "the right way" by the states that the agent has. (Thus we standardly insist that "the reasons one has" need also to be "the reasons why"

¹ English usage of 'having', 'being', and 'giving' in association with 'reasons' and 'justification' is varied and nuanced. So our remarks are meant to capture some distinctions commonly marked by in these terms, and not to codify usage.

one believes as one does.) Saying something about what can count as the right way is part of the business of this paper. However, this much can be said to begin with: in some important sense, the causal processing involved in being justified must “reflect” the relations implicated in what it is to have justification; it must involve counterfactual dependencies between cognitive states that parallel the contentful support relations.

II. A Proto-theory of Human Belief-Generating and Belief-Sustaining Process as an Epistemological Commonplace.

So the central notion of “being justified” in holding a given belief is commonly recognized as having to do with the sorts of dependency relations exhibited in an agent’s psychological processing. Accordingly, epistemologists commonly have come to dabble in psychology. At the very least, their positions regarding what it is to be justified commonly imply views regarding what human cognition should be like—what it must be like, if we are to have justified beliefs. Insofar as this literature has reflected on apparent cases of justified believing, and has thus naturally tended to presume that we do at least have some justified beliefs, epistemologists have also presupposed a view regarding what human cognition can be, and sometimes is, like. Such psychological commitments are appropriate to epistemological theorizing: in view of the connection between being justified and undergoing the appropriate psychological processes, epistemologists who seek to account for how we can be, and are, justified, are necessarily committal with regard to the psychology of us creatures who are supposedly justified.

Unfortunately, the psychological implications of accounts of justification remain largely inarticulate in the work of many more traditional epistemologists. As a result, the psychological plausibility of those same accounts remains likewise ill-attended-to. While one can tentatively tease out the apparent psychological implications of various discussions, it must be acknowledged that epistemological inquiry could be put on a more straightforward footing, were the implications more forthrightly articulated and addressed by various authors. What, for example, is “the mirroring relation” that must obtain between the contentful relations articulated in the epistemological account and the psychological processes? Because such questions are not articulately addressed, and are at best hinted at, the psychological realism or plausibility of various epistemological accounts is obscure. Because that relationship is quite generally acknowledged to be important, these accounts become problematic. At best, they are inarticulate or incomplete. At worst, inappropriate or unworkable psychological implications are being conveniently underplayed or swept under the rug.

However, the inarticulate character of epistemologists’ psychological dabbling is only a small part of the problem. Much epistemology seems to be wedded to, or at least to tend towards, one substantive general picture of

what the relevant psychological processing is like (or to a set of closely related pictures). That is, there is a proto-theory of the epistemically relevant psychological processing that (more or less implicitly) seems to be taken for granted in much of the discussion of the competing epistemological positions such as foundationalism, coherentism, and contextualism. These assumptions seem to us to have engendered several problematic limitations within epistemological perspectives. Our task in this section is to articulate the proto-theory, in order to challenge it in later sections.

As a first pass, the proto-theory can be formulated as follows:

- PT**
1. What psychological states are relevant to justification—and then, what must be causally operative in the processes that make for our being justified—are accessible to the agent. This is roughly to say that the relevant information possessed by the agent could be accessed in a fashion adequate to reconstructive articulation of those processes in the context of giving justification.
 2. In cases where an agent is justified in holding a given belief, the belief-fixing processes (the generating and sustaining processes) are *occurently isomorphic* to the structure of content relations featured in the agent's having justification. This is to say, the relevant psychological states are *occurrent*, and as such they causally conspire to generate or sustain the belief in a way that mirrors the way in which they contentfully conspire to “support” the belief.

It must be emphasized that PT is intended to encapsulate a central tendency to which epistemologists have been drawn. We will soon wish to distinguish some prominent variations on its themes.

Some readers will think that PT (or at least clause (1)) simply characterizes epistemological internalism. But, insofar as internalism is the view that what makes for justification must be accessible to the agent who is justified in holding a given belief, this is not quite right.

First, in certain respects, PT is *weaker than internalism*. It only insists that the epistemically relevant psychological states—those whose appropriate causal involvement makes for justification—be accessible. This does not of itself entail that what makes for justification be accessible. For example, one might hold PT and still think that what made for the justification of a given belief is its being generated in a reliable fashion. That is, one could be a reliabilist, and thus an externalist of a stripe, and hold PT.² Of course, such a reliabilist would be committed to a view about what human cognitive processes are like, and particularly to what the human processes that tend to

² Alston (1989) provides one example of such a position.

the production of true beliefs are like—for that reliabilist would be committed to the relevant processes employing only accessible/articulable states.

Second, while internalism may incline one to holding PT, it is not just internalism *per se* that would do so, but internalism plus the recognition that being justified involves being produced or sustained by the appropriate cognitive processes. Together, these entail that the relevant processes be comprised of accessible states (and their causal interactions). Because it has become quite common to recognize the connection between justification and causality, and because internalist sympathies are widespread, one would expect that PT is widely held.

One helpful way of thinking about PT, and for recognizing its pervasiveness, is to notice its kinship with a common epistemological approach to thinking about, or analyzing, what it is to be justified: the context of giving justification has commonly served as a model for the kinds of causal relations that are constitutive of justification. But, to the extent that epistemologists have relied on what we would say in justification-giving, their approach presupposes that the sorts of states that can be articulated in justification, and the sorts of relations between states that can there be pointed to, exhaust what is important for being justified. We are not saying that having-justification or being-justified has been conflated with giving-justification—for many epistemologists who take the approach in question would distinguish these as done above. It is to say that the sorts of things that could be adduced in “the context of justification” (that is, the context of justification-giving) are thought to exhaust what is relevant to having and being justified—only that being justified is a matter of just these things playing appropriate causal roles, being appropriately interdependent. However, when giving-justification is used as a model for being-justified in this way, we get an important constraint on our thinking about being justified: while being justified is still recognized as not requiring articulation, it can only be seen as turning on what is articulable.

To avoid a misunderstanding, we should note that, while PT holds that the epistemically relevant psychological processes of belief-generation and belief-maintenance mirror what would figure in a full articulate giving of justification, this is not to say that any agent who is justified can actually fully articulate that justification. That all items in it are accessible, and that proper parts of it are thus articulable, does not entail that the agent can fully articulate the full justification. (Although this latter claim may become tempting as a result of holding PT, it is not entailed by PT.)

It is worthwhile to reflect on (and decompose) this use of justification-giving as a model of being justified. This will help us to appreciate how it can seem so natural, and thus will account for the naturalness of associated epistemological points. To this end, we set out the giving-as-model approach (or **GM**, for future reference):

1. Justification-giving is thought to reflect the (normative) structure of justification-having.
2. Since (in the case of a justified belief) the causal structure of belief-fixation is isomorphic to the normative justification had for the belief, justification giving (in giving access as above to justification-having) reflects the occurrent causal structure of the (justifying) belief-fixation process. (Of course, since justification-giving may be incomplete or elliptical, it may incompletely reflect the causal structure constituting being-justified.)
3. Although someone whose beliefs *are* justified need not necessarily be able to give the justification in its entirety (since the full justification might be very complex and cumbersome, and articulating it might then become cumulatively intractable), still it has seemed natural to assume that the full justification would amount to, or turn on, “more of the same” sort of states that are accessible in giving-justification. That is, it is commonly assumed that inarticulability of the full justification, where it appears, is a result of the difficulty of keeping track of all the information and states that are relevant; thus it is not a matter of there being inaccessible information and states entering into the justification-constituting processes.
4. If we are to be able to give the justification we have (and thus reflect the states that make for justification) we must not simply have access to those states, but also must be at least “sensitive to” the contentful relations between them that make for their counting as reasons. That is, while we need not be able to articulate what contentful relations make for their serving as reasons, we must at least be sensitive enough to these matters to identify the relevant states as reasons.

So:

5. (Clause (1) of PT.) The states that figure in being-justified are limited to states that are accessible (This results from 1, 2, and 3.)

So:

6. The causal structure of the justification-constituting belief-fixating processes is consciously accessible (if not readily fully articulable). (This results from 1–4.)

Thus, in sum (drawing on all the above):

7. Potentially articulable justification *is* the justification one has, and one's belief-fixation process (that makes for being justified) is occurrently isomorphic to this potentially articulable justification (that one has).

Again, we should emphasize that GM is not intended as a putative argument for PT. Rather, it is intended as a formulation of a very common way of thinking of giving justification as providing a model for what it is to be justified. As represented above, we find that PT arises naturally in the context of this model.

While we can point to examples where PT (and GM) seems to be at work, documenting its working in a decisive and fully representative fashion would be exhausting at best. After all, the suggestion here is that PT is a pervasive *assumption* that is honored in a range of epistemologies. Documenting its working in any one case requires careful attention to nuances of what is said and not said in the relevant texts. Such fine-grained textual analysis for a representative sample of authors is clearly beyond what can be undertaken here. In place of such demonstration we rely mainly on setting out PT and situating it with reference to a pervasive way of thinking about justification. We hope that doing so has already made fairly evident the naturalness of PT (and GM) to those accustomed to thinking along lines that are standard in contemporary epistemology. We can, however, mention several examples where PT seems operative. These examples are provided by way of illustration; they are not intended as generally decisive evidence of the pervasiveness of PT.

Paul Moser (1989) provides a fairly clear example of PT. Much of Moser's discussion of "justification" has to do with the concept of "epistemic justification" as the "evidence" one has for a proposition, with one's "justification" as "the truth-indicators" one has (1989, p. 35). As Moser recognizes, a different notion of justification is in play when we talk of someone "being justified in believing" a given proposition, and it is this second matter that is directly required for a belief's being a bit of knowledge (pp. 45, 151–58). (This second notion is our central concern here.) To begin with, to be justified in believing a given proposition, one is required to have justification-as-evidence for that proposition. And Moser's analysis of justification-as-evidence is characterized by a particularly robust internalism—one that requires that the states that make for justification be more than simply accessible. On Moser's analysis, if one has justification for P at t, the justifying states must either be occurrent states of awareness at t, or dispositions resulting from the earlier operation of such occurrent states. His analysis of what it is to justifiably believe something then comes to inherit this robust internalism. This much is clearly stronger than, and implies, PT(1). It suggests something of the "occurrent isomorphism" of PT(2)—a

suggestion that is borne out by Moser's insistence that the causal dimensions of "being justified in believing" turn on a "causal significance" of an *awareness* of an "association relation" between the proposition believed and its "justifying evidence." (pp. 156–57), where that association relation is itself established by a causal interaction of awareness-states (pp. 141–42). Being justified in believing thus turns on a causal interaction of *awareness* states—paradigmatic occurrent states—a causal interaction that is isomorphic to the content relations making for justification as evidence. (To this, Moser adds the flourish that the causal relations are themselves mediated by further contentful occurrent states—awarenesses of the evidential relations between the agent's belief and the evidence possessed. Here, it seems, Moser views cognitive processing down to the most detailed level relevant to epistemic justification as a structure of occurrent contentful states; and again, it is required to be isomorphic to the normative structure of having justification.) In all this, the central model informing Moser's analysis seems to be that of occurrent states interacting causally in a manner mirroring the evidential relations described in his analysis of justification-as-evidence. This amounts to a particularly clear example of PT(2).

Robert Audi's (1993) work is characterized by a laudable concern for a psychologically realistic epistemology, with the result that psychological presumptions are commonly relatively close to the surface in his discussions. Thus, his dabbling is more articulate than most epistemologists'. As a foundationalist, Audi's concern and focus is commonly on the generation of perceptually induced beliefs, and it is here that his discussion seems most psychologically committal. His "modest foundationalism" posits the following sort of sequence (a sort that is fairly typical of what foundationalists commonly take as epistemically relevant processing). First a *perceptual state* occurs; using an example of Audi's (p. 130): working in my study in the evening, I "see a headlight beam cross my window." This immediately induces in me a *perceptual belief* "that a car's light is moving out there." This perceptually induced belief is taken to be epistemologically basic and to arise "directly" out of the perceptual state. This is to say that, from the point of view of epistemologically relevant psychological processes, the link between the perceptual states and the perceptual belief is a direct causal one. Finally, on the basis of this belief and other beliefs that I antecedently possess, I may inferentially generate yet further beliefs—*non-direct and non-basic beliefs*—for example, "that someone has entered my driveway." The generation of these further beliefs need not be simple—they can be the result of the interaction of many informational states (including beliefs) in addition to the perceptual beliefs in their etiology. To relate another example, I perceive a "distinctive patter outside my open window," immediately generating the perceptual belief "that it is raining," which may inferentially spawn the belief that the seats on my convertible Maserati are getting wet. The force

of what is and is not mentioned in such examples is to delimit what are the epistemically relevant states and cognitive processing to perceptual states, occurrent beliefs, and their causal interactions. (Nonoccurrent beliefs are apparently justified insofar as they are the stored results of, or the sustained results of, processing involving such elements.) The force of calling the perceptual belief “direct” is not *simply* to insist that other beliefs do not feature in its generation, but also to insist that, with respect to what is epistemically important, it comes immediately from the perception (without epistemically relevant mediating states).³

For our purposes here, it is important to note how the epistemological action at the junctures that most concern Audi has been limited to states that are paradigmatically occurrent—perceptions and occurrent beliefs. Yet when we reflect on these examples, we readily recognize the role for information that the system has, information that has been learned and somehow brought to bear—but information that is not carried in paradigmatically occurrent states. (We will later say more about how such information is involved and why it is causally relevant.) Depending on what exactly we take to be the perceptual state, there will be a great deal of information somehow employed by the agent in either getting to the perceptual state itself or in getting from it to the perceptual belief. When the perceptual state is taken somewhat sparsely—say as “a distinctive patter”—we are less likely to notice the role of information antecedently possessed by the system. (But, we wonder, does “distinctive” simply mean discriminable—as in a noticeably distinct or discriminable tone—or does it mean distinctive or characteristic of rain-drop-impact sounds? The latter is not a “sparse” perceptual state.) When the perceptual state is taken as more informative—say as perceiving “a headlight beam across my window”—we are more likely to notice the role of information antecedently possessed by the system and employed in their generation.

³ The present point receives further support when reflecting on Audi’s discussion of the “epistemic dependence” of beliefs on other beliefs (or other states). He relies on a distinction between positive and negative dependence— x is positively dependent on y only if x is the result of the presence of y ; x ’s being negatively dependent on y is a matter of the realization of x turning on the absence of y . Audi seems to see epistemic justification as having dependency patterns that reflect both positive and negative psychological dependencies between states. Thus, he insists that a perceptual belief has a measure of (defeasible) epistemic justification deriving from the perceptual state from which it directly springs (and apparently only from that state). This positive epistemic dependency (of justification) results from the belief’s being positively psychologically dependent on the perceptual state. This justification can be defeated by incoherence with other beliefs, but he insists, this makes the justification negatively dependent on incoherence—it turns on the absence of incoherence. Importantly, here, beliefs and other information-carrying psychological states are dropped out of the epistemic picture when considering the generation of perceptual states and perceptual beliefs. Perceptual beliefs are said not to be positively dependent on such states. At least at the level relevant to epistemology, the perceptual beliefs arise directly out of perceptual states, which, from an epistemic point of view, are the beginning of the story (they just happen).

Of course, moving from a sparse perceptual state such as the perception of a discriminable sound to the belief that there is rain falling, or getting from the perception of a momentary light patch on my shades to the belief that there is a moving auto in the vicinity, employs quite a bit of learned information. (That is why we manage it, our dogs may manage it, but small infants do not.) What is important is how such information gets conveniently left out of the story Audi provides. It does not feature in the sort of internalist reconstruction of evidential support relations that Audi is ready to provide; it gets omitted from the psychological story being told. What we then get is a story in which the psychological goings-on are occurrently isomorphic with the normative account of relations between accessible (or, indeed, accessed) cognitive states.

Laurence Bonjour's coherentism (1985) provides an interestingly different example of the way PT has informed epistemology. Of course, Bonjour is an internalist. And, while internalism is distinguishable from PT(1), it certainly commits one to PT(1). What is more difficult to see, and more interesting when seen, is how Bonjour honors PT(2). Bonjour is clearly committed to a kind of normative psychology—to a normative account of the sorts of cognitive processes that make for justification. On the basis of this normative psychology he is led to despair of our being actually justified, as opposed to our being merely potentially justified—in effect, he despairs of our actualizing, or instantiating, the sort of processing called for. It is important to see that something like PT(2) is implicated in Bonjour's very demanding normative psychology.

Reflecting on the various sorts of considerations that can come up in the context of justification, Bonjour is led to his holistic-coherentistic understanding of epistemic support relations. On his view, the justification of any one belief turns on other beliefs to such a degree that, ultimately, the justification of one's entire belief set is primary, while local justifications for individual beliefs are secondary and derivative. Now, what sort of cognitive processing is required for an agent's cognitive processing to appropriately parallel or mirror the contentful structure of coherentist support relations? Bonjour rightly insists that merely having a globally coherent belief set is not enough for an agent's being justified. Somehow, the agent's cognitive processing must produce or maintain beliefs in a way that is conditioned by coherence relations, in a way that mirrors those relations. One might say, the agent's processing must "take account of" global coherence. But just what do such requirements amount to for Bonjour? One very natural assumption within the context of recent epistemology is already witnessed in Moser and Audi: that cognitive processing would need to mirror the contentful structure of epistemic support relations by way of a causally interactive structure of occurrent states, states that jointly exhaust whatever content contributes to justification. Bonjour adheres to this "natural" line. Since the coherentist

evidential structure involves one's global belief-set, Bonjour presumes that an epistemic agent would need to have cognitive processes that somehow *survey*, that is, *explicitly look over*, his or her own full belief-set, noting coherence or the lack thereof, and then update the agent's beliefs appropriately. In his words, he supposes that for cognitive processes to mirror global-coherentist support relations in a fully appropriate, "actual justification"-conferring way, they must work by making "fully explicit" the agent's "entire belief system" and its coherence relations (1985, pp. 151–53). This demand clearly reflects PT(2), as it requires that the full-belief set (being normatively relevant) becomes, in some appropriately encompassing fashion, occurrent and thereby causally operative. It is then upon recognizing that actual believers cannot so survey their entire belief-set that Bonjour is led to conclude that actual believers are, at best, only potentially rather than actually justified in their beliefs. That is, while one's beliefs might happen to be so contentfully interrelated as to conform even ideally with the coherentist standards Bonjour codifies, because agents cannot make all their beliefs explicit (and thus occurrent) as a whole, the agent's processing cannot access this fact of coherence and thereby causally induce the agent to continue to so believe as a result. One might be potentially justified, in that, were such processing possible, one would be justified by one's fine belief-set. But, for Bonjour, one cannot be actually justified thereby.

In what follows, we set out reasons for thinking that PT is problematic and that repudiating it will advance epistemology. But first, in order to get a fair understanding of the situation in contemporary epistemology, it will be useful to keep in mind several close relatives of PT. Recall that PT characterizes a central tendency informing epistemological theorizing. However, different philosophers will tend to hold onto and insist on respecting different components of PT to different degrees. In particular, some epistemologists may feel that the first clause of PT should be a strong constraint on their thought, while not feeling a similar pull for PT(2). Others may embrace the second clause, while entertaining alternatives to the first. So, it will be worthwhile to explore plausible weaker versions of the proto-theory common to mainstream epistemological thinking, versions that arise in the ways just indicated.

A wary proponent of something like PT might wonder whether important states have been missed—states that are not necessarily consciously accessible, but otherwise operate causally in the way envisioned by PT. That is, without the assumption that all the epistemically relevant states are consciously accessible, one might still hold to the second clause of PT. An epistemologist doing so would be holding to the same basic view about what it is for an agent's causal psychological processing to appropriately "mirror the normative structure of support relations"—mirroring is still a matter of occurrent isomorphism, on this line. Call this variant **PT-occurrent iso-**

morphism, or PT-oi. Like PT itself, PT-oi insists that the epistemically relevant psychological states play their causal role as occurrent states. However, whereas with PT itself one looks for an isomorphism between **articulable** psychological processing and epistemic support-relations, with PT-oi one allows for the possibility that there may be epistemically relevant processing going on that is occurrent but not articulable.

Here the comparison with classical artificial intelligence seems apt. In classical AI accounts, such as the account of visual processing in Marr (1982), much information processing operates on representations that the agent is not conscious of and could not raise to consciousness. (In Marr's account, this is true for much of what happens in the early stages of vision.)⁴ It is plausible that such information can, and often does, constitute part of the evidential basis for certain beliefs—e.g., visually induced perceptual beliefs.⁵ But although the information can be unconscious, according to classical AI, still such representations can only have a causal role insofar as they are tokened, and thus occurrent, in the relevant processor. So AI work like Marr's suggests a model that extends the picture of occurrent, tokened, interacting states to apply it to the sub-basement of epistemologically relevant processing. The further occurrent states are "more of the same" in how they affect belief-generation and belief maintenance, except for being sometimes unconscious and not articulable. On this line of thought, justification-giving remains an appropriate model for the psychological processes that make for being-justified, even though the inaccessibility of certain epistemically relevant processes significantly expands the extent to which one can be justified in holding a belief without having the capacity to **fully articulate** one's justification. Pollock (1986) seems to point in the direction of the PT-oi position.

⁴ Marr's computational theory of vision treats the human visual system as a device that receives pairs of retinal images as inputs, and produces as output a visual representation of the three-dimensional distal environment. The process breaks up into several sequential stages. Initially the system constructs a representation of certain changes and structures in the retinal image. Then the system uses this representation to construct a "primal sketch" of the surface outside the viewer. Next, on the basis of the primal sketch, the system constructs a "2.5D sketch" of the viewed surface, explicitly representing information about contour and rough depth in a viewer-centered coordinate frame. Finally, on the basis of the 2.5D sketch the system constructs a "3D model representation" of shapes and their three-dimensional spatial orientations, relative to an object-centered coordinate frame. Stages in this process need not be accessible to conscious awareness, and some stages—the primal sketch, for instance—most probably are not. (Here and below, we draw upon the nice summary of Marr's theory in Segal 1989.)

⁵ The passage from each stage to the next is regarded as an inference, yielding an overall inference comprising a chain of component inferences. Marr writes, "the true heart of visual perception is the inference from the structure of an image about the structure of the real world outside" (p. 68). In terms of content, each successive representation in the sequence is indeed evidentially supported by the representations preceding it.

On the other hand, the sorts of concerns for accessibility reflected in the first clause of PT are clearly widespread. Many epistemologists might wish to adhere to these while thinking that the second clause is overly constraining. Such epistemologists might find inspiration for developing a particularly natural variant on PT in the psychological notion of “procedural knowledge.” The basic psychological story here is as follows. In many contexts, how we approach a cognitive task evolves as we acquire expertise. Initially, we might deal with a cognitive task by consciously representing all pertinent, task-specific, information and employing this information in accordance with fairly general-purpose rules or procedures. In the psychological literature, information that is thus occurrently and consciously represented is often called “declarative”; thus, the idea is that novices often attack cognitive tasks by means of general-purpose rules or procedures that operate on declarative representations of task-specific information. However, with experience dealing with a particular class of common cases falling within the application of a general procedure, we might develop a special-purpose procedure for treating these cases, with some of the information appropriate to the special class of cases incorporated within, or accommodated within, the procedure itself. In employing such a procedure, some of the computational tasks involved in the general procedure come already done when the procedure is implemented. As a result, these subtasks do not need to be redone as the special-purpose procedure is put into play. Accordingly, not all the information on which the general-purpose procedure operated needs to be explicitly, occurrently represented; i.e., it need not all be declarative information. Instead, some of it can be implicitly taken account of by executing the special-purpose procedure. In the psychological literature, information that is thus implicitly accommodated is often called “procedural,” in contrast to declarative information.

Thinking in terms of such proceduralization could lead one to loosen up in connection with PT(2) as follows. As a thought experiment, suppose that we start with an epistemic system that conforms to epistemic standards fully in keeping with PT. With respect to PT(2) in particular, we would then have a system that employs as occurrent representational or informational states all the sorts of states that are epistemically relevant. The epistemic system manages to operate on those states in the content-appropriate ways specified “without skipping a beat.” So far, thinking along the lines of PT, we should be very happy with this system. However, the question now arises whether something less than, but related to, this full conformity with PT would also be unexceptionable. In particular, it seems reasonable to allow for certain forms of learned development within that system in the procedures it implements as well as the informational base on which it operates. One sort of development that seems reasonable to allow is proceduralization. That is, it seems reasonable (indeed epistemically desirable) that the system would come to compact its epistemic computational tasks in the way described above,

where feasible. This would be a matter of acquiring an epistemic expertise in a class of cases. But, when this is done, the system ceases to conform strictly to PT(2)—the states on which it operates and the transitions on those states no longer will be fully occurrently isomorphic to the normatively appropriate content relations under which, *ex hypothesi*, we started. In its present causal transformations involving occurrent states, the system is now taking short-cuts. Still, the results seem in keeping with the spirit, if not the letter, of PT. In keeping with the spirit of PT, with an emphasis on PT(1), what is important is that this system, which has come to deal so compactly with certain classes of cases, *could* have performed these epistemic tasks “longhand” in full conformity with PT had this been called for somehow, and that the information relied on in its processing is accessible.⁶

We can then christen a slightly weakened variant on PT—call it **PT-proceduralization** (or PT-p) which insists on PT(1), and insists that the system could have fully complied with PT(2), while allowing for proceduralization so that in a given episode the processing in a system need not be fully occurrently isomorphic with the content relations specified in a set of epistemic ideals.^{7, 8}

⁶ Information that is procedural rather than occurrent during processing is not conscious during processing, since all conscious mental states are occurrent states. Nonetheless, procedural information can be consciously accessible anyway. Roughly, such information is consciously accessible just in case it reliably can be recovered in conscious, occurrent, form under the relevant kinds of post-hoc elicitation circumstances—for instance, in contexts where one is seeking to articulate one’s justification for a given belief.

⁷ The notion of procedural information can be understood in various ways, some broader in scope than others. In the narrowest sense, procedural information is information that formerly was employed in declarative form within the given cognitive system, and now is accommodated implicitly in the system by virtue of proceduralization that occurred during learning; a “longhand” procedure has been replaced by a corresponding “shorthand” procedure. In a somewhat broader sense, procedural information is information that is accommodated implicitly and could instead have been accommodated explicitly in a corresponding longhand procedure—*whether or not* the cognitive system has ever actually employed the longhand procedure itself. This broader construal is the one we have built into our characterization of PT-p. (One can envision an even broader usage; see note 9 below.)

⁸ As noted above, Moser seems to hold PT. However, it is easy enough to envision a way of getting something like PT-p out of his view with a somewhat legislative reading. Recall that Moser’s account of justification-as-evidence employs the notion of a disposition in a central fashion (as much of our evidence will typically be things like beliefs, which are understood as dispositional states resulting from episodes of assenting to a proposition). If epistemically relevant dispositions are limited to such evidential states, and this is certainly the paradigmatic fashion in which they feature in Moser’s discussion, then Moser seems to hold to PT. The reason is that such dispositions function as stored bits of declarative knowledge. The notion of proceduralization involves a distinction between procedural and declarative knowledge—as proceduralization involves a modification in procedural knowledge that lessens the role of declarative knowledge in treating classes of cases. If, however, Moser were to allow for acquired dispositions that involve not simply the declarative knowledge he focuses on, but also procedural knowledge, then

PT-p is more thoroughly in keeping with the spirit of PT than is PT-oi. For, PT-p retains the accessibility requirement PT(1), while backing away from the full-fledged occurrent isomorphism requirement PT(2) in a way that still respects the general spirit, though not the letter, of occurrent isomorphism. But PT-oi, in eschewing the accessibility requirement, thereby acknowledges that consciously accessible states may well be just the tip of a larger iceberg of epistemically related psychological states—a theme that will loom very large later in this paper. In short, the requirements imposed by PT-p constitute only a modest weakening of those imposed by PT, whereas the requirements imposed by PT-oi constitute a more substantial weakening.

This being so, anyone who is prepared to fall back from PT to PT-oi should also be prepared to fall back a bit further, to a proto-theory that allows proceduralized shortcut-alternatives for processing that the system can perform in a way that is fully occurrently isomorphic to the content relations specified in a set of epistemic ideals. Call this view (**PT-occurrent isomorphism**)-**proceduralization** (or PT-oi/p). Just as PT-p cleaves to the spirit, though not the letter, of PT itself, so likewise PT-oi/p cleaves to the spirit, though not the letter, of PT-oi.^{9, 10}

Moser could be seen as embracing the more moderate PT-p. One can point to aspects of Moser's (1989) that suggest an openness to such a position, but whether he would allow for PT-p in place of PT is probably indeterminate from that text.

⁹ Furthermore, it also would be in the spirit of PT-oi to allow proceduralized shortcut-alternatives even under a more liberal construal of the notion of 'procedural information' than we have adopted in the text. This construal would allow the possibility, for instance, that (i) some epistemically relevant information is accommodated implicitly, rather than being explicitly represented during processing, (ii) an algorithm that procedurally employs this information is hardwired into the system, and (iii) the system lacks the capacity to run a corresponding longhand algorithm employing an explicit representation of the relevant information.

Marr's theory of early vision suggests just such a possibility. Marr emphasizes that the visual system's inferences from its retinal stimulations to the nature of the distal causes of those stimulations depend on certain assumptions. One such assumption is that "the visible world can be regarded as being composed of smooth surfaces whose spatial structure may be elaborate" (Marr, p. 44); another is that "the items generated on a given scale tend to be more similar to one another in their size, local contrast, colour, and spatial organization than to other items on that surface" (Marr, p. 46). Given Marr's general account of the input-output function computed during early vision, it is entirely possible—even likely, for reasons of speed and resource efficiency—that the algorithm that computes this function accommodates such assumptions implicitly; it is also very likely that the algorithm is hardwired. So presumably this procedural algorithm constitutes the *only* way that the visual system can compute the input-output function; the system does not, and cannot, computationally generate a 3D visual representation by executing an alternative algorithm that employs explicit, occurrent, representations of the relevant background assumptions.

In order to simplify subsequent discussion, in the body of this paper we will largely ignore the liberalized construal of 'procedural information', and the corresponding liberalized version of PT-oi/p, described in the present note; but the reader should keep them in mind. See also note 13 below.

While PT seems to be the proto-theory of choice, or at least of ready assumption among contemporary epistemologists, PT-oi, PT-p, and PT-oi/p provide weaker variants that might be attractive to the wary once psychological commitments begin to be articulated as done here. However, our objections to PT are also telling against these kindred alternatives. The psychological understanding that inspires our misgivings are described in the next section. With this rather different picture of human belief generating and sustaining processes in hand, the final section will sketch certain transformations in our epistemology that seem indicated.

III. An Importantly Different Picture of Our Belief-Generating and Belief-Sustaining Processes.¹¹

Certain problems and developments in cognitive science provide reason to believe that the human cognitive processes that typically produce and sustain beliefs diverge in important ways from PT and its kin. The alternative picture emerges in the wake of a family of recurrent, 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 we describe the alternative conception suggested by the 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. Finally, we underscore the ways that this alternative approach calls into question PT and kin.

III.1. The Moral of the Frame Problem: Belief Fixation as Essentially Morphological.

In the closing pages of Fodor (1983), it is argued that certain problems in classical cognitive science look to be in-principle problems, and hence that the prospects for understanding human cognitive processes like belief fixation within the framework of classical cognitive science are very bleak indeed. These problems continue to plague the computational approach to the mind, and suggest the need for a radically different approach.

¹⁰ The most straightforward understanding of Pollock’s (1986) sees him as embracing PT-oi as described earlier, and some of Pollock’s discussion suggests elements of PT-p, as he makes much of the notion of the internalization of procedural knowledge (pp. 129–32). However, the importance of PT-p elements within Pollock’s position is rather difficult to gauge, as he seems sometimes to suggest an individuation of norms—and thus cognitive process types—that is too course-grained to allow a precise fix on his position regarding these matters. In any case, the PT-oi elements in this thought, and his appreciation for proceduralization, might lead one to read him as an adherent of PT-oi/p.

¹¹ The discussion in this section draws freely on material in Horgan and Tienson (1994, 1996) and Horgan (1997a). Some points we will make about psychological processing, and their potential import for epistemology, are discussed by Herman Stark (1993, 1994, 1995).

The main claim of Fodor (1983) 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. There is good evidence, Fodor argues, that human input systems, including those that mediate speech perception, exhibit modularity. Where classicism has gotten somewhere, he says, is in understanding such modular subsystems, which by their nature delimit the class of relevant information.

Classicism has made very little progress in understanding *central* processes, however. Belief fixation—the generation of new beliefs on the basis of current input together with other beliefs—is a paradigmatic example. Updating of the overall belief system in light of currently available new information is a closely related example. Fodor argues convincingly that 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.

Fodor articulates these considerations in terms of the analogy between belief fixation in human cognition and scientific confirmation. Concerning central cognitive processes like belief fixation, he says, “it seems reasonable enough that something can be inferred about them from what we know about *explicit* processes of nondemonstrative inference—viz., what we know about empirical inference in science” (p. 104). Scientific confirmation, “the nondemonstrative fixation of belief in science,” has two crucial features. It is (in Fodor’s terminology) *isotropic* and *Quineian*. He says:

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.... (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 (p. 107).

Isotropy brings in the whole of current theory: any bit of actual or potential information from any portion of the 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” (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 a matter of inductive inference from the information provided by input systems and the information in memory, 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 respects, 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) (pp. 112–13).

When one considers the sorry history of attempts in philosophy of science to construct a theory of confirmation, the prospects for understanding central processing within the classical computational paradigm look very discouraging indeed:

Consider...the situation in the philosophy of science, where we can see the issues about fixation of belief writ large. Here an interesting contrast is between deductive logic—the history of which is, surely, one of the great success stories of human history—and confirmation theory which, by fairly general consensus, is a field that mostly does not exist. My point is that this asymmetry, too, is likely no accident. Deductive logic is the theory of validity, and validity is a *local* property of sentences. Roughly, the idea is that the validity of a sentence is determined given a specification of its logical form, and the logical form of a sentence is determined given a specification of its vocabulary and syntax. In this respect, the level of validity contrasts starkly with the level of confirmation, since the latter...is highly sensitive to global properties of belief systems.... The problem in both cases is to get the structure of the entire belief system to bear on individual occasions of belief fixation. We have, to put it bluntly, no computational formalisms that show us how to do this, and we have no idea how such formalisms might be developed.... In this respect, cognitive science hasn't even *started*; we are literally no farther advanced than we were in the darkest days of behaviorism.... If someone—a Dreyfus, for example—were to ask us why we should even suppose that the digital computer is a plausible mechanism for the simulation of global cognitive processes, the answering silence would be deafening (pp. 128–29).

These are wise words. Let us underscore their wisdom by dwelling just a bit on the depth of, and the apparently in-principle nature of, the difficulties encountered by attempts to model global cognitive processes computationally. Take, first, the Quineian aspect of belief systems. Simplicity and conservatism are properties of (or relations between) belief systems that depend upon the formal, semantic, and evidential relations among *all* of the beliefs in the system(s). A computational system would have to somehow survey the entire stock of beliefs, in a manner that tracks all the multifarious

interconnections among the beliefs, and somehow derive a *measure* of net overall simplicity and net overall conservatism from these local features. As Fodor said, “We have...no computational formalisms that show us how to do this, and...no idea how such formalisms might be developed....” (Recall Bonjour’s problems.)

In addition, when new information comes in from the input modules, the central system would have to find, from among the vastly many competing, incompatible ways of revising the whole system to accommodate this new information, a mode of revision that maintains overall simplicity and conservatism better than most of the others. All this would have to be done via *tractable* computation, executable quite rapidly. Not only do we have no computational formalisms that show us how to do this; 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.

Now consider isotropy—the potential relevance of anything to anything. Fodor’s *definition* of isotropy suggests that the problem is that there is *too much* that may need to be considered, and that there is no known manageable way to bring what is needed into consideration. This is typical of the way frame type problems are perceived. But there is actually an even deeper problem concerning relevance, a problem suggested by Fodor’s talk of local versus global properties and by his comparison with the historical failure of confirmation theory: viz., that *relevance itself is a Quineian as well as an isotropic property*. How is (e.g., confirmatory) relevance to be determined computationally? Certainly, some observations are necessary to establish general beliefs like the belief that almost all crows are black. But the belief that many black crows have been observed and that no non-black ones have been observed does not, by itself, support an inference concerning crows in general or concerning any particular crow. For a cognitive system to make such an inference, the predicate ‘black’ must be *projectable* for the system relative to the predicate ‘crow’. But projectability of predicates is a global feature of a cognitive system. The projectability of one predicate relative to another for a cognitive system depends upon features of a great many other relevantly (!) similar predicates within the cognitive system. Thus, belief fixation depends upon projectability, and in the real world, projectability of any single predicate for a cognitive system depends upon relations among large numbers of cognitive states of that system.

Fodor, in the passages bemoaning the lack of an available computational formalism, is telling us that 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?

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 somehow, 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—the isotropic and Quineian aspects—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), the 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), 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. They are not accomplished by computationally manipulating explicit, occurrent, representations of all relevant information. Nor, therefore, 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.

III.2. Dynamical Cognition and 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” in a smooth and direct way, so that an adequate account of real human cognition would turn out to be just a straightforward extension of the given model. 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 a positive 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. This nonclassical framework for cognitive science is described at length in Horgan and Tienson (1996). 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.

A connectionist system, or neural network, is a structure of simple neuron-like processors called nodes or units. Each node has directed connections to other nodes, so that the nodes send and receive excitatory and inhibitory signals to and from one another. The total input to a node determines its state of activation. When a node is on, it sends out signals to the nodes to which it has output connections, with the intensity of a signal depending upon both (i) the activation level of the sending node and (ii) the strength or “weight” of the connection between it and the receiving node. Typically at each moment during processing, many nodes are simultaneously sending signals to others.

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,” as connectionists like to say—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. Because learning involves the system’s undergoing weight changes that render its representation-forming dispositions appropriate to the content of what is learned, learning is the acquisition, “in the weights,” of new morphological content.

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 hypothetical. The set of all possible states of the physical system—so charac-

terized—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. 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 training procedures employed in connectionist modeling bring about a controlled co-evolution of these two factors.) Once a network has been successfully “trained up” to perform an information-processing task, 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 they call the *dynamical cognition* framework (the DC framework). This alternative approach offers an answer, in principle, to the question, “How could the holistic, Quineian/isotropic, aspects of cognitive processes like belief fixation be accommodated automatically and morphologically, rather than by computationally manipulating explicit representations of all relevant information?” The answer is this:

In principle, Quineian/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 Quineian/isotropic information.

As we remarked above, we know of no extant connectionist models that come anywhere close to dealing with frame-type problems of the kind routinely encountered in real human cognition, and we know of no good reason to think that current models are likely to “scale up” in any straightforward way. Nevertheless, the DC framework offers a general, and not implausible, answer to the extremely daunting question of how the holistic aspects of belief fixation could be handled in an essentially morphological way. The classical computational approach is no longer “the only game in town” in cognitive science; and, whereas frame-type problems look quite intractable from within the classical framework, the DC framework suggests the outlines of a solution.¹²

¹² What about the well known argument in Fodor and Pylyshyn (1988), to the effect that connectionist systems must either (i) exhibit the fatal inadequacies of traditional associationism (by eschewing language-like mental representations that are subjected to structure-sensitive processing), or (ii) merely constitute an implementation of the classical computational approach? The DC framework, as described by Horgan and Tienson, avoids both horns of this dilemma. On one hand, the DC framework endorses a (nonclassical) “language of thought,” with structure-sensitive processing of language-like mental representations. But on the other hand, it deviates from classicism by denying that cognitive processing conforms to *programmable rules* for manipulating these representations on the basis of their syntactic structure. Cognitive state-transitions can fail to conform to such representation-level programmable rules even if they are realized by state-transitions that do conform to programmable rules at sub-cognitive, implementational, levels of description (as do the sub-cognitive state-transitions of the connectionist systems that are simulated on standard computers). Basically, this is because the realization relation from cognitive to mathematical states can be so complex and subtle that the property of tractable computability, for state transitions, fails to “transfer upward” from sub-cognitive levels to the cognitive level of description. For further discussion of these matters, see Horgan and Tienson (1996, pp. 63–67) and Horgan (1997b).

III.3. PT Rejected.

As we argued in section III.1, the apparent moral of frame-type problems in computational cognitive science is that the holistic, Quineian/isotropic, aspects of cognitive processing are not subserved by processes that update beliefs (and other informational states) in a manner that is occurrently isomorphic to holistic normative-justificatory relations among these items. This kind of occurrent isomorphism between normative-justificatory relations and psychological processing is just not possible. Rather, the holistic aspects of cognitive tasks like belief fixation are primarily subserved morphologically: cognitive transitions are automatically appropriate to large amounts of implicit information and to holistic normative-justificatory relations involving that information. The holistic aspects of cognitive processing *must* be subserved morphologically; this is an essential aspect of Nature's "design solution" to the problem of avoiding frame-like breakdowns in human cognition, and to the closely related problem of managing significant inductive reasoning without intractability. The DC framework for cognitive science explains, in broad terms, how this essentially morphological mode of belief-fixation is possible.

This picture of cognitive processing is importantly different from the picture provided by PT. PT requires, as a condition of a belief's being justified, that it be causally generated and sustained by processes that are occurrently isomorphic to the normative-justificatory support relations that this belief bears to other intentional states of the cognitive system, and that these processes be consciously accessible. But according to the alternative picture, certain crucially important, holistic, normative-justificatory relations among the cognitive system's intentional states are not—and cannot be—subserved by occurrently isomorphic psychological processes; instead, they are, and must be, subserved by processes in which much of the epistemically relevant information is accommodated morphologically. Thus PT(2) is violated.

This kind of cognitive processing is also at odds with the psychological picture provided by each of the fallback variants of PT, viz., PT-oi, PT-p, and PT-oi/p. Since PT-oi retains condition PT(2), the remarks in the preceding paragraph apply directly, *mutatis mutandis*. PT-p and PT-oi/p both impose the following requirement as a condition for a belief's being justified: if the cognitive processing that generates or sustains the belief relies upon morphological content, then this processing must be a proceduralized shortcut version of a form of processing that (1) could have occurred in the cognitive system, and (2) would be occurrently isomorphic to normative-justificatory

support relations. This condition too is violated, if indeed holistic evidential factors can only be accommodated morphologically, not occurrently.¹³

What about conscious accessibility of epistemically relevant information and of evidential support-relations, as required by PT and by PT-p (although not by PT-oi or PT-oi/p)? In characterizing PT-p in section II, we pointed out that even information that is accommodated procedurally during belief-fixation sometimes can be accessible anyway: sometimes the agent has the capacity to represent the information occurrently and consciously after the fact, in subsequent cognitive processing aimed at explicitly articulating the justification for the belief. The same point applies to information that is accommodated morphologically during belief-fixation: in principle, some such information can be accessible for purposes of subsequent justification-giving.¹⁴ However, states of conscious awareness are all *occurrent* mental states; hence, any information that is accessible to consciousness must be information that can be explicitly represented in occurrent cognitive states. But on the psychological picture we have been describing, the holistic, Quineian/isotropic, evidential relations involved in belief-fixation can only be accommodated morphologically, rather than being explicitly represented in occurrent cognitive states. For the most part, then, such evidential relations are inaccessible to consciousness; facts that cannot be explicitly represented at all cannot be explicitly represented *consciously*.¹⁵ In short, the reasons why

¹³ This condition is also violated, it should be noted, by a form of processing that is allowed for under the liberalized construal of PT-oi/p that we mentioned in note 9—viz., processing in which certain epistemically-involved information is accommodated implicitly in a hardwired “shortcut” algorithm, and the system cannot execute a corresponding “longhand” algorithm employing that information in explicit, occurrent, form. However, in cases where a cognitive capacity has been produced by the evolutionary installation of a hardwired *shorthand* algorithm for computing a given tractably-computable cognitive transition-function, that same capacity presumably could have been produced instead by the evolutionary installation of a corresponding *longhand* algorithm for computing the same function. A cognitive system employing the longhand algorithm would be, as one might say, a *biologically feasible variant* of the actual cognitive system. So all the members of the PT family of prototheories, including the liberalized construal of PT-oi/p, impose the following requirement as a condition for a belief’s being justified: if the cognitive processing that generates or sustains the belief relies upon morphological content, then this processing must be a proceduralized shortcut version of a form of processing that (1) could have occurred in the cognitive system *or in a biologically feasible variant of it*, and (2) would be occurrently isomorphic to normative-justificatory support relations. This weakened condition too is violated by the picture of cognitive processing we have described in this section, because according to the DC framework, for *any* biologically feasible cognitive system with a large-scale set of beliefs, holistic evidential factors can only be accommodated morphologically, not occurrently.

¹⁴ This is to be expected, since procedural information is just a special case of morphological information.

¹⁵ To be sure, holistic evidential factors often can be occurrently and consciously represented in a sketchy way. Often we can articulate specific facts that bear on the cogency of a given inductive inference, and often we can say something too about the all-things-

this alternative picture of belief-fixation violates the conditions on being-justified imposed by PT(2) are also reasons why the picture violates the conditions imposed by PT(1) too.¹⁶

It is worth adding that there are additional reasons for being dubious about PT's accessibility requirement, having to do with the causal generation of perceptual beliefs. It is plausible that much epistemically relevant information—for instance, about ambient light-levels—is accommodated in perceptual belief-fixation without being either consciously noticed or accessible for subsequent conscious retrieval.¹⁷ To some extent, of course, this point can be acknowledged and incorporated by the fallback positions PT-oi and PT-oi/p. Given the essentially morphological nature of belief fixation, on the other hand, it is a plausible conjecture that certain epistemically important information that gets unconsciously taken into account in the rapid, automatic, generation of perceptual beliefs is accommodated in a way that (i) is morphological rather than occurrent, (ii) is not a proceduralized shortcut for perceptual processing that is occurrently isomorphic to the relevant justificatory-support relations, and hence (iii) is *essentially* morphological. (Think, for instance, of recognizing a partially obscured object, or of understanding sentences pronounced in a very thick foreign accent.¹⁸) I.e., it is plausible that such perceptual belief formation can only be accomplished via reliance on morphological content. To the extent this is so, such processing violates not only what is required for being justified by PT and PT-p, but also what is required by the fallbacks that do not demand accessibility, PT-oi and PT-oi/p.

In closing section III, let us comment on the overall dialectical structure of our argument against PT and its kin. We maintain that the frame-type problems encountered in classical cognitive science provide a strong reason to maintain that the holistic, Quineian/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

considered plausibility of the inference. But normally, what is articulable in this way falls far short of a full-fledged reconstruction of the complete evidential basis for the inference, with all the Quineian/isotropic factors fully spelled out. What is consciously accessible is just the tip of the evidential iceberg.

¹⁶ As a referee has pointed out, Polanyi (1958) already argued that unconscious psychological processes figure in the fixation of beliefs, including justified beliefs. It is an interesting question, which we will not pursue here, to what extent Polanyi's arguments point toward a conception of belief fixation like the one we have set forth here, which conflicts radically with PT(2) in addition to conflicting with PT(1). Prima facie, Polanyi's notion of tacit knowledge has affinities with morphological content.

¹⁷ See, for instance, Henderson (1994b, 1995).

¹⁸ Such examples suggest that perceptually generated spontaneous beliefs often are not "modular" in Fodor's sense, and that often their generation too involves holistic, Quineian/isotropic, factors.

is not essential here. Even those who are thoroughly dubious about connectionism could accept the argument from frame-type problems to the conclusion that processes like belief-fixation must *somehow* be essentially morphological—and hence that PT and kin must be repudiated.

Furthermore, even if one has doubts about the argument from frame-type problems to the claim that belief fixation is essentially morphological, one still should take seriously the implications of this claim for epistemology. For at the very least, the considerations we have canvassed in this section do show that the truth of the claim is a viable *empirical possibility*—and hence that the falsity of PT and kin is also a viable empirical possibility. This already is reason enough to make the claim relevant to epistemology.

IV. The Large-Scale Features of an Alternative Approach to the Structure of Justification

Psychological theses of the sort just presented do have epistemological consequences. As noted earlier, being justified (as opposed to merely having justification) turns on the causal dependency relations exhibited in an agent's belief generating and sustaining processes; thus, accounts of being justified purport to provide accounts of how our cognitive processes should work. To the extent that epistemological standards applicable to a system (“oughts”) should be attuned to capacities and possible capacities of that system (to “cans,” and this will soon be argued for), accounts of being justified also must be in tune with how our cognitive processes can work. As a result, empirical findings to the effect that we are so constituted that our cognitive processes work, indeed must work, along very different lines from those urged in an epistemological account tend to undermine that account. That is, the sort of results surveyed in the previous section serve as constraints on appropriate epistemological accounts. A deeper appreciation of these constraints, and of the relation between epistemological “oughts” and psychological “cans,” can be had by reflection on the nature of appropriate epistemological standards.

Epistemic standards for how we ought to reason are standards for when we will have performed our epistemic tasks well, and for when we then justifiably hold certain beliefs; such standards must make sense in terms of our epistemic ends. It is such ends—principally the end of producing a systematic and true belief-system—that distinguish the associated epistemic standards from other norms, say of prudential rationality generally. Epistemic standards characterize appropriate ways of pursuing these ends. Of course, what is appropriate in the pursuit of given ends is partially a function of the capacities, or possible capacities, of the relevant systems.¹⁹ Thus, the

¹⁹ Rambo may perhaps most deter neighborhood crime by intimidating would be malefactors—doing pushups in his front yard dressed in full battle gear, and periodically firing automatic weapons into the air, as recommended by the NRA—but most of us do

epistemic norms that are appropriate to us humans have to do with how *we* can best (or at least reasonably effectively) pursue our epistemic goals.²⁰ To formulate appropriate standards, then, we must keep in mind not just the sorts of content relations that might obtain among the information-carrying states that systems such as ourselves can have, but also *how* such systems can process such information in content-appropriate ways. In these respects, the appropriate epistemic norms must be tailored to the agents for which they are to serve as standards. This is why, in epistemology as elsewhere, “ought implies can.” The central point is that appropriate standards for strong epistemic justification, sometimes called warrant, must characterize effective “engineering solutions” for the systems in question—typically humans—with respect to epistemic tasks needful in the pursuit of our general epistemic goals.²¹ An effective engineering solution, if there is one, will be a process by which the system in question can manage such a task. One might even, with Plantinga (1993b), say that our epistemic standards have to do with the “proper functioning” of a cognitive system, but we would caution against Plantinga’s tight association of proper functioning with a “design plan.”²²

In view of these points, we can formulate the implications of the last section. We have noted that broadly empirical results regarding the nature of our psychological capacities, and how we may manage epistemic tasks, serve as a constraint on epistemological accounts. Now we find that considerations from cognitive science strongly suggest that our capacities are such that we cannot undergo the sorts of processes required by epistemic theories embracing the PT family of prototheories (PT, PT-p, PT-oi, and PT-oi/p). Thus, the PT family of prototheories make for inappropriate standards for being justified, for such standards would not be appropriately tailored to us as epistemic agents. In effect, they would require psychological processes to mirror evidential support-relations in ways to which actual belief-formation processes generally cannot conform. Happily, there are modes of processing

not have the capacity to present such an imposing front and would do better setting up a neighborhood watch program.

²⁰ Henderson (1994a, 1994b) develops this understanding of appropriate epistemic standards in writing of “epistemic competence.” Of course, the general approach to epistemic standards is related to Goldman’s (1992a, 1992b) notion of “strong justification,” and somewhat more distantly to Plantinga’s (1993) notion of “proper functioning.” Our thoughts have also been influenced by Cherniak (1986).

²¹ We should add that the needed notion of an “engineering solution” does not require a literal or figurative engineer in the etiology of the processing-capacities in question

²² One way that the “design plan” association goes wrong is in associating warrant or strong justification with standards set by a literal designer—a deity or systems engineer—or an analogous process—evolution setting a genetic plan for a class of standard organisms. For example, we believe that Swamp Thing, were there such a thing, could not only have beliefs and desires, but could also share our epistemic ends and have justified beliefs. What is important here is that, from the point of view of our epistemic goals, the designerless processing that Swamp Thing employs be effective.

that allow us to manage our epistemic tasks. These then provide more effective and appropriate engineering solutions—and understanding them should lead us to reconceive of how cognitive processes must mirror support relations in order for beliefs to be justified.

It will occur to some to try to evade the constraints employed above by insisting that epistemological accounts are to provide a characterization of what cognitive processing would have to be like for an agent to be justified in holding a given belief, and that such accounts can be indifferent to whether we humans (or any other physically realizable system, for that matter) are, or even could be, justified in holding any belief at all. Suppose one insists that PT (or something closely related) characterizes the processing that is necessary for being justified, being indifferent to whether we can conform to PT. On the basis of the previous section, one who adopted such a position would need to conclude that human beings (and like finite systems) cannot be justified. In honor of an influential proponent of such a line, we might call this “pulling a BonJour,” recalling BonJour’s recognition that humans cannot survey their full belief set for coherence (as his account would require) and his reluctant conclusion that we are consequently never actually justified in our empirical beliefs (1985, p. 152). Pulling a BonJour is implausible and unsatisfactory. For, as we have seen, the principle that epistemological “oughts” should be attuned to “cans” is rooted in the nature and point of epistemic standards.

Although we must reject those elements of epistemologies that suggest PT and kin, we can yet learn from such theories. One way of conceiving of what they can teach us is as follows. Commonly, philosophers have attempted to use justification giving as a model for being justified, as described earlier (GM). At the very least, this provides some insight into some of the information that is relevant to having and being justified. However, reflecting systematically on the kind of information that is epistemically relevant to a given belief, but very unsystematically on the sorts of processing that might be sensitive to such information, epistemologists have commonly fallen into thinking in terms of models of justification that seem impossible for cognitive systems like ourselves to satisfy. They have been abetted in so doing by a tradition that makes internalism particularly natural (Plantinga 1993a). Still, their reflection seems to have uncovered tasks with which a desirable epistemic system would need to deal. We are then put in an awkward position. On the one hand, we need to honor the epistemic desirability of dealing with the tasks in question by insisting that being justified requires systems to deal with the information whose relevance has just been recognized. On the other hand, we must avoid making unreasonable demands on creatures such as ourselves. However, the problem dissolves when we better appreciate how we might manage the epistemic tasks in question. We can both continue to recognize the importance of the information and tasks

that seem needful epistemically and, at the same time, be psychologically reasonable in our epistemic demands—provided we are willing to abandon the common assumption that the tasks must be addressed by processing that conforms to the PT family of prototheories. Thus, in the space that remains, we sketch an understanding of the sort of epistemology that emerges from taking seriously the psychological views advanced in the previous section.

One of the most striking positive epistemological lessons to emerge is that the accessible states that are the focus of traditional epistemology, and that exhaust the epistemologically relevant states according to PT and PT-p, are themselves just a proper part of what is in play as we manage epistemologically needful tasks and deal with epistemically important information. In doing a good epistemological job, our belief generating and sustaining processes must deal with information without that information being represented in occurrent states. They do this by employing information in the form of morphological content, and much of this is not accessible. As explained, such processes are crucial for our tractably managing the isotropic and Quineian elements of belief fixation. Accordingly, an account of how we can be justified in holding beliefs will need to reflect the way in which we can and must employ inaccessible as well as accessible information in managing our epistemic tasks. An account of proper epistemic functioning, an account of our epistemic competence, must recognize this full range of epistemic action.

Recognizing the full range does not preclude legitimate epistemic concern for, and focus on, certain components of that processing. In particular, it does not imply that accessible states do not remain a worthy subject for epistemic focus. To the extent that such states and transitions among them are an important component of our epistemic competence, as they surely are, there are reasons for making them one special focus of our epistemology. Nothing we have said here remotely indicates that the component of our cognitive processing involving occurrent and accessible states is merely epiphenomenal. As indicated in the previous section, such reasoning will be informed by inarticulate and commonly inarticulable processes turning on morphological content (as articulate induction depends on not fully articulable management of the isotropic and Quineian elements of belief fixation). But, such nonautonomy does not imply epistemic insignificance. Such processing causally influences the beliefs we possess and may function to further our epistemic interests in various ways, compatible with the points we have made regarding the importance of inarticulate and not fully accessible processing.

Sometimes, for example, processing with a salient articulate or accessible component may address tasks that are given parallel and complementary treatment in largely inaccessible processing. For example, expert chess players evince and report relatively quick and generally effective capacities for

recognizing situations and candidate strategies. The processes involved presumably employ much morphological content and seem largely inarticulate and inaccessible. On the other hand, these experts also report the usefulness of complementary articulate spot-checking of strategies, serving as a filter on some of the more prominent candidate strategies. In a similar way, while that component of our processing that is accessible cannot provide for all our epistemic tasks (and generally will itself depend on much inaccessible processing), still, it commonly has the advantage of affording us fairly robust self-checking and self-regulation. For example, to some extent, we can monitor ourselves in cases where we generalize from samples. We may check certain cases to insure that we have not been moved by samples that we have little reason to believe are representative. We can articulately regulate ourselves to be more attentive to sample bias. In this fashion, we may perhaps shape somewhat some of the inarticulate processes on which we must rely, while those processes will in turn continue to inform the accessible component of our processing.

The voluntarism debates associated with Alston (1985) are beside the point when thinking of accessible-level processing as just done. We may not so much choose what to believe as how to reason at this accessible level. Yet, to the limited extent that we can do this much, we have a purchase for epistemic improvement—and thus a strong motivation for epistemic focus. While not yielding a full understanding of what makes for justifiably or warrantably holding various beliefs, our understanding of epistemic competence at the accessible-level may certainly have special epistemic payoff.

Another reason for devoting some special attention to articulable processes arises when we think in terms of social epistemology—our interests in the cooperative production of a system of true beliefs. If we are to learn from each other, and to help each other learn, we must be able to access each other as sources of information, and to access packages of information that are provided. If we are to work together, we will want to inform each other's reasoning in various respects. But, in order for multiple-person interactions to be effective in transmitting reliable information (as opposed to scurrilous gossip, unreliable rumor, narrow-minded prejudice, and the like), these interactions will need to conform to certain epistemic norms governing accessible psychological phenomena—norms that determine, for instance, when it is reasonable to request from another person an explicit justification for a claim the other has made, and what counts as an adequate response to such a request.²³ Reflecting on the myriad ways in which we manage such things

²³ Some of these concerns are reflected nicely in Miller (1995). Alston's reflections (1989, pp. 234–36) on the source of the concern for accessibility seem to anticipate roughly the reasons for that concern we mention here.

gives ample reason for making the giving and eliciting of reasons—and *a fortiori* articulable processes—one focus of our epistemology.

We then are led to at least two legitimate epistemological foci within a general account of our epistemic competence. It must be emphasized that these are complementary, not competing, foci within a unified account of our full epistemic competence. First, there is the full set of epistemically relevant states and processes, the wide focus. Second, there is the focus on accessible states and the articulable component of our cognitive processes. The analogy of an iceberg suggests itself here. The wide-focus processes are the whole ball of ice. But, just as only a small portion of an iceberg is visible from above the water, so only a portion of those processes are classically accessible; call these “exposed processes.”²⁴ The exposed processes cannot be understood without appreciating the wide-focus processes, of which they are a component, much as the behavior of the exposed ice cannot be understood without appreciating the glacial mass of which it is a part. We may have reasons for giving special attention to the exposed ice—for example, it serves as a ready indicator of the presence of the rest, and it may provide bases for study and intervention. We have indicated some epistemic parallels—for example, that exposed cognitive processing may provide special opportunities for self-regulation. Adopting these metaphors, we call ours *iceberg epistemology*.

Among the most intriguing aspects of iceberg epistemology is the way in which various themes that are traditionally associated with competing epistemological accounts naturally come to be seen as complementary when we think in terms of varying our focus. We then can learn from and find partial vindication of themes dear to the hearts of coherentists, foundationalists, even contextualists.

Thus, consider (the full set of) wide-focus processes. We have argued that these must be recognized in order to understand how we cognitively manage the isotropic and Quineian elements of our belief-fixing processes. In doing so, we come to appreciate just how important morphological content is to our epistemic competence. This allows us to salvage elements of coherentism. In effect, coherentists such as BonJour (1985) reflected on the kinds of information that are relevant to the justification of various beliefs, and quite properly were led to appreciate the in principle exceedingly ramified character of those relevance relations. Dealing with such relevance relations, somehow taking account of them, and somehow taking care that one’s set of relevant beliefs is itself in order, were then recognized as epistemic tasks of the first importance. These tasks are associated with all, or most, more particular epistemic tasks. Such characteristic coherentistic themes receive vindication.

²⁴ An “exposed process” is the analog of being roughly at or above the water line. The analogy is apt in that to be exposed is not to be accessed, but to be accessible. To be above the water line is to be readily visible, not necessarily to be glimpsed.

On the other hand, BonJour, for example, falls into thinking that these tasks would need to be managed by some form of occurrent surveying of one's entire set of beliefs. He rightly finds this intractable, and ends in a quiet skepticism. We can now see that BonJour's identification of coherentist epistemic tasks was on track, and it was his understanding of how we might pull these off that was too limited. This is something of a partial vindication of coherentism. What is vindicated is the epistemic importance of beliefs being conditioned by information that is so ramified that one's global informational possession, and its global features, seem implicated.²⁵ One thing that is not vindicated is the traditional coherentist idea that all the justificatory action is a matter of relations between belief states (traditionally conceived). More generally, the limitation of epistemically relevant states in keeping with PT(1) is rejected. We can understand how we might manage to employ the mass of relevant information that BonJour's coherentism would suggest as relevant, but we can only do this when giving up the suggestion that this information is all encoded in traditionally understood belief-states, or even in accessible states more generally.²⁶

A different structure seems characteristic of the exposed epistemic processes. For, while this processing will be conditioned by the processing that is not exposed, it will itself have a less ramified structure that aptly might be described as quasi-foundationalist. Recall Audi's characterization of the generation of justified empirical beliefs: he describes a simple development, with perceptual states directly generating perceptual beliefs (basic

²⁵ A more standard formulation would be to say that "one's global belief-set, and its global feature, seems implicated." However, we have argued that our cognitive system possesses and employs much information in the form of morphological content. It is an open question just how much of this is properly spoken of as a matter of beliefs we possess. The nonstandard formulation is intended to bracket this question. However, as the text reflects, it is fairly clear that the relevant informational content is not limited to traditionally understood belief-states.

²⁶ The structure of support relations within the wide-focus processes need not be wholly a matter of global coherence relations. When we flesh out our understanding of these processes, foundationalist themes may find some place along side the important Quineian/isotropic features we have discussed. Such a partial vindication of foundationalist themes seems likely to the extent that we would count particularly rudimentary input states serving as inarticulable bases, conditioning other states while themselves being largely unconditioned by other informational states, as providing such a vindication. If this line of thought were developed, one would likely end with a very nontraditional intermingling of coherentist and foundationalist themes: on the one hand, all normatively appropriate beliefs, and much (or even all) normatively appropriate morphological content, would receive some epistemic support from their place within the global system of psychological states; while, on the other hand, certain kinds of non-belief psychological states—in particular, certain occurrent states that figure causally as inputs to perceptual processing, provide some degree of non-coherentistic basis for other states (e.g., to certain cognitively spontaneous beliefs). This sort of possibility is similar in spirit to the approach to the structure of justification suggested in Haack (1993); see also note 28 below.

beliefs), followed by the more complex generation of nonbasic beliefs. Earlier, we noted that much information that rightfully influences epistemic processing is omitted from this story about the generation of perceptual states and perceptual beliefs. However, taken as a proper part of the epistemic story, one focused on exposed processes, there may be little wrong with this. The dependency relations within the exposed processes seem commonly to have just such a structure. That is, in the model cases, where the unexposed processes do not throw up special considerations to attend to, the exposed processing will commonly have such a direct phenomenological character.

However, this vindication of foundationalism is partial (as it was for coherentism). The reason is not just that the exposed processes, with which this vindication is concerned, can only be part of the epistemological account. Additionally, even at this exposed level, foundationalist themes are imperfectly vindicated. Foundationalism is commonly understood as an answer to the regress argument. It is recognized that the justification of one belief often turns on other beliefs, and that these beliefs must themselves be justified if they are to provide justification. Accordingly, a belief that is justified in terms of others seems to initiate a regress of justification. In the classical argument, what is at issue is the structure of this justificatory regress. It is argued that it cannot be either circular or infinite. So it is concluded that there must be basic beliefs, stopping places, whose justification does not turn on other beliefs. This is the central tenet of foundationalism. However, if there are to be such basic beliefs, whose status as such is not mysterious, they must constitute a homogeneous class (or small set of homogeneous classes) having properties that give them their status. The class of perceptual beliefs (as characterized by Audi and many others) has provided the most obvious candidate. The idea is that such beliefs are of a sort typically arising directly from subdoxastic perceptual states—the belief that there is a brown chair in front of me, that there is a North American Goatsucker on the overhanging branch, or that a car is entering my driveway. It is thought that, when spawned by perceptual states, beliefs with such content are justified independently of other beliefs. While this may characterize the typical structure of justificatory processing within the exposed processes, it ultimately only characterizes the typical cases—and there will be extraordinary cases. As a result, the foundationalist themes only get partial vindication.

To see why the homogeneity theme of foundationalists is vindicated only as a rough approximation, recall that wide-focus processing influences the exposed processing. This produces an important diversity in the structure of exposed processes having to do with perceptual beliefs. In an inarticulate way, we may be sensitive to various relevant information that conditions our generation of perceptual beliefs. In most cases, the standard or normal cases, this information may remain largely inarticulate. However, in some cases, “caution flags” may be thrown up by non-exposed processes, so that, even

within the exposed-processing, the perceptual belief may arise justifiably only with the midwifery of further beliefs. For example, as a noted ornithologist, Willard may occasionally find himself with a belief that there is North American Goatsucker on the overhanging branch. In most cases this may arise “directly,” in that, focusing on exposed processes, we would correctly say it arises out of antecedent perceptual states without the involvement of other doxastic states. In other cases, however, a perceptually spawned belief with just that content may turn on other beliefs operating within Willard’s exposed processes, beliefs whose place in that processing is fostered by the non-exposed processing. For example, knowledge of the range of the Goatsucker, and of one’s rough location, may enter into a quick judgment of the likelihood of encountering such a bird—and this may play a crucial role in generating Willard’s justified belief that there is one on that branch.

In any case, as the result of conditioning by unexposed processes, the structure of justificatory exposed processing will turn out to be a rather rough-hewn or irregular affair. The class of stopping places will be rather more heterogeneous than foundationalists would expect; and, insofar as one considers exposed processes only, some stopping places will seem somewhat arbitrary.²⁷ A given perceptual belief will serve as a stopping place in one instance (e.g., the belief is cognitively spontaneous, and the agent experiences no accompanying sense of uncertainty), while another perceptual belief—even one with the same content and held by the same agent—will be in need of further support within exposed processes (e.g., because an accompanying sense of uncertainty, a “caution flag,” has been thrust into consciousness by unexposed processes). Insofar as standard cases fit the foundationalist mold, it seems proper to call the resulting normative structure of exposed processing *quasi-foundationalist* or *ceteris-paribus foundationalism*.

It is worth noting that one line of thought associated with what is called contextualist epistemology has been that, while the regress of justification comes to an end, the regress stoppers are a heterogeneous lot. Timmons (1993) refers to this as structural contextualism, and Henderson (1994b) argues that an appropriate theory of epistemic competence will vindicate structural contextualism. The above remarks on the partial vindication of foundationalism point in just this direction.

Of course, exposed states and processes are just the tip of the iceberg of epistemically involved psychological states and processes. Thus, beliefs that are quasi-foundational and contextually basic *with respect to exposed processes* can still possess—and in general will possess—epistemic support from other psychological states *within the full iceberg*. For one thing, such beliefs typically will receive epistemic support by way of their situatedness in a system of pervasive, holistic, mutual support among items in the full

²⁷ These themes are developed in more detail in Henderson (1994b).

iceberg—just as each entry in a crossword puzzle receives support from the overall mutual fit of all the entries with one another. Coherentistic themes, applicable to the iceberg as a whole, thus link up nicely with foundational and contextualist themes that apply to exposed states and processes. In addition, the complete iceberg might well include occurrent non-belief psychological states that contribute fairly direct epistemic support (perhaps via the mediation of morphological content) to beliefs that are quasi-foundational at the exposed level—just as an entry in a crossword puzzle receives fairly direct support from an associated clue.²⁸

In summary, we began by pointing out that philosophers have come to recognize that an agent's being justified in holding a given belief requires that the belief arise or be sustained by the right sort of cognitive processes—ones that “mirror” the epistemically important structure of support relations within the information the agent possesses. This recognition, which reflects the pre-theoretic concept of being justified, leads to the difficult question of just what sort of mirroring is needed.²⁹ We then set out a family of answers to this question that we believe has influenced and has been presupposed in much contemporary epistemology—the PT family. PT and its kin purport to characterize what epistemically virtuous psychological processing is, and must be, like. At least in the full dress version, for a cognitive process to mirror normative support relations in the right way, that process must turn on accessible states causally interacting in the form of occurrent states, and this interaction of occurrent states must be isomorphic to the structure of content relations featured in the agent's having justification. Variations on PT emphasize certain elements and weaken others. However, in an important respect, the variants each hold to the central idea that epistemically requisite mirroring is a matter of occurrent isomorphism, even when they do not require processing that fully realizes such occurrent isomorphism. That is,

²⁸ The crossword puzzle metaphor is due to Haack (1993), who defends a view about the global structure of justification that she calls ‘foundherentism’. Foundherentism resembles traditional coherentism in its emphasis on the holistic nature of epistemic support relations, but it rejects the traditional coherentist idea that beliefs are the only kinds of psychological states involved in such relations, and also rejects the idea that all epistemic support is (or is derivative from) holistic coherence relations among the set of epistemically involved states (cf. note 26). Certain non-belief states, such as occurrent experiential states, can play a role too, analogous to the role played by the clues in a crossword puzzle.

From the wide, whole-ball-of-ice, perspective within iceberg epistemology, such occurrent states are likely to occur below the level of consciousness, or anyway to occur so fleetingly that they cannot be subsequently remembered. For discussion of cases where fleeting, subsequently inaccessible, occurrent states of awareness contribute epistemic support (probably via morphological content) to a cognitively spontaneous perceptual belief, see Henderson (1994b).

²⁹ In our view, the answer to this question is left somewhat open by the pre-theoretic concept of being justified, and thus is partly an empirical matter. In this respect the concept of being justified resembles natural-kind concepts, like the concept *water*.

even in those variants that do not require processing that is occurrently isomorphic to the normative support relations (PT-p, PT-oi/p) there is at least the requirement that the processing employed is, in effect, a shorthand version of occurrently isomorphic processing which could have been instanced instead. We then set out lessons emerging out of recent cognitive science; these tell us that human cognitive capacities crucially depend on processes that are radically different from the sorts of processes enshrined in PT and kin. To cope with the Quineian and isotropic aspects of cognition, we do and must employ much information in the form of morphological content. If our norms for being justified are to be appropriate to the sorts of cognitive systems that we are, they must allow for our managing our epistemological chores by relying on this sort of processing. Thus, presuming PT has foisted on epistemologists a distorted understanding of the way in which an agent's cognitive processes should mirror the dependency relations between informational states that commonly are recognized as epistemically significant. The processes we describe would evince causal or counterfactual dependencies between informational states that do mirror contentfully appropriate support relations, but do so without occurrent isomorphism (and without what would merely be a shorthand stand-in for occurrent isomorphism). This, we think, is all that should be demanded in the way of mirroring—all that can appropriately be demanded in our epistemic norms. Finally, on this line of thought, there will be multiple foci of epistemic interest. And these differing foci allow the partial vindication of epistemic approaches that previously have been commonly seen as competing. A wide focus accommodates coherentist themes. A narrow focus leads to a quasi-foundationalist, structural contextualist, view. Both foci are appropriate, and the supported views are not competing. Beliefs that are contextually basic and quasi-foundational within the tip of the iceberg can be epistemically supported within the total iceberg—by virtue of their position within an epistemically coherent whole, and perhaps in other ways as well.³⁰

³⁰ We thank Robert Audi, Jonathan Kvanvig, Michael Lynch, Herman Stark, Mark Timmons, John Tienson, several anonymous referees, and several university audiences for discussion and/or comments.

References

- Alston, W. 1985: "Concepts of Epistemic Justification," *Monist* 68: 57–89.
- _____. 1989: "An Internalist Externalism," in W. Alston, *Epistemic Justification: Essays in the Theory of Knowledge*. Cornell University Press.
- Audi, R. 1993: *The Structure of Justification*. Cambridge University Press.
- Berg, G. 1992: "A Connectionist Parser with Recursive Sentence Structure and Lexical Disambiguation." in *AAAI-92: Proceedings of the Tenth National Conference on Artificial Intelligence*. AAAI Press/M.I.T. Press.
- BonJour, L. 1985: *The Structure of Empirical Knowledge*, Harvard University Press.
- Cherniak, C. 1986: *Minimal Rationality*, M.I.T. Press.
- Fodor, J. 1983: *The Modularity of Mind: An Essay in Faculty Psychology*, M.I.T. Press.
- Fodor, J. and Pylyshyn, Z. 1988: "Connectionism and Cognitive Architecture: A Critical Analysis," *Cognition* 35: 183–204.
- Goldman, A. 1992a: "Strong and Weak Justification," in A. Goldman, *Liaisons*, M.I.T. Press.
- _____. 1992b: "Epistemic Folkways and Scientific Epistemology," in A. Goldman, *Liaisons*, M.I.T. Press.
- Haack, S. 1993: *Evidence and Inquiry: Towards Reconstruction in Epistemology*. Blackwell.
- Henderson, D. 1994a: "Epistemic Competence," *Philosophical Papers* 23: 139–67.
- _____. 1994b: "Epistemic Competence and Contextualist Epistemology," *Journal of Philosophy* 91: 627–49.
- _____. 1995: "One Naturalized Epistemological Argument Against Coherentist Accounts of Empirical Knowledge," *Erkenntnis* 43: 199–227.
- Horgan T. and Tienson, J. 1994: "A Nonclassical Framework for Cognitive Science," *Synthese* 101: 305–45.
- _____. 1995: "Connectionism and the Commitments of Folk Psychology," *Philosophical Perspectives* 9: 127–52.
- _____. 1996: *Connectionism and the Philosophy of Psychology*. M.I.T. Press.
- Horgan, T. 1997a: "Connectionism and the Philosophical Foundations of Cognitive Science," *Metaphilosophy* 28: 1–30.
- _____. 1997b: "Modelling the Noncomputational Mind: Reply to Litch," *Philosophical Psychology* 10: 365–71.
- Marr, D. 1982: *Vision*. W. H. Freeman.

- Miller, R. 1995: "The Norms of Reason," *Philosophical Review* 104: 205–45.
- Moser, P. 1989: *Knowledge and Evidence*. Cambridge University Press.
- Plantinga, A. 1993a: *Warrant: The Current Debate*. Oxford University Press.
- _____. 1993b: *Warrant and Proper Function*. Oxford University Press.
- Polanyi, M. 1958: *Personal Knowledge*. Harper and Row.
- Pollock, J. 1986: *Contemporary Theories of Knowledge*. Rowman and Littlefield.
- Pollack, J. 1990: "Recursive Distributed Representations," *Artificial Intelligence* 46: 77–105.
- Segal, G. 1989: "Seeing What Is Not There," *Philosophical Review* 98: 189–214.
- Sejnowski, T. and Rosenberg, C. 1987: "Parallel Networks that Learn to Pronounce English Text," *Complex Systems* 1: 145–68.
- Smolensky, P. 1990: "Tensor Product Variable Binding and the Representation of Symbolic Structures in Connectionist Systems," *Artificial Intelligence* 46: 159–216.
- _____. 1995: "Connectionist Structure and Explanation in an Integrated Connectionist/Symbolic Cognitive Architecture," in C. MacDonald and G. MacDonald, eds., *Connectionism: Debates on Psychological Explanation*, Vol. 2. Blackwell.
- Stark, H. 1993: "Rationality is not Algorithmic," *Kinesis* 20: 16–30.
- _____. 1994: "Connectionism and the Form of Rational Norms," *Acta Analytica* 12: 39–53.
- _____. 1995: "Rationality Without Rules," Doctoral Dissertation, University of Memphis.
- Timmons, M. 1993: "Moral Justification in Context," *Monist* 76: 360–78.