

EPISTEMIC NORMS AND THE SELLARSIAN DILEMMA FOR FOUNDATIONALISM¹

Joe Cruz
Williams College

1.

Foundationalists and coherentists disagree over the structure of the part of the mental state corpus that is relevant for epistemic achievement (Bonjour, 1999; Dancy, 1989; Haack, 1993; Sosa, 1980; Pollock and Cruz, 1999). Given the goals of a theory of epistemic justification and the trajectory of the debate over the last three decades, it is not difficult to see how structural questions possess a kind of immediacy. In order to undertake an epistemic evaluation of a belief, one intuitive and appealing strategy is to investigate the reasons for that belief to determine whether it is epistemically positive, where the reasons are typically other beliefs. This demands that we must in turn determine whether the reasons for the belief are themselves justified. A regress looms (and thus a regress argument is in the making), and foundationalism and coherentism propose proprietary views on the structural relations between beliefs with an eye toward resolving it.

The foundationalist claims that there is a set of basic beliefs — or, at any rate, basic cognitive states — which do not require reasons to explain their epistemically positive nature because of some special characteristic(s) that they have. The epistemic credentials of beliefs that are not foundational are due to a traceable lineage from basic beliefs via a basing relation that must be illuminated by the foundationalist. Thus, the

structure proposed includes relations that are asymmetrical, with some beliefs or cognitive states providing epistemic support for other beliefs without themselves needing epistemic support.

In contrast to foundationalists, coherentists maintain that no beliefs are by themselves epistemically positive. By their lights, every belief relies in some way on other beliefs for its epistemic status. On one reading of coherentism, beliefs are epistemically positive based on a lineage in a structure that may ultimately loop back onto itself. There may be no need, however, to trace reasons in a way that is circular. One might instead claim that a belief is epistemically positive in case it is a member of a coherent belief corpus without pursuing particular reasons in a linear fashion. What is crucial about the coherentist proposal about structure is that it rejects the claim that any relations are asymmetrical.

The distance between foundationalism and coherentism should not be overstated. Every credible foundationalist proposal makes room for the importance of broadly speaking coherence relations², and some familiar coherentist views take measures to make sure that non-perceptual beliefs do not always overwhelm perceptual beliefs.³ Still, foundationalism and coherentism are at least rivals with respect to what I will call *the structural question*.

I presume that this story as it stands is a complete commonplace, so I will not revisit its details even though there are many subtleties and complications.⁴

¹ I thank Yves Bouchard, Tomoji Shogenji, and Jonathan Vogel for valuable comments on this paper.

² These relations are sometimes handled through a theory of epistemic defeat, where beliefs can serve to defeat one another, thus pushing the entire structure in the direction of coherence.

³ Recall Bonjour's *cognitively spontaneous beliefs* in his (1985).

⁴ For instance this portrayal leaves out any discussion of *infinetism*, namely the possibility that the regress

2.

The structural question must be answered in order for an internalist theory of justification to be successful.⁵ Once the structural question is highlighted through the regress argument, our next step is to decide how to go about answering it. In a straightforward sense, answering the structural question will demand a mature foundationalist or coherentist theory, as only these can ultimately sway us to one account or the other. Let us step back, though, from particular foundationalist or coherence relations, and coherentist theories and ask ourselves — so to speak — where we need to look for an answer to the structural question.

The standard approach involves investigating the mental states possessed by the epistemic agent for whether they plausibly suggest and sustain either a foundationalist or coherentist structure. Since the crucial point of disagreement between the two views seems to be whether there are candidates for structurally asymmetric beliefs or cognitive states, it seems reasonable to search carefully for these. Put this way, and in spite of its historical pedigree and intuitive appeal, there is immediately some reason to worry that foundationalism faces difficulties. It is claiming something surprising and extraordinary about some mental states, namely that they can confer justification but are not in need of it. The potentially crippling problem with foundationalism is that it postulates mental states that require a mysteriously dual nature. And it seems completely up for grabs

continues on infinitely. While that option has sometimes been regarded as hopeless, some prominent defenses of it can be found (e.g., Klein, 1999).

⁵ I am presupposing a minimal kind of internalism here. On well-known externalist theories of justification, the controversy between foundationalism and coherentism does not appear with the same drama that it does on an internalist view. Some have argued that externalism can be thought of as a kind of foundationalism (Sosa, 1980), while others have claimed that reliabilism, at least, is compatible with either foundationalism

whether there are any of these. To put the issue as Bonjour does in *The Structure of Empirical Knowledge*,

If ... intuitions or direct awarenesses or immediate apprehensions are construed as cognitive, at least quasi judgmental..., then they will be both capable of providing justification for other cognitive states and in need of it themselves; but if they are construed as noncognitive, nonjudgmental, then while they will not themselves need justification, they will also be incapable of giving it (1985, p. 69).

In an earlier paper Bonjour is not nearly so gentle. He characterizes the foundationalist's structurally asymmetrical states as "contrived and ad hoc," and points out that the basic move, if acceptable, would solve a range of outstanding philosophical problems: "...Thus we would have semi-events, which could cause but need not be caused; semi-explanata, which could explain but need not be explained; and semi-beliefs, which could justify but need not be justified" (1978, reprinted in Sosa and Kim, 2000, p. 270). Patently, semi-events and semi-explanata are unacceptable, so why, Bonjour asks, should we accept the possibility of structurally asymmetrical cognitive states?

More than anyone, perhaps, it was Sellars who drew epistemologists' attention to this difficulty, and I will follow others in calling it the Sellarsian dilemma for foundationalism (Pollock, 2000). There can be no serious doubt that the Sellarsian dilemma has crystallized a central discussion in the theory of justification, and as a result the foundationalism/coherentism debate from the early 1970's on is extremely instructive. The Sellarsian dilemma is as clear as it is pressing. In light of its power, foundationalists have heroically sought to describe some mental state that fits the bill.⁶ To the degree that one is skeptical about that project, coherentism looks to be an attractive option. Indeed,

or coherentism (Goldman, 1986).

⁶ This is a project that is being currently pursued. Bonjour - who has given up his earlier allegiance to coherentism - argues that the content of a conscious mental state is known infallibly, and can be used to

this may be understating things. Some philosophers have seemed to see the Sellarsian dilemma as a decisive blow to foundationalism and as the primary positive reason for coherentism.⁷ This makes coherentism a fundamentally reactionary strategy that is invoked in light of the failure of foundationalism to plausibly halt the regress. We then have a dialectical situation where the apparent presumption is in favor of foundationalism if some way can be found to solve the Sellarsian dilemma. I cautiously follow this line here in thinking that if we can blunt the dilemma, coherentism becomes otiose. I say ‘cautiously’, however, because the work of some coherentist — Lehrer (1990) and Thagard (2001) perhaps — has at times not so much seemed a reaction to perceived failures of foundationalism but rather as originating in considerations that have an independent plausibility. If some coherentists are right to think that their view possesses an advantage over foundationalism in some other respect than not facing the Sellarsian dilemma, I will accept a modest characterization of my project here as that of attempting to show that coherentist approaches that do rely on the Sellarsian dilemma to defeat foundationalism gain no foothold.⁸ I maintain that the Sellarsian dilemma can be met.

3.

One immediate challenge is sorting out the status of the structural question vis-à-vis what looks to be a *different* issue, namely whether there can be cognitive states that are justifiers not in need of justification. The close connection between these issues is a

justify other beliefs. This is closely related to Fumerton's view.

⁷ One has to add that infinitism and skepticism are intolerable alternatives, and that those two combined with foundationalism and coherentism exhaust the possibilities with respect to the structural question.

⁸ I actually subscribe to a stronger thesis because I would be surprised if foundationalism could not achieve all the things that coherentism does. That is an argument for another time, though.

result of the way the foundationalism versus coherentism debate usually proceeds. Once the structural question is made clear via the regress argument, the search begins for the cognitive states that might address the Sellarsian dilemma. In a way, the structural problem disappears, or is at least deferred, in favor of a new challenge.

Perhaps this is as it should be. After all, deflecting the regress argument does not *by itself* fully answer the question of what makes a belief justified or rational. An answer to the structural question cannot be a complete theory of epistemically positive beliefs because it will not offer insight into what abstract property or properties must be had by a cognitive state in order for it to (1) transmit justification and (2) to transmit justification for a *particular belief*. These topics are as fundamental to a theory of justification as issues concerning the structural relations between cognitive states relevant to epistemic achievement. Indeed, once the structural question gives way to the Sellarsian dilemma, (1) and (2) above can be properly recognized. Sellars and others are, of course, arguing that the foundationalist is not going to find non-ad hoc cognitive states that shed light on (1) and (2).⁹

My claim is that not all of the avenues with respect to the structural question have been properly investigated. In my view, the particular pattern of investigation found in the literature has tended to put foundationalism unnecessarily on the defensive. Notice again that the structural question has traditionally invited scrutiny of the mental states implicated in epistemic achievements. I called this the standard approach above, but in order to relieve it of some its polemical advantage and in order to give it a more instructive label, let us dub it the *mental state approach*. The mental state approach has

⁹ For an articulate summary of why prominent foundationalist proposals seem ad hoc, see Michael Williams (2004).

an aura of inevitability. What else might there be to do than to find out whether there are mental states that have a dual nature and thereby provide a resource for a structurally asymmetrical foundationalism to escape the Sellarsian dilemma? Here is my suggestion. Suppose we attempt to independently certify kinds of inferences that intuitively generate justified beliefs *without an antecedent commitment as to what mental states can figure in those inferences*. On this view, if the inferences - or, as I prefer, the norms - are ones that seem to us to yield justified beliefs, then *whatever* mental states appear in those norms will be appropriate components of an epistemological theory. The thought is that we can change the priority of assessment of the components of a theory of justification that are supposed to shed light on the structural question. If we let the norms direct the inquiry - and let us call this the *norm approach* - the foundationalist need not find herself crippled by the Sellarsian dilemma.¹⁰

While perhaps surprising, the possibility of an alternative to the mental state approach will not, I hope, arouse undue suspicion. After all, when investigating structure, in general it is possible to investigate either the nodes or entities in the structure, or the connections that give shape to the structure. Either one will tell you something important about the nature of the structure in question, but privileging one approach over the other may create mysteries where there are none to be had.

Think about the debate over psychological explanation in the wake of the collapse of Behaviorism. Identity theories seemed attractive in light of their materialist credentials, but foundered in part because the putatively identical mental state/physical

¹⁰ Something like this view may implicitly be at the root of Pollock's work on these issues. This would explain why he has never shown any interest in tackling the Sellarsian dilemma head-on. At any rate, this is how I now understand what we were up to in Pollock and Cruz (1999).

state types were too narrow to appear in plausible psychological generalizations, say, across species. Functionalism seemed, and to many still seems, like a good idea because it turned attention away from the pesky mental states and to the information processing relations without much regard for the vehicles of the process. This seems to be an instance where relations in a structure (in this case, psychological explanations) prevailed over nodes in the structure, at least with respect to the theoretical goals of the philosophy of psychology.

We would not want to take this analogy too far, but I contend that the structural question in the foundationalism/coherentism debate exhibits a parallel sensitivity to the approach. More specifically, endorsing the norm approach and concentrating on epistemic norms may show that certain kinds of mental states simply must and do figure in the formation of justified belief. There need not be any danger that those mental states will have occult, dual natures. It is the norms that bridge antecedent mental states to beliefs, and their independent certification will relieve any air of mystery. If I am right, the mental state approach to the debate has claimed an illicit privilege, and has thereby given an undeserved tone of finality to the coherentist's argument against foundationalism. To be sure, more needs to be said about the norm approach to the structural question before any of this is at all plausible. So I turn to a discussion of my conception of epistemic norms.

3.

What are epistemic norms? This much is clear: Epistemic norms do not merely describe the procedures that people actually use in belief formation. A faithful description

of those procedures originating in psychology or sociology may be relevant to our understanding of epistemic norms, but they apparently cannot tell us how we ought to form our beliefs. As a first pass, we may say that epistemic norms specify how epistemic agents ought to manage their intellectual lives.

Beyond this banal starting point, things are far murkier. Epistemologists apparently craft epistemic norms in order to describe good inference. In a reasonably non-committal sense, then, I take epistemic norms to be recipes that capture a peculiarly epistemic propriety between cognitive states. Consider some concrete instances: Some epistemologists have proposed that there is an epistemic norm regarding the good inferential relationship between perceptual states (or beliefs about perceptual states) and beliefs about the way the world is. One way of putting this norm is, “If a person has a clear sensory impression that *x is F* (or of *x*'s being *F*) and on that basis believes that *x is F*, then this belief is *prima facie* justified” (Audi, 2001, p. 43). Another way is: “If *S* believes that he perceives something to have a certain property *F*, then the proposition that he does perceive something to be *F*, as well as the proposition that there is something that is *F*, is one that is reasonable for *S*” (Chisholm, 1966, p. 47). A third is: “Having a percept at time *t* with the content *P* is a defeasible reason for the cognizer to believe *P-at-t*” (Pollock and Cruz, 1999, p. 201).

Other formulations of this norm vary widely (cf., Bonjour, 2000, p. 30; Plantinga, 1993, p. 99; Sosa, 1980, p. 22), but they all seem to be after the same thing, namely the claim that sensing that *p* is a fallible but good basis for an inference to the belief that *p*. Incidentally, there is not enough material to generate the Sellarsian dilemma yet, since the norms described above do not claim that clear sensory perceptions or beliefs about

perceptions or percepts are foundational. That is as it should be, as I am here merely attempting to get clear on what norms are, and this should be possible in a way that is independent of the foundationalism and coherentism debate. If the norms did claim that percepts are foundational, and if we were following the mental state approach, then Sellarsian worries would surge.

Norms having to do with perception are not the only kind found in the contemporary epistemological literature. Similarly formulated recipes with respect to, for instance, memory abound. The following is a plausible epistemic norm having to do with memory: when it appears to S that she remembers that p, then S is (highly defeasibly) justified in believing that p (cf., Audi, 2000; Pollock and Cruz, 1999). One can also find discussion of inductive norms. Induction is a particularly interesting case, because the norms involved in it are not so simple to state as the ones involved in perception and memory. Goodman taught us that we need to include a projectibility constraint in our inductive norms, so no recipe of the form *discovering that most F's are G is a prima facie reason for think that all F's are G* will be adequate. This goes to show that the conditions that must be met for a justified belief will in some cases be fairly complicated.¹¹

We can isolate a generic schema for thinking about epistemic norms. The epistemologist claims that inferences are apropos when a mental state ψ_n (a belief) is based on the content of prior mental states $\psi_1 \dots \psi_{(n-1)}$. For empirical knowledge,

Being in mental states $\psi_1, \dots, \psi_{(n-1)}$ is a good (but defeasible) basis for an inference to

¹¹ This observation is the basis of the liaison between epistemology and artificial intelligence research in Pollock's work (Pollock, 1995; Pollock and Cruz, 1999). Since epistemologists will sometimes be incapable of following the interactions between complicated norms, it is useful to encode the norms in an artificial reasoner. The epistemologist will then have a 'summary' of the interaction of the norms that can be

mental state ψ_n in the cognizer.

The epistemologist is relying on the possibility of filling this schema for all the norms that we possess. We should not underestimate just how challenging this project of articulating norms will be. Under my reinterpretation, a central task *of* epistemology is giving a plausible set of epistemic norms.

Some epistemic norms might be single-step recipes, and will therefore fill out the conditions for forming a particular belief with a single ψ on the left-hand side of the generic norm schema. I understand these one-step norms as the limiting case of the application of norms involving many steps that are justification preserving. When we move away from mundane cases of justified belief to rarified cases involved in, say, scientific inference, the epistemic norms governing a conclusion may be extremely complex and involve a long chain of justification-preserving steps.

Epistemic norms need not be causally or consciously explicit in the epistemic agent. When epistemic norms are formulated, they are explicit to the philosopher doing the formulating and they may or may not have a psychological reality in the epistemic agent. This leaves open the possibility that the underlying realization mechanism is the system of causes that can be understood normatively, which would be a kind of normative naturalism (one that I am sympathetic to. cf., Pollock & Cruz, 1999). Perhaps we would then be inclined to say that the epistemic norms supervene on the epistemic agent's psychology, but I take that to be a more partisan commitment than is necessary for our purposes. Note, too, that the schema does not identify a relation that is *merely*

assessed for intuitive reasonableness. For the specific discussion of induction, see pp. 234-238.

causal. If ψ_n were merely causally related to $\psi_1, \dots, \psi_{(n-1)}$, norms would be instances of causal (perhaps counterfactual-supporting) analyses of the relations between mental states. This would leave out any role for the contents of the various mental states involved in the inference. Presumably it is the contents of the mental states that persuade us that the norm yield justified beliefs.

A final point that should be made about epistemic norms is that they express steps that are epistemically permissible rather than obligatory. An epistemic agent is not required to draw all the justified conclusions on the basis of her mental states that are consistent with her epistemic norms. In this respect (and perhaps in others), epistemic norms are different from moral norms.¹² The sense of permissibility at work here does have an element of the claim that epistemic reasoning is ‘optional’, but it should be kept in mind that, as epistemic norms, epistemic agents ought to reason in accordance with them. An instrumentalist account of normativity might be invoked here to claim that, if justified beliefs are an agent’s goal, then she should reason according to her norms. Categorical accounts would, naturally, offer a different spin on the origin of the normative force of epistemic norms.

How can this account of norms be normative? In order to answer this question, one would apparently want to be clear on what is distinctive about the normative. One way this distinctiveness has been pursued is to treat normativity as about the possibility of error. In order for a recipe of any sort to be normative, it must be possible to go wrong. When we borrow a conception of normativity from moral discussions (especially Kantian

¹² This might not be entirely correct. Perhaps epistemic agents are required to draw especially obvious conclusions that are consistent with their norms; this might be the point behind discussions of epistemic responsibility.

accounts), we are apt to think that departing from the norm is possible only as a result of the rational will. But that is not what is fundamental to a norm. What is fundamental is that there can be departures from it. The possibility of departure does not specify the reason for the departure. Employing one's rational will to depart from a norm may be one way to do it, but there are others. For instance, one may depart from a norm due to exceeding cognitive resource constraints. I can fully understand the norm of good long division. After all, I can describe it: *Inspect the first digit of the dividend and determine how many times the divisor fits into it. If it does not fit in at all, inspect the first two digits (treating them as a single number now) and determine how many times the divisor fits into it. Write down that value, and keep track of any remainder.* And so on. This seems to be a norm that I possess *par excellence*. But when I am faced with performing long division in my head, I may very well run up against resource constraints that prevent me from following the norm. Indeed, it is simple to induce departures from the norm in me. Just offer me a dividend with enough digits, and I will tend to forget the running tally of numbers that I achieve in performing the division. Allowing me the use of pencil and paper will make this effect go away, again confirming that I possess the norm. Or, to put it in a way that preserves our theoretical neutrality, it confirms that I may be described by that norm.

What is meant, then, when it is said that I ought to reason in accordance with my norm for long division? In the view on offer here, what is meant is that that norm has command over the goal of division, i.e., it specifies what doing division entails. It is not, of course, the only norm that comprehends division. The hope is that this account will generalize for epistemic norms.

5.

We are now in a position to connect this account of norms to my alleged alternative conception of the foundationalism/coherentism debate. What is the difference between foundationalism and coherentism when we take the norm approach? It looks like coherentism is the view that all mental states that appear among the antecedents in an epistemic norm must themselves also appear as the consequents of an epistemic norm.

This seems like a dramatically undermotivated stipulation. To see why, consider: Epistemic norms are belief-forming recipes that yield justified beliefs. Thus, the intuitive epistemic credentials of the output identify epistemic norms. In this way, we connect our account of norms directly to the central methodology of epistemology. We identify an epistemic norm by focusing on the outcome of a recipe for belief formation. If the outcome is non-accidentally epistemically laudable, then we have before us a plausible candidate for an epistemic norm. There is nothing about these recipes that stipulates that they only take beliefs as ingredients. The antecedent mental states will not necessarily be beliefs in part because the epistemologists will be as permissive as her intuition demands in her reflections on epistemically normative relations. So, whereas the mental state approach faced the Sellarsian dilemma, the normative approach does not.

But this seems too simple, and perhaps the philosophical equivalent of a parlour trick. I am claiming that the fundamental methodology of epistemology certifies epistemic norms on the basis of the epistemic propriety of their output given a certain input. Could the methodology be so neutral as to the nature of the input, and so cavalier in the face of the mysteries of the Sellarsian sort that surround the input? I do maintain

that the certification of norms is neutral with the respect to the input in the sense that, though it must be a mental state bearing content, it need not be a *belief*. Now, the crucial issue is whether, in order to be contentful, the input is in need of justification. If it is, we are lapsed into the Sellarsian dilemma. I will return to this issue in a moment, but before I do I wish to highlight something else about the approach I am advocating.

Traditionally, the generic schema for norms has been driven by the view that there is something about the antecedent mental states that makes them plausible justifiers for the consequent mental state. For instance, Sellars (1975) talks about *transmitting reasonableness* as a capacity that a mental state has to have in order to figure in epistemic achievements. The idea is that those cognitively ‘downstream’ contentful states inherit the epistemic laudableness of the ‘upstream’ contentful states. A natural question that arises here is just how mental states are able to do this. Again, the way that the Sellarsian dilemma has been framed seems predicated on thinking that it is the mental states that have something that can be transmitted or inherited or caught (in the contagion sense of ‘caught’) by other mental states. On my view, this question is misguided. The property of a justified belief that is important for justification has to do with the epistemic norm involved in its production, not with the nature of the belief itself. What is crucial is how the cognitive agent arrived at the belief, not anything about the belief.

Thinking otherwise is to confuse two separate issues. One issue is the matter of justification, proper. This is the central concern of the recent history of epistemology. The other issue is *content inheritance*. By content inheritance I have in mind the process through which a downstream mental states takes on aspects of the content of the upstream mental states that produce it. For example, the appearance state (if there should be such

states) that there is a red patch before me might well lead to the belief state that there is a red patch before me. The two states share content. Indeed, it is irresistible to think that the belief state got its content from the appearance state. How this happened and how to understand the general case of content inheritance is an interesting issue, but the important thing for my purposes here is that content inheritance has nothing to do with whether the inheritor state is justified. The question of justification lies, I urge, along a different vector, namely one that implicates the epistemic norm that generated the inheritor state. I submit that proponents of the Sellarsian dilemma have mistakenly insisted on a single theoretical quantity to account both for content inheritance and justification. If one were persuaded that that must be the case, then a Sellarsian suspicion might seem reasonable: How could a state with such a such a content — with which its justification is inextricably bound — confer its epistemic properties without being in need of justification itself?

But if justification is understood as a property had in light of being the product of a genuine epistemic norm, then it is the norms that we must focus on, not the states. To be sure, several central mysteries remain. We are owed a theory of content inheritance and broader view of the capacity of norms to generate justified beliefs. I do not, of course, offer these here. My point is simply that the Sellarsian dilemma is not *inevitable* for foundationalism. It was the Sellarsian dilemma that made foundationalism seem ad hoc in its postulation of dual-nature mental states, however. On the normative approach, it is now coherentism that seems ad hoc and mysterious in its stipulation regarding the allowable antecedents in epistemic norms.

References

- Audi, R. (2001) *The Architecture of Reason*. New York: Oxford UP.
- Bonjour, L. (1999) The Dialectic of Foundationalism and Coherentism. In *The Blackwell Guide to Epistemology*, J. Greco and E. Sosa, eds. Malden: Blackwell.
- Cruz, J., Pollock, J. (2004) The Chimerical Appeal of Epistemic Externalism. In *The Externalist Challenge*, Schantz, R. (ed.). Dordrecht: de Gruyter.
- Chisholm, R. (1966) *Theory of Knowledge*. Englewood Cliffs: Prentice-Hall, Inc.
- Goldman, A. (1986) *Epistemology and Cognition*. Cambridge, MA: Harvard UP.
- Haack, S. (1993) *Evidence and Inquiry*. Oxford: Blackwell.
- Lehrer, K. (1990) *Theory of Knowledge, 2nd edition*. Boulder: Westview Press.
- Pollock, J. (1995) *Cognitive Carpentry*. Cambridge, MA: MIT Press.
- Pollock, J. (2000) Comments on Fumerton and Bonjour. In *The Return of Old Fashioned Foundationalism*, DePaul, M., ed. Lanham: Rowman & Littlefield.
- Pollock, J. and Cruz, J. (1999) *Contemporary Theories of Knowledge, 2nd edition*. Lanham: Rowman & Littlefield.
- Sellars, W. (1975) Epistemic Principles. In *Action, Knowledge, and Reality*, H. Castaneda, ed. Indianapolis: Bobbs-Merrill, pp. 332-349.
- Sosa, E. (1980) The Raft and the Pyramid. In *Midwest Studies in Philosophy Vol. V..Epistemology*, eds.
- Thagard, P. (2001) *Coherentism in Thought and Action*. Cambridge, MA: MIT Press.