

Journal of Philosophy, Inc.

Epistemic Competence and Contextualist Epistemology: Why Contextualism is Not Just the Poor Person's Coherentism

Author(s): David K. Henderson

Source: *The Journal of Philosophy*, Vol. 91, No. 12 (Dec., 1994), pp. 627-649

Published by: Journal of Philosophy, Inc.

Stable URL: <http://www.jstor.org/stable/2940759>

Accessed: 20/02/2009 16:57

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=jphil>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit organization founded in 1995 to build trusted digital archives for scholarship. We work with the scholarly community to preserve their work and the materials they rely upon, and to build a common research platform that promotes the discovery and use of these resources. For more information about JSTOR, please contact support@jstor.org.



Journal of Philosophy, Inc. is collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Philosophy*.

<http://www.jstor.org>

THE JOURNAL OF PHILOSOPHY

VOLUME XCI, NO. 12, DECEMBER 1994

EPISTEMIC COMPETENCE AND CONTEXTUALIST EPISTEMOLOGY: WHY CONTEXTUALISM IS NOT JUST THE POOR PERSON'S COHERENTISM*

This paper provides one argument for a contextualist epistemology of empirical knowledge. Contextualism is a neglected and generally underdeveloped alternative to the more familiar epistemic approaches of foundationalism and coherentism. The argument I present is an example of the manner in which naturalizing epistemology can contribute to fairly traditional epistemological issues.

I

As stage setting, I must explain how I conceive of epistemological standards and of the implications of psychological results for such standards. Epistemology is properly a normative discipline: it should provide normative standards for the direction of our reasoning in inquiry. Since inquiry is itself reasoning directed to producing that classic epistemological value, true beliefs, the normative force of our standards may readily be understood as inherited—as a function of their promise to facilitate the production of comprehensive true theories.¹ Epistemology is thus to provide us with a set of hypothetical imperatives that meliorate our epistemic practice: standards for *how* we can effectively pursue the production of true systems of belief. To provide the appropriate hypothetical

* I wish to thank Mark Timmons, Terry Horgan, and John Teinson for helpful comments on earlier versions of this paper.

¹ This formulation is overly schematic in several respects. First, it must be acknowledged that not all true statements are created equal, even when the production of true theories of the world is taken to be the central epistemic value. Second, under this crude umbrella characterization, there is room for a range of related individual and social epistemic goals. As a result, the psychological focus employed here is an oversimplification. These reservations are nicely accommodated in Philip Kitcher, *The Advancement of Science* (New York: Oxford, 1993). Such qualifications seem particularly marginal in their impact on the argument of this paper, however, and are thus set aside.

imperatives, epistemology must be naturalized; for it must deal with how we, as a class of cognitive systems, in a sort of environment, can effectively be put to work producing true beliefs.²

In keeping with these observations, I propose that normative epistemology be directed to the production of what I call theories of *epistemic competence*.³ These are distilled empirical analyses of how we do, or could, reliably manage certain cognitive-epistemological tasks. The basic idea is that we can select from empirical analyses of our cognitive dispositions, accentuating just those dispositions which are particularly truth conducive. When such information is culled to provide a normative model of what is effective in our present practices, we have a model of our *actual epistemic competence*. The resulting standards reinforce our epistemically desirable practices. When such information is employed to construct a model of yet more epistemically effective agents that we could become, we produce models of more *ideal epistemic competencies*. Such standards direct us in how we can meliorate our practice.

It is worth emphasizing that such models of competence (whether of actual or ideal competence) provide standards appropriate to us, as human cognitive agents (in the environment in which we find ourselves). This conforms to the motivation for naturalized epistemology. It also rules out epistemological ideals calling for certain classes of methods that have been commonly prescribed in more traditional epistemology. The *central constraint* is straightforward: our epistemological standards—our ideals for human epistemology—must recommend or prescribe methods or processing that (1) can be (thoroughly) implemented by humans, and (2) when so implemented, are truth conducive. To say that a method or process is (thoroughly) implementable by humans is to say that humans can come to have cognitive dispositions to the processing in question—that, barring interruptions or defeating background conditions, they would then evince that processing whenever occasioned. This formulation is intended to recognize that the relevant dispositions would conform to generalizations containing ineliminable *ceteris paribus* clauses, as is characteristic when dealing with psychological systems.⁴

²These themes and the meliorative concerns that motivate such naturalizing of epistemology are dealt with by Kitcher, "The Naturalist Returns," *Philosophical Review*, CI (January 1992): 53-114.

³In my "Epistemic Competence," forthcoming in *Philosophical Papers*.

⁴Horgan and Tienson, "Representations without Rules," *Philosophical Topics*, xvii (Spring 1989): 27-43; "Soft Laws," *Midwest Studies in Philosophy* (1990): 256-79; *Connectionism and the Philosophy of Psychology: Representational Realism without Rules* (Cambridge: MIT, forthcoming); and my *Interpretation and Explanation in the Human Sciences* (Albany: SUNY, 1993), ch. 8.

This central constraint does not merely require that (absent defeating conditions) humans could *sometimes* implement the processes. This would be too weak, for a method that is sometimes intractable, even in benign (endogenous and exogenous) conditions, is not a method to be generally prescribed. Thus:

- (A) The central constraint rules out a standard that calls for a method whenever, within the set of problems to which that method applies, there are problems for which the processing cannot be implemented, even in ideal conditions for the systems in question.

Associated with the idea that ought implies can, (A) is reflected in several influential discussions⁵ in which classes of epistemic norms are criticized for being intractable.

Violations of (A) are principally significant because they would mandate that epistemic agents possess dispositions that they neither do nor can have. This is to misrepresent the sort of processing by which the agent, as a member of a sort of cognitive system, does, or even could best, inquire in a truth-conducive manner. A particularly serious class of such misrepresentation involves standards that mandate processes that would be commonly intractable in practice; this obtains when problems arise commonly or routinely for the system in its environment for which the system could not implement the processing recommended. We can distinguish two sorts of down-to-earth violations of the central constraint:

- (B) The central constraint rules out any standard that calls for a method whenever, within the set of problems with which the system would routinely deal, there are sets of problems for which the processing cannot be jointly (across-the-board) implemented, even in ideal conditions for the systems in question.
- (C) The central constraint rules out any standard that calls for a method whenever, within the set of problems with which the system would routinely deal, there are individual problems for which the processing cannot be implemented, even in ideal conditions for the systems in question.

Now, (B) has informed several recent epistemological discussions. For example, some foundationalist accounts would require us to keep track of chains of justifying beliefs, as this seems required in order to ensure that our beliefs are well grounded. It is clear from psychological investigations, however, that we are very poor at keeping track of justificatory chains. Gilbert Harman argues that our characteristic

⁵ Christopher Cherniak, *Minimal Rationality* (Cambridge: MIT, 1986); Gilbert Harman, *Change in View* (Cambridge: MIT, 1986).

performance should not be seen as violating any proper epistemic standard, we should not keep track of such justificatory pedigrees (*op. cit.*, ch. 4). The reason is fairly simple: we could not keep track of all that. For belief sets of any interesting size, certainly for beliefs sets of the size that we commonly need in our everyday rounds, keeping track of myriad chains of linear justification (even closely approximating what the foundationalist would see as ensuring proper grounding for all member beliefs) would simply be intractable—to demand such violates (B). In effect, the pedigree certification problems faced within any reasonably sized set of beliefs would become *jointly intractable*.

Christopher Cherniak argues for caution in employing standard logical methods as epistemic standards. Drawing on complexity theory, he points out that, while we may metatheoretically know of a decision procedure for a class of problems, still the application of such a procedure to particular problems can be intractable for us (or even for any physically realizable system) (*op. cit.*, ch. 4). Cherniak discusses how this sort of intractability characterizes the tests for truth-functional consistency familiar from elementary logic. His discussion illustrates the force of (C), if we plausibly suppose that the onset of intractability turns on the size of our belief set, and that the onset readily affects sets within the range characteristic of human belief systems.

These constraints, and particularly (C), will feature centrally in my arguments.

II

I now briefly sketch the epistemological approaches of foundationalism, coherentism, and contextualism. This will help clarify what I intend to show. We can begin by recalling Aristotle's classic regress problem, which, of course, was formulated in an argument for foundationalism. Suppose that we are set the task of justifying a given belief *B*. It is common to think of this as a matter of pointing to other beliefs from which we can infer *B*. Of course, if *B* is to be inferentially justified, it would seem that these other beliefs must themselves already be justified. Here, we conceive of justification as possessing an essentially linear structure in which justification is passed on from certain beliefs already possessing it to the beliefs to be justified. When we ask how the justifying beliefs themselves are justified, we initiate a regress of justification. The problem, classically, is to explain how such a regress can be satisfactorily structured, and thus to characterize the structure of justification. It seems clear that such inferential justification can neither go on infinitely, appealing ever to further beliefs, nor circle back onto itself. So, if justification is

essentially linear, and so it is supposed for purposes of the regress argument, the regress of inferential justification must terminate in beliefs having a significant fund of noninferential justification. On this view, if skepticism is to be avoided, there must be beliefs that have justification without the inferential aid of other beliefs. It is the hallmark of foundationalist approaches to purport to identify a distinct, homogeneous class of such basic beliefs.

A central problem with foundationalism is that there turns out to be a paucity of candidate epistemic types on which defensibly to confer the status of basic belief. The arguments for this conclusion cannot be rehearsed here.⁶ I simply accept their conclusion and look to the remaining two approaches. For the most part, my argument for contextualism simply presupposes the arguments against foundationalism as a telling background, it then provides one reason for preferring contextualism to coherentism.

While coherentism has traditionally been something of a dialectical foil for foundationalists, it has come to be a serious competitor; to my mind, it is clearly the superior position. Coherentists reject the linear model of justification that is central to both foundationalism and the regress argument. On this view, the justification of our comprehensive belief system is prior to the justification of particular beliefs, and this systemic justification is a matter of several sorts of coherence relations among the member beliefs. This is nicely captured in Laurence Bonjour's model for the justification of individual beliefs, which requires systemic justification along the way. According to [coherentism], the fully explicit justification of a particular empirical belief would involve four distinct main steps or stages of argument, as follows:

- (1) The inferability of that particular belief from other particular beliefs and further relations among particular empirical beliefs.
- (2) The coherence of the overall system of empirical beliefs.
- (3) The justification of the overall system of empirical beliefs.
- (4) The justification of the particular belief in question, by virtue of its membership in the system.⁷

The coherence relations of step (2) include, but are far richer than, simple relations of consistency within the system of beliefs. They

⁶The arguments are commonly tailored to the sort of basic belief candidate entertained. One particularly general and forceful argument is to be found, however, in Laurence Bonjour, "Can Empirical Knowledge Have a Foundation," *American Philosophical Quarterly*, xv (January 1978): 1-13, and *The Structure of Empirical Knowledge* (Cambridge: Harvard, 1985).

⁷*The Structure of Empirical Knowledge*, p. 92.

include the various sorts of relations (explanatory, evidential, and so forth) canvassed in contemporary understandings of the unified sciences.

While coherentists see nonlinear global justification of belief systems as the central justificatory task, they also recognize that, with this task roughly in hand, there is a place for episodic linear justifications of particular beliefs. In such contexts, the globally justified theory serves as an understood background on which to draw. Thus, to see the justification of individual beliefs as linear is, at best, somewhat elliptical, for such justification ultimately relies on the nonlinear justification of a belief system. Accordingly, coherentists are able to recognize analogues to the foundationalist's terminating regress of justification: local justificatory episodes do "come to an end," terminating in a class of beliefs that, for purposes of that episode, seem simply accepted.

This core coherentist model can be elaborated to provide an account of empirical input. Again, BonJour's treatment seems appropriate. The basic idea is that, while there are beliefs that *arise noninferentially*, these must still be *justified inferentially*. These (characteristically rich and "forceful") beliefs are termed "cognitively spontaneous beliefs" by BonJour. Their coherentist mode of inferential justification would be a case of the sort of local linearity described above: the cognitively spontaneous belief is vindicated by being inferred from the agent's globally justified system. Although the belief may arise noninferentially, it is only epistemically certified when the appropriate inferential justification features in its retention. The rudimentary model for the justificatory inferential processes would take the following form:

- (1) I have a cognitively spontaneous belief that *P* which is of kind *K*.
- (2) Conditions *C* obtain.
- (3) Cognitively spontaneous beliefs of kind *K* in conditions *C* are very likely to be true.

Therefore, my belief that *P* is very likely to be true.

Therefore, (probably) *P* (*ibid.*, p. 123).

This form of argument, together with the more general four-step model, comprise the simplest coherentist model of empirical knowledge and justification. I have drawn on BonJour in setting out the coherentist position, for his account seems to me to indicate the direction in which any plausible fundamentally coherentist account would need to be developed.

Contextualist epistemologies retain elements of both foundationalist and coherentist epistemologies, but each is significantly altered.

They can and do allow that the sort of systemic relations of mutual support featured in coherentist accounts play a central role in the justification of systems of belief. Contextualists insist, however, that these desiderata are never, and could never really be, applied to our belief system globally. A strictly global application (to the whole of our belief set)—the central coherentist ideal—is simply not an ideal for contextualists at all. Such global coherentist ideals are eliminated for two sorts of reasons. On the one hand, it is not a fitting standard for any finite cognitive system, and particularly not one appropriate to systems with capacities in the range characteristic of human cognizers. On the other hand, because our epistemic standards are themselves conditioned by our other beliefs, putting the whole system “up for grabs” at once would undercut all epistemology, and thus be self-defeating.⁸

Similarly, central elements of foundationalism are appropriated and transformed in contextualist epistemology. In particular, along with the foundationalists, contextualists recognize that, as an ideal as well as in practice, linear episodes of belief justification can properly terminate in beliefs that really are not susceptible to, or helped by, further (presumably coherence-theoretic) justification. Contextualists contend, however, that these regress-stopping beliefs do not comprise an epistemically homogeneous class with a distinct characteristic that gives them a measure of justification. Imagine belief tokens being fully characterized in epistemically relevant terms that would interest various foundationalists. Contextualists insist that one belief may stand in need of no further justification, and serve as a regress stopper, while a second belief, characterizable in just the same way as the first, may stand in need of justification. To foundationalists and coherentists, this proposal will look to be inconsistent. The foundationalist will insist that, if a belief of a (fully characterized) epistemic sort can serve as a regress stopper anywhere, such beliefs must always and everywhere possess some measure of justification on their own. The coherentist will insist that, if beliefs of a sort need justification anywhere, this merely reflects the fact that cases of such beliefs, when justified, are at least backed by “implicit justifications.” Accordingly, coherentists hold that such beliefs, when not explicitly justified, are only approximately justified, only justified

⁸This second objection is most forcefully formulated in Michael Williams, *Groundless Belief* (New Haven: Yale, 1977), and “Coherence, Justification, and Truth,” *Review of Metaphysics*, xxxiv (December 1980): 243-72. It may also be reflected, however, in Dudley Shapere, “Method in Philosophy of Science and Epistemology,” in *The Process of Science*, N. Nersessian, ed. (Dordrecht: Kluwer, 1987), pp. 1-39; and in Kitcher, “The Naturalist Returns.”

against a background of presupposed coherence-theoretic justification. When the background justification called for by coherence-theoretic ideals is not evinced, as it would not be in contextualist practice, the beliefs seem only “potentially justified,” if justified at all.⁹ I explain below how the seemingly unstable contextualist proposal can make epistemic sense.

In much of the contextualist literature, there is the further suggestion: that when and where a belief token can serve as a regress stopper is ultimately a (perhaps negotiated) social matter. This theme is commonly associated with the Wittgensteinian heritage of much of contextualist thought. I, however, am not convinced of the force or exact nature of these claims. I leave them aside for purposes of this paper. The contextualism I defend here might best be termed *structural contextualism*,¹⁰ for I am interested in it as an alternative model of the structure of epistemic justification, a model that is distinct from coherentism and foundationalism in the ways just mentioned.

There is a tendency among advocates of coherentism and foundationalism to conceive of contextualist practice as something of a poor-kin approximation to the more mainline positions, particularly coherentism.¹¹ After all, what contextualists seem to envision roughly accords with what one would get if one were approximating coherence-theoretic justifications to the extent that such could reasonably be implemented by finite agents who are subject to varying pressures to settle on beliefs and act accordingly. Are the contextualist ele-

⁹Sophisticated coherentists such as Bonjour seem to soften their ideals by formulating them in terms of the accessibility, rather than implementation, of coherentist evaluation (*The Structure of Empirical Knowledge*, pp. 92, 102, 123). But, they clearly keep an eye on the more simpleminded model while spinning out their moderated coherentism. So much so that, on sophisticated variants, even understandable deviations from the simple ideal come still to be conceived as problematic (pp. 152-3).

¹⁰Various common contextualist themes are carefully distinguished in Timmons, “Moral Justification in Context,” *The Monist*, LXXVI (July 1993): 360-78, and in “Outline of a Contextualist Moral Epistemology,” in *Moral Knowledge*, W. Sinnott-Armstrong and Timmons, eds. (New York: Oxford, forthcoming).

¹¹For example, Bonjour classifies Williams’s avowedly contextualist account in “Coherence, Justification, and Truth” as a kind of unambitious coherentism (an understanding that is fostered by Williams’s initial discussion of his view using the label ‘coherentism’ which is later dropped). Admittedly, coherence theoretic aspects of contextualist views have occasionally led those sympathetic to contextualist epistemology to conceive of it as having such a secondary status—as a mere approximation to other ideals. This is not true of Williams’s own contextualism, for he argues that the very basis of our epistemic standards is incompatible with those most basic standards being thoroughgoing coherentist ideals (which are holistic in a problematic manner). This argument marks similarities in coherentist and contextualist understandings, while insisting that the former are internally unworthy at the level of ideals, not just unworkable in practice.

ments in our proper cognitive practice the results of playing coherentist hunches that further inquiry would likely produce (at most) marginal revisions in the beliefs being evaluated? It is the central contention of this paper that contextualist practices are emphatically *not* to be so understood. To so understand them would be to see them as excusable and understandable deviations from what is properly the ideal model—as practically acceptable flaws. In contrast, I insist that the proper ideals are contextualist, not coherentist.

III

Now, were the simple coherentist model outlined above taken as a model of what processing an agent must go through in order to be justified, it would not be at all plausible as an epistemic standard for human beings. It would be obviously overambitious. The general justificatory schema seems to require that agents could have an explicit grasp of their full system of beliefs so that, within one round of processing, an agent could survey those beliefs and check them for the global features comprising coherence. These demands run afoul of constraint (C), however. To begin with, even the suggestion that one's entire belief set could be surveyed is problematic, a point that does seem to trouble Bonjour.¹² A fortiori, the suggestion that the agent has the capacity to gauge the coherence of her belief set is problematic, particularly when limitations on the speed and power of human cognitive endowments are kept in mind. The point is reflected in Cherniak's discussions, mentioned above, and is pressed with particular force by Hilary Kornblith.¹³ It is at least equally clear that the simple model for justifying cognitively spontaneous beliefs violates (B). After all, the justificatory arguments schematized there would be fairly demanding. To characterize real agents accurately, the premises suggested would require an agent to notice and articulate often subtle differences in conditions making for reliability of judgments regarding different subject matters. While some actual agents might be able to produce such sophisticated arguments in some limited range of cases, it is palpable that, even within that range of cases (if it is at all significant), explicitly producing such arguments would be cumulatively intractable.

The simple model never was proposed as anything but an starting place that captured the central core of coherentism. Bonjour, for example, introduces it as a model of "fully explicit justification," not epistemological justification *simpliciter*. Accordingly, it is natural to

¹² *The Structure of Empirical Knowledge*, pp. 102, 152.

¹³ "The Unattainability of Coherence," in *The Current State of the Coherence Theory*, J. Bender, ed. (Cambridge: Kluwer, 1989).

think that moderating elaborations lie in the direction of dropping, with replacement, any suggestion that the indicated arguments must be explicit in justification. Finding a satisfactory replacement poses a particularly fascinating problem for coherentism, however. In moderating coherentist demands on justification, we cannot simply insist that we are espousing an ideal that will be satisfied if the agent “has on hand” the relevant beliefs, regardless of the agent’s cognitive processing or cognitive dispositions to mimic the processing called for. Such a “wimpy” coherentism would neglect the epistemological concern for *how* we can manage well our epistemological chores. The trick to be turned is then that of demanding neither too much nor too little. The only promising course seems to be that of building on Bonjour’s suggestion that the arguments indicated in the simple model must at least be *accessible* to an agent who is epistemically justified. If justification is not strictly a matter of what is done (tokened) cognitively, it could yet be a matter of what is done plus what further could be done.

Consider a *sophisticated coherentist model* that mandates (a) cognitive processing that implements in a limited fashion the simple model, backed by (b) dispositions to ever more nearly approximating it, were the matter pursued further. The central idea is that, while performance or tokening of ideally coherentist processing may be reined in for any number of situational reasons, still, as long as the dispositions (or underlying competence) giving rise to the particular processing themselves conform to the simple model, and thus would produce ever more “perfect” processing were the epistemic agenda pursued further, then all is fine. On this model, the simple coherentist model remains the central epistemic ideal or standard—only now it is an ideal standard for epistemic dispositions (that are always incompletely implemented), not for processing. This way of understanding coherentism is suggested by several (but not all) passages in Bonjour’s *The Structure of Empirical Knowledge*.¹⁴ The suggested sophisticated model is fully in keeping with my view that normative epistemology should give central place to a theory of competence. (Although this particular account of epistemic competence will ultimately turn out to be psychologically implausible in ways that violate the constraints on a proper theory of competence.)

¹⁴In one particularly notable passage, Bonjour seems to repudiate the sophisticated model to which he otherwise seems strongly attracted. He does so by suggesting that “elements” that are “tacit” or “implicit” in common justificatory contexts may not be accessible to the agent (*The Structure of Empirical Knowledge*, p. 152). But perhaps this explains why Bonjour would also occasionally despair of agents having more than “potential justification.”

IV

The sophisticated coherentist model is ultimately unacceptable. The dispositions that it recommends may approximate those which would be featured in the proper theory of epistemic competence. But, at best, they only approximate them by paralleling them up to a point. Ironically, this is *mutatis mutandis* what is commonly said in criticism of contextualist accounts. Coherentists (and foundationalists) commonly insist that contextualist standards merely codify the approximations to the ultimate epistemic ideas that one finds in practice. Contextualists are thought to confuse delineating ideals of competence with describing performance, with the result that they come to call for dispositions that merely approximate the more correct coherentist (or foundationalist) ideals. Coherentists thus come to view contextualism as what I term a poor person's coherentism. But, recalling constraints (A)-(C), we can see that it is sophisticated coherentism that cannot provide the proper ideals. If anything, contextualist standards are on the mark, and do not merely describe a poor person's approximate implementation of coherentist ideals.

My reasons for thinking that the simple coherentist model was unworkably demanding are, implicitly, general objections to the sophisticated coherentist model as well. For my central reservation was that the simple model prescribed practice that itself required dispositions that no human could have, violating constraints (B) and (C). Since the sophisticated model drops the demand that performance faithfully and fully implement dispositions that conform to the simple model, but retains the demand that the agent possess such dispositions, the earlier criticisms are directly applicable. This provides several compelling reasons for thinking that the sophisticated model violates (B) and (C), and is thus unacceptable. But I want to develop a distinct (but related) reason for thinking that it violates (C)—one that focuses on the viability of such coherentism as an account of our perceptual competence and empirical knowledge. My objection is not simply that the power and speed of human cognitive capacities are not up to the task of providing on demand the envisioned justificatory arguments for perceptual beliefs, although that is probably so, but that the processes that give rise to perceptual beliefs involve sorts of processing that do not render uniformly accessible the premises that would be needed for such argument. Our perceptual competence fits very poorly with the coherentist demands at this point.

It is important that epistemologists appreciate the fundamental processes that human beings can, do, and must employ in percep-

tion. I believe that once we focus on what really is the underlying competence involved in the production of cognitively spontaneous beliefs, we find that, while humans might occasionally be able to produce what look like (limited) coherence—theoretic justifications for such beliefs, often they are not able to do so. The reason for this inability is deeply rooted in human perceptual competence itself: it does not uniformly provide the articulable raw material that the coherence-theoretic model demands. The dispositions through which humans do and can admirably (even ideally) manage their perceptual affairs often leave no basis (or place) for articulating the coherence-theoretic arguments. There is no blemish in this for our empirical knowing. In effect, what we find is that a proper theory of human perceptual competence reveals that the sophisticated coherentist model violates constraint (C), and it also leaves us with a model that is recognizable as a version of contextualism.

Cognitive science provides two very general ways of coming to understand and explain the dispositional capacities of cognitive systems. On the one hand, there is a sort of analysis found in rather traditional artificial intelligence and much influential cognitive psychology. The sort of analysis I have in mind treats the cognitive system as characterizable in terms of a set of transformation rules stated over representations. To characterize and explain the sophisticated dispositions of the cognitive system, this approach would have us write a program that, in finer and finer terms, accounts for the system's working. The approach is given a particularly fine exposition by Robert Cummins's¹⁵ discussions of functional analysis.

Recently, cognitive science increasingly has come to focus on a second (sometimes complementary, sometimes competing) general approach: *connectionism*. Rather than analyze dispositions in terms of programs, connectionist cognitive science describes networks of a sort that are in certain limited, but arguably crucial respects, like our brains.¹⁶ The connectionist descriptions account for how networks of the sort in question manage tasks (recognizing a certain class of physical objects, for example). This is to account for the realization of a dispositional property of the connectionist system.

¹⁵ *The Nature of Psychological Explanation* (Cambridge: MIT, 1983).

¹⁶ For some perspectives on respects in which the networks connectionists describe are like, and differ from, our brain's networks, see Paul Churchland, "On the Nature of Theories: A Neurocomputational Perspective" and "On the Nature of Explanation: A PDP Approach," in his *The Neurocomputational Perspective* (Cambridge: MIT, 1989); and Paul Smolensky, "On the Proper Treatment of Connectionism," *Behavior and Brain Science* (1988): 1-74.

Connectionist *models* account for a given capacity by describing the components of the realizing system—the nodes in the network and the connections between them—and their capacities. The important capacity of a node is its tendency to take on certain activation levels, given certain levels of input coming through the connections it has with other nodes in the network. The central capacity of an individual connection is termed its “weight”; this is its disposition to pass along (in one direction) a more or less diminished signal between the two nodes it connects. Activation of various nodes in the input layer of a network will then spread through the net, as the activation of any individual node in the interior of the net results from (a) the sum of excitatory and inhibitory signals passed to it through connections and (b) its tendency to become activated given certain levels of incoming summed signals. Ultimately, and quickly, the result will be a pattern of activation in an output layer of the network. Here the input-output disposition of the full system is accounted for in terms of a relatively narrow range of dispositional features of components.¹⁷

There is a contrast to be drawn between many capacities best dealt with using connectionist theory and the capacities typically described in functional analyses. Functional analyses allow us to understand a capacity by decomposing the task performed into a set of simpler steps. On the other hand, networks often seem to realize their capacity in a single step, as the network simply settles into a pattern of activation in reaction to prompting input. It seems likely that one approach may be particularly useful in accounting for certain capacities (which are realized in a stepwise manner), while the other shines in accounting for different capacities (which are realized more immediately).¹⁸ Such contrasts should not overshadow the fact that both these analyses allow us to understand how systems manage certain tasks. Accordingly, they both provide apt bases from which to draw in constructing a theory of epistemic competence that will allow us to describe how we do, and prescribe how we could, manage cognitive tasks.

¹⁷Typical connectionist *theorizing* about the ability of models to perform tasks is a rather more abstract matter. From the point of view of connectionist models, connectionist theory anticipates robust multiple realizability of system-level tendencies.

¹⁸While Dan Lloyd—in “Leaping to Conclusions: Connectionism, Consciousness, and the Computational Mind,” in *Connectionism and the Philosophy of Mind*, pp. 444–59—probably overgeneralizes or overextends this idea of accomplishing cognitive tasks in a single step, his discussion provides one interesting development of the idea in connectionist terms. The idea, as the associated idea of sensitivity to unarticulated information, plays an important role in my view of perceptual competence.

The sorts of processing for which a connectionist analysis promises to be central within a theory of competence—those that seem to be managed in a single step—are diverse. Paul Smolensky characterizes them as residing in an “intuitive processor” that is “presumably responsible for . . . a huge portion of human behavior: Perception, practiced motor behavior, fluent linguistic behavior, intuition in problem solving and game playing—in short, practically all skilled performance” (*op. cit.*, p. 5). He notes that connectionism seems particularly fitted for representing the workings of this intuitive processor. What is important for our purposes is that among the activities and dispositions for which connectionism seems to provide particularly promising accounts are those Bonjour would describe as the production of cognitively spontaneous beliefs.

In a connectionist network, information is “stored in the weights”: rather than being stored in the form of tokens of syntactical sentences inscribed at some memory location and awaiting recall to the central processing unit, the learned information on which the network relies is “widely distributed,” being contained in the many weighted connections within that system. Learning then involves adjusting the connection weights. To employ one much discussed example, a very simple system may be trained to “perceive” explosive mines at sea (as distinct from rocks and various sea creatures) in response to sonar echoes. This will involve setting the connection weights so that the network automatically produces the desired (positive or negative) responses when given various sonar profiles as input. One can begin with a network whose weights are set at random. Of course, the output of such an untrained system is arbitrary and worthless—garbage. In training up such a system, it is exposed to a problem set of echoes and its “positive” and “negative” responses are compared to what is known to be the correct answer in each case. The system’s weights are then adjusted very slightly whenever it gives an incorrect response to a member of the problem set. Over time, the system learns to be quite a good perceiver of mines.

Such trained-up systems come to have a number of intriguing characteristics. The system comes to be able to perform well on cases that it has not previously encountered. (In our example, the simple network responds with impressive reliability to sonar profiles differing from those it encountered in training.) Thus, in some sense, the system generalizes beyond the samples that it is given in training. In doing this, the system apparently acquires both a set of paradigmatic cases (paradigmatic mine-representations) and an exquisitely sensi-

tive multidimensional similarity matrix.¹⁹ This bodes well for the suggestion that our own perceptual competence is to be understood in connectionist terms, for much categorization and perceptual recognition seems to turn on just such information. This is the sort of information which allows one to recognize objects of a given sort (pigs, for example) from various angles, in various postures, differing lighting conditions, and so forth. The characteristic flexibility of our recognitional capacities extends to our ability perceptually to recognize individuals and instances of sorts of objects from "presentations" that are far short of optimal (say when the pig perceived is partially obscured by trees, fences, farmers, or trucks). Again, such a facility is also characteristic of connectionist systems working on the basis just sketched: their performance in recognitional tasks degrades fairly gracefully as the input they are given is of a partial or degraded quality. Finally, the characteristic processing of connectionist systems would be quite rapid, as it is a direct function of the number of layers of nodes (not of the number of nodes) involved, the speed of signals across connections, and the time each node requires to react to input with a level of activation and output. A fairly minimal time would typically be required for the several layers of interconnected nodes to settle into a pattern of activation. Thinking of the nodes as neurons, we would then have some explanation for the rapidity of our recognitional experiences, and for how we may employ vast quantities of information so quickly and "automatically." These points highlight the promise of connectionist accounts of our perceptual competence.

Because connectionist systems rely on information that is distributed across the weighted connections of the network, and commonly nowhere encoded or represented in a token of some sentence,²⁰

¹⁹Talk of paradigms and similarity is familiar from studies of perception, see E. Smith and Douglas Medin, *Categories and Concepts* (Cambridge: Harvard, 1981). The sense in which one can write of paradigmatic cases and of similarity in the present context is articulated in Churchland, "On the Nature of Theories" and "On the Nature of Explanation."

²⁰There may be an important respect in which the distributed representations that are a system's pattern of activation may have a componential syntax. See Horgan and Tinson, "A Neoclassical Framework for Cognitive Science," *Synthese* (forthcoming), "Levels of Design in Neoclassical Cognitive Science," *Philosophy* (forthcoming), and *Connectionism and the Philosophy of Psychology*. Perhaps something similar can be said for the "standing information" stored in the weights of the system. Particularly in this latter case, however, it seems that this possibility should give little comfort to the coherentist. For while there may be a nonstandard respect in which there would then be "belief-sized sentence" stored (now distributively) within the system and informing perception, and while these may reflect a respect in which the system is sensitive to information of a sort that would also turn

such a system may rely on information in its processing without making that information “explicit” in any plausible sense. Accordingly, perception, or the generation of cognitively spontaneous beliefs, can proceed in a manner that is acutely sensitive to information that is never articulated.²¹ Indeed, this seems just what is likely to make for the “automatic” or “spontaneous” phenomenology of such productions. Further, these observations suggest something of great epistemological significance: a connectionist system can possess information that is automatically used in responding to input, yet not really “rehearsible” or evaluable by the system itself. While the system may employ information that is distributed across its network (along with much other information in the form of input) the system itself may have no capacity to “decompose” much of the information that it uses into discrete, belief-sized bits which might then be employed in the sort of coherence-theoretic vindication entertained above. If connectionism provides the best account of the human perceptual competence, as I think is likely, then it also is very likely that a well-trained human perceptual system may be attuned to “clues,” and in a very significant sense may deal with information, some of which it could not articulate.

out to be useful in coherentist justificatory argument, there is no reason to think that such information is available in an articulable form as would be required for the availability of such argumentation. As we have seen, the point of writing of “accessibility” seems to be to posit a disposition to argumentation. This seems to be intimately connected with the idea that the epistemic system would have various standing states (standing beliefs) to which it has internal access, that it can recall, transforming them into occurrent states (occurrent beliefs) that would then feature in processing (justificatory argument). There is no reason to think that the distributed representations relied on by the system in perceptual recognition could quite generally be made explicit in the sweeping fashion called for. Presumably this can sometimes be done, but there seem little reason to think that it is feasible for any given piece of “needed” information.

²¹ For a handle on how this is possible, consider: as the system is trained up, the weights on its various connections are adjusted, with the result that the system comes to be disposed to behave in (sometimes intricate) content appropriate ways. One can think of the learned information stored in the weights as giving rise to a particular dynamical system. The tendencies of such a dynamical system to evolve from one overall activation state to another can be conceived in terms of a topography for a very high-dimensional space (where each point on that “landscape” characterizes a distinct overall activation state of the system, and where the system tends to evolve from one state point to another as a ball rolling down the gradients of that landscape). The picture one then gets is of information distributed across the weights of the network automatically coming into play as the system settles, or evolves, into a particular state after having certain input. For less compressed discussion, see Horgan and Tienson, “Levels of Design in Neoclassical Cognitive Science” and “A Neoclassical Cognitive Science”; and Churchland, “On the Nature of Theories” and “On the Nature of Explanation.”

Consider the following scenario: you are driving along a familiar six-lane city street. You are in the middle lane of the three lanes going your direction. Your eyes are, of course, surveying the passing scene as you listen to National Public Radio's *Morning Edition*. You tend to foveate on traffic that might most immediately impact on your auto (vehicles in those adjacent lanes heading the same direction as you, those in the nearest oncoming lane, and those at driveways and intersections). (Although you may not be aware of this, for that matter.) Suddenly you come to generate a series of cognitively spontaneous beliefs. You believe that, several lanes of traffic to your left, two cars are involved in a traffic mishap. You believe that one car, a bluish-silver one, attempted to change lanes (or perhaps turn) and that it thereby intersected the course of a recent vintage van of some reddish-brownish color. The accident occurred at something like an eighty or eighty-five degree angle from your line of travel. You seemed to get a "good (but quick) look at it," however, as your attention was momentarily diverted its way by the sound of tires sliding across pavement. Because you knew that drivers who could negatively impact on your own vehicle might be distracted by the same events, and might act irresponsibly, you did not indulge in an extended visual examination of the scene—you drove on. A news story about underpaid philosophers comes on, your attention quickly shifts. You do not reflect on the new beliefs until later. All along, you feel confident that you got a good look at the events in question.

So, you believe that a recent vintage reddish/brownish van was involved. You also believe that a bluish car did the precipitating lane change. If pushed, you would probably recognize gaps in your explicit knowledge which would undermine any attempt to present a BonJour style vindication of the beliefs in question. For example: the fostering observations occurred in the early morning period when light levels are changing rapidly. You know that the amount of light could effect the reliability of your color perceptions. But you did not at the time entertain an articulate belief regarding the light level. (Had you considered the matter at the time, you could have simply made further observations regarding light level, to fill in, but you did not.) Your uncertainty regarding light levels is a matter of not having "really noticed," rather than having forgotten. Obviously, the reliability of your perceptual system in the light levels that did obtain is relevant to any argument regarding the likelihood of the veracity of your color beliefs. You know this much, and this renders the coherence-theoretic vindica-

tion inconclusive.²² But you do know that the color beliefs seem “comfortable,” whatever that means.

Reflecting on this and variant cases, I believe we can appreciate how our perceptual systems may have been sensitive to light levels without being able, in each case, to articulate information on which it relies. On other occasions you might seek to identify the color of an object, only to find that you automatically feel uncertain and readily attribute that to the inadequacy of the light for the discriminations sought. You might say, “it is perhaps a blue or a green, it seems something like that, but I really cannot tell because the light is so poor.” What this indicates is that, in the course of your various past color judgments with feedback, you have come to be a system that is sensitive to light levels even though you typically do not consciously notice light levels. Your information on light levels typically is unconscious, and in the ordinary cases is and remains inarticulate and inarticulable. (You are probably more inclined to notice—articulate information about—such factors when they frustrate your observations; that is, in the nonordinary cases.) You may then employ information regarding light levels, and a wealth of other factors, without either articulating that information, or being able to articulate it (in any form that would facilitate the production of coherentist justificatory argumentation).

Perhaps such inarticulate sensitivity is most striking in the way our perceptions accommodate partial obstructions, “filling in” what is obstructed or missed. In our example, you know that the extent to which your vision of the vehicles in question was obstructed by other vehicles is surely statistically relevant to the reliability of much that you came to believe. Further, there is certainly a sense in which your visual system had information on this matter. In important respects, it compensated for the obstructions as you settled into your perceptual recognition of the ill-timed lane change. But any judgment you make on the extent of obstruction is likely to be a hopeless reconstruction, at least if one demands the sort of detail that would be necessary for judgments of reliability.

Of course, in such cases, one could push to “recall just what one saw”—as if one were replaying from memory the record of the input,

²²You know that your observations were taken during a period in which a parameter crucial to reliability was undergoing rapid variation. Given what you have available for explicit argumentation, and what you know to be relevant, you could at best argue for the conclusion that your color beliefs are likely true, *provided* that the light levels were adequate. This is not to argue for the conclusion that the resulting cognitively spontaneous color beliefs are probably true.

now reviewing the vivacity of the colors and the exact obstructions. I suspect, however, that in important respects the “justificatory” arguments that are accessible to us subsequent to perceptual processing are like eyewitness reports of an event: when we push ourselves to recover information that was largely unconscious—inarticulate, but used by our perceptual system—we commonly get a convenient reconstruction. Such reconstructions may have notably little to do with what actually happened (in the event or processing). If so, this is rather damning to the coherentist arguments envisioned. For those arguments have to do with actual producing processes—with what happened in the inception of the relevant cognitively spontaneous belief—and with the parameters determining their reliability. If our arguments regarding such processing typically draw on a somewhat distorted cartoon reconstruction, perhaps a somewhat self-serving one, then they could have little justificatory force. The fact that we could make something up is of little help in epistemology.

The upshot of all this is that emerging results of cognitive science strongly indicate that our actual perceptual competence is to be understood in generally connectionist terms. Because such connectionist competence may, and here very plausibly does, allow us, as a cognitive system, to employ information in appropriate, sensitive ways without being able to articulate it in a form that would make it accessible for the sort of argument entertained in the sophisticated coherentist model, a crucial problem arises for sophisticated coherentism: the dispositions that enable our generation of perceptual beliefs—our actual perceptual competence—are such that there will very commonly be cognitively spontaneous beliefs that (even individually) cannot be provided the sort of justificatory argumentation whose accessibility is demanded in coherentism. This is also to say that any cognitive system that would need to generate its cognitively spontaneous beliefs in such a primarily connectionist manner (more particularly, any system that is limited to connectionist processing as the core of its perceptual processing, where that processing does not uniformly employ readily articulable beliefs regarding those various factors making for its own reliability), is a system whose dispositions—possible as well as actual—are incompatible with the sophisticated coherentist model as an ideal for its empirical knowing.

The sort of perceptual processing I have sketched above seems so deeply ingrained in the human organism that at most marginal modifications (or supplementation) seem feasible. Of course, we might be taught to notice certain limited features of the perceptual situation, so as to make information on these features more readily acces-

sible for subsequent justificatory argument. But, such *marginal modifications* in our perceptual processing will always remain partial in a fashion that does not vitiate the point already made: there will remain routine cases in everyday life for which the problem of justifying the individual cognitively spontaneous belief will be unsolvable on coherentist ideals due to the inability of the system to articulate the relevant information. If all that is possible are such marginal changes, then, in violation of (C), there are, within the set of problems with which the perceptual system would routinely need to deal, individual problems for which the coherentist processing cannot be implemented, even in ideal conditions.

The conclusion is worth emphasis: insofar as our actual perceptual competence is as described here, and insofar as only marginal modifications are feasible, then, by (C), as well as by (B), the sophisticated coherentist model is not a proper ideal for human cognizers. That is, we should not use the sophisticated coherentist model as our theory of ideal epistemic competence.²³

v

It is time to state clearly just how the above discussion is to provide an argument for contextualist epistemology. To begin with, it should be clear that it does not, at least as it stands, vindicate all the claims associated with contextualism (even in the limited form set out in my initial section), nor is it intended to do so. Rather, it is intended to complement other discussions that point to a contextualist position.²⁴ It complements these presentations by focusing on contextu-

²³ It is worth noting that, even were we to entertain, for purposes of argument, the view that human cognition is more flexible than I have suggested, there remain reasons to question the epistemic desirability of available alternative sorts of processing. The place for such connectionist processing in ideal competence should be evident when we recall that such processing accounts for epistemically important features of our perceptual systems: their speed, ability to deal with novel cases, graceful degradation in the face of degraded input, ready (phenomenologically automatic) implicit use of relevant information. These features are crucial to the efficient and reliable generation of cognitively spontaneous beliefs, on which our empirical epistemic life depends and seemingly cannot be duplicated in a radically alternative fashion within human cognitive systems.

²⁴ Prominent here are Williams, "Coherence, Justification, and Truth" and *Groundless Belief*; the obvious writings of Ludwig Wittgenstein, *Philosophical Investigations*, G. E. M. Anscombe, trans. (New York: Macmillan, 1958) and *On Certainty*, Denis Paul and G. E. M. Anscombe, trans. (New York: Harper and Row, 1969); and pieces such as David Annis, "A Contextualist Theory of Epistemic Justification," *American Philosophical Quarterly*, xv (July 1978): 213-9. I would also emphasize, however, the contextualist thrust of several other works from mainstream analytic philosophy. For example, W. V. Quine's use of Otto Neurath's figure of piecemeal boat repair and his associated suggestion that part of our webs of belief are central in a way that generally puts them outside of run of the mill epistemic suspicion should clearly blunt the general coherentist direction of his thought in a fash-

alist elements in our handling of that hallmark of empirical knowledge: perceptual beliefs. In effect, it supports contextualist epistemologies by indicating respects in which our epistemic competence in regard to perception seems best understood in contextualist terms. In a naive way, foundationalism has traditionally seemed to hold out the most promise for accounting for the epistemic status of perceptual beliefs. Coherentists such as Bonjour, however, have presented decisive objections to foundationalism. Still, they have recognized that any epistemology must ultimately account for *empirical* knowledge, and that, to do this, the epistemological approach must generate an account of the epistemic place for perceptual beliefs. Now, I have argued that coherentism is inadequate to this task.²⁵ Accordingly, coherentism as well as foundationalism is unacceptable.

The positive side to my argument in this paper is that, unlike the other two approaches, contextualism can embrace the connectionist account of perceptual competence that is emerging from recent cognitive science. It can accommodate this model as an epistemic standard (or ideal) at two levels: both as a model of our actual epistemic-perceptual competence and as the basis for a model of fully ideal human epistemic-perceptual competence. Thus, in comparison with foundationalism and coherentism, contextualism is better suited to providing an acceptable epistemology of empirical knowledge.

Let me emphasize the contextualist elements in the perceptual competence outlined above. To begin with, it is central to contextualism that epistemic justification is always circumscribed by a presupposed background that is properly taken for granted, not because it has already received articulate justification itself, but because there is no epistemic basis for, nor desirability to, proceeding further in the particular case. As we have seen, perceptual competence seems informed by much information that, in many common cases, cannot be articulated. In such “run of the mill” cases, there is obviously no basis for challenging what cannot be articulated. Further, it seems

ion that is congenial to contextualism. Kitcher’s discussion (in “The Naturalist Returns”) of the epistemic need to rely on unquestioned parts of our bodies of handed-down knowledge also is congenial. Such discussions tend to feature the more global contextualist aspect of our epistemic practice, while I seek to focus on contextualist aspects of our epistemic practice with respect to perception.

²⁵At least this holds insofar as the coherentist epistemology conceives of justification in terms of the accessibility of justificatory reasons and arguments—that is, insofar as justification, in the traditional epistemological fashion, is taken as characterizable in terms of articulable argument. This qualification, and the alternatives to which it points, are developed in my “One Naturalized Epistemological Argument against Coherentist Accounts of Empirical Knowledge,” unpublished MS.

that the cognitive “mill” in these cases can be working ideally—that is, according to an ideal epistemological competence for perceptual processing. From the point of view of appropriate epistemic ideals, then, there is no blemish in the unexamined, and unexaminable, reliance on this information. To demand more than this contextualist resting place (even the availability of more) is to violate constraint (C) on proper epistemic ideals.

It is strictly more appropriate, however, to conceive of contextualism as presenting a model of the ideal form of articulate/articulate justification. The observations in the preceding paragraph can be recast in keeping with this aspect of contextualism; we then concentrate on the stopping places in articulate justification—as contextualists say, (linear) justification just comes to an end in a class of beliefs that are not (in the particular cases) subject to, in need of, or epistemically improved by, further justification. Nor are they “self-justified.” In our account of perceptual competence, simple perceptual beliefs *can* serve this role. (But this is not to say that all such simple perceptual beliefs do or can. Much depends on the particular case—as sketched below.) Of course, the generation of such beliefs can rely on much important information. But, insofar as it is commonly inarticulate information, any standard demanding that argumentative vindication in a context of justification be accessible is inappropriate. In such cases, then, such simple perceptual beliefs commonly serve as contextually appropriate stopping places, neither needing, nor susceptible to, epistemic improvement.

A second hallmark of contextualist epistemology is the insistence that the proper resting places in the context of epistemic justification are a diverse lot. When we think of epistemological justification being pursued ideally far, to the ideal exhaustion of epistemic appeals, we might say, the contextualist believes that we arrive at a nonhomogeneous set of beliefs that neither need nor can be justified further. (In contrast, the foundationalist envisioned a homogeneous set of basic beliefs—homogeneous with respect to some intrinsic characteristic that is what makes them somehow already justified, and thus regress stoppers. On the other hand, the coherentist thinks that the justificatory task will burgeon to encompass the subject’s entire set of beliefs considered in a decidedly nonlinear fashion.) When we reflect on the articulate (or articulable) beliefs that serve in individual cases as stopping places in the justificatory contexts, as this is sketched above for cases of perceptual beliefs, we find agreement with the contextualist model. In one case, the perceptual belief that a blue vehicle changed lanes might

serve as a stopping place, subject to no further justification. In another case, a token of the same belief type might be held justificatory hostage until supporting beliefs are obtained or marshaled. Perceptual beliefs are sometimes stoppers, sometimes not. Sometimes the articulate and articulable basis for judgment will end with a given perceptual belief, in other cases, our perceptual system will throw up reservations—further justificatory problems. “The car looks to be blue,” we may think, “but the light makes it hard to tell.” A quick glance at familiar objects ready to hand may reassure us of our discriminative ability in the conditions at hand. (Or it may not.) This would be to rely on a rather diverse set of beliefs—as coherentists would be quick to detail for us.

DAVID K. HENDERSON

University of Memphis