

Contemporary Debates in Epistemology ed. Matthias Steup and Ernest Sosa. Boston: Blackwell, 2005, pp. 156-167.

**NON-FOUNDATIONALIST EPISTEMOLOGY:
HOLISM, COHERENCE, AND TENABILITY**

Catherine Z. Elgin

Much epistemology assumes that cognitive success consists in knowledge, where knowledge is justified or reliable true belief. On this conception, since propositions are the contents of beliefs and the bearers of truth values, they are what is known. If this is right, the sort of justification of interest to epistemology seems to be the justification of individual propositions. A linear model of justification is almost inevitable. To justify a given proposition is either to infer it from already justified propositions or to show how belief in it emerges from reliable belief forming mechanisms. *S* is justified in believing *p* on the basis of *q*, and *q* on the basis of *r*, and so on. Holists contend that this picture is misleading. They maintain that epistemic acceptability is, in the first instance, acceptability of a fairly comprehensive system of thought, comprised of mutually supportive commitments. The priority in question is epistemological, not historical. There is no contention that people come to believe a theory before coming to believe the various claims that comprise it. The point is that regardless of the order in which they are acquired, claims are justified only when they coalesce to constitute a tenable system of thought. The acceptability of individual sentences, as well as methods and standards, is derivative, stemming from their role in a tenable system.

The challenge for such an epistemology is to explain how systematic interconnections give rise to justification, how the fact that deliverances dovetail affords reason to believe they are true. Some philosophers hold that the coherence of a sufficiently comprehensive constellation of claims makes them true [Blanshard, 1939; Rescher, 1973]. This strikes me as implausible, but I will not argue against it here. The position I want to investigate is that coherence is the source of

epistemic justification, not the ground of truth. But if truth is independent of what we believe, why should mutual accord among our beliefs be indicative of truth? What is the connection? To avoid begging questions, it is perhaps better to begin by focussing not on the justification of beliefs, but on the justification of deliverances, these being representations that present themselves as candidates for belief. If we are concerned with justification, we should not limit ourselves to assessing the status of what we actually believe, but ask which of the things that could in given circumstances be believed should in those circumstances be believed. Deliverances, as I use the term, include perceptual inputs, fixed or transient beliefs, passing thoughts, and so forth.

Perhaps things will become clearer if we consider a case. Yesterday Meg's Latin book was stolen from her locker. Three students may have witnessed the theft. None of them is very reliable. Anne is given to proving theorems in her head, and tends to be oblivious to her surroundings when preoccupied with a tricky proof. To compensate for her habitual distractedness, she draws plausible inferences about mundane events, and often does not notice whether her opinion is due to observation or to such an inference. Ben frequently forgets to wear his glasses. Like Anne, he draws plausible inferences about events around him, and tends not to remember having done so. Chauncy is simply a liar. Presumably he knows when he is speaking sincerely, but given the fluency and frequency of his lies, nothing he says is trustworthy. Not surprisingly, the social circles of the three students do not intersect; none would deign to speak to the others. When questioned about the theft, Anne and Ben report what they think they saw, but confess that they are not sure what they actually witnessed and what they inferred. Chauncy insists that his report is accurate, but in view of his record, his claim is suspect.

Individually, none of the reports would count for much. Had only one of the witnesses been present, the most we could reasonably conclude would be that the thief might fit the description. But all three reports agree, and agree in alleging that the thief had an unusual appearance: He had spiked green hair. This makes a difference. Even though individually each report is dubious, and the probability of a green haired textbook thief is low, the fact that the

three reports provide the same antecedently improbable description inclines us to believe it. Their accord evidently enhances the epistemic standing of the individual reports [Lewis, 1946, p. 346]. We seem to have more reason to believe each of them in light of the others than we have to believe them separately. The question is why? How can multiple statements, none of which is tenable, conjoin to yield a tenable conclusion? How can their relation to other less than tenable claims enhance their tenability?

Given the unreliability of the witnesses, we might expect them to be wrong about the thief. But we would not expect them to all be wrong in the same way. The fact that they agree needs an explanation. If they were in cahoots, the explanation would be straightforward: They conspired to tell the same tale. But not being on speaking terms, they are probably not co-conspirators. If the description they provided fit a relevant stereotype, then a penchant for plausibility could explain their accord. But green spiked hair is far from any stereotype one might harbor for a textbook thief. So despite Anne's and Ben's propensity to draw inferences based on plausibility, their descriptions of the thief do not seem to result from such an inference. Evidently the best explanation of the agreement is that the reports are true.

It is not just our ability to exclude obvious alternatives that leads us to credit the allegation. A variety of collateral considerations support it. Some bear directly on the content of the claim. Dan dimly recalls seeing an odd looking stranger lurking in the hallway. The custodian thinks he saw a container of hair dye in the trash. Although the tentativeness of these reports makes them less than wholly credible, they are suggestive enough to buttress the eyewitness testimony. Other collateral considerations concern the witnesses and their circumstances. Book thefts are observable events, so there is nothing inherently dubious about a claim to have seen someone steal a book. The light and the sight lines were such that the witnesses could have seen what they report. The witnesses are adept at recognizing furtive adolescent behavior. None was subject to psychological experiments with implanted memories. None was on drugs. And so on. Separately, these factors count for little. Either their credibility is low or their bearing is slight. But they weave together to make a solid case. This suggests that

the epistemic tenability of the several reports and the conclusion they sanction derives from their mutual supportiveness.

Although our focus is on the status of the allegation, it is the really account as a whole that is or is not acceptable. Many of the relations of justification are reciprocal. The allegation is acceptable only if (at least most of) the rest of the constellation of supporting considerations is. But since the eyewitnesses are unreliable and the contentions of the collateral witnesses are tenuous, the acceptability of the testimony likewise depends on the acceptability of the allegation. The epistemic status of the allegation is inseparable from the status of the rest of the story. Some of the background information may be separately secured, but to a considerable extent, the various components of the story stand or fall together.

The thesis of the sort of epistemological holism that I want to consider is that epistemic justification is primarily a property of a suitably comprehensive, coherent account, when the best explanation of coherence is that the account is at least roughly true. The epistemic justification of individual claims derives from their membership in a justified account. There is no universally accepted criterion of coherence. But at least this is required: The components of a coherent account must be mutually consistent, cotenable and supportive. That is, the components must be reasonable in light of one another. Since both cotenability and supportiveness are matters of degree, coherence is too. So if it can be shown that epistemic justification is a matter of coherence, there remains the question of how coherent an account must be in order for it to be epistemically justified. Before facing that worry, though, other challenges need to be met. At least two worries immediately arise. The first is that coherence is too demanding an epistemic requirement. The second is that it is not demanding enough.

Even where we take ourselves to be on solid ground, contravening considerations are not uncommon. Mrs. Abercrombie, the aging geometry teacher, says that during the relevant period she saw a young man sporting a green hat. A green hat is not green hair, so her report conflicts with the reports of the other witnesses. Ms. Mintz, the hall monitor, insists no one was in the corridor at the time of the alleged theft. Mr. Miller, the classics teacher, disputes the allegation

on the grounds that students do not want Latin books enough to steal them. These reports are clearly relevant to and at odds with the account I gave. If we incorporate them into my account, we render it incoherent. But we seem to have no legitimate reason to exclude them. The problem is this: The discussion so far suggests that the credibility of the various claims comprising an account depends on how well they hang together. If so, the failure of other, equally relevant information to cohere threatens to discredit the account.

Although true, this is not so daunting as it appears. The immediate threat of incoherence comes from assuming that we must take seemingly contravening considerations at face value and incorporate them into an account as they stand. But we need do no such thing. Rather, we assess contravening considerations just as we do the rest of our evidence. Recall that we did not take the eyewitness reports at face value. We initially deemed them suspect because our background information indicated that the informants are unreliable. The credibility of the reports increased because of their agreement with one another and the support provided by collateral information. That agreement gave us reason to think that the general unreliability of the witnesses did not affect the standing of these particular reports. Contravening considerations are subject to similar assessments. Mrs. Abercrombie, being near-sighted and woefully out of date, cannot even imagine that a green thatch on someone's head might be his hair. That being so, her characterization of the suspect as wearing a green hat seems close enough to count as supporting rather than undermining the original allegation. Although Ms. Mintz flatly disputes what others have said, there are reasons to doubt that her claim is true. Since the three eyewitnesses saw each other in the corridor during the period when Ms. Mintz denies that anyone was there, her contention is dubious on independent grounds. Since she occasionally goes AWOL to smoke a cigarette, there is reason to suspect that she was absent when the theft occurred. Mr. Miller's argument cannot be so easily discredited. But the book is gone. Meg put it in her locker when she arrived at school. It was not there when she returned. Even if Latin books are not attractive targets for teenage thieves, the book's having been stolen may better explain its absence than any available alternative would. Just as other considerations compensate for the improbability of a

green haired thief, other considerations compensate for the improbability of a Latin book thief. In determining the acceptability of a claim, we assess the considerations that afford evidence pertaining to its tenability. This is not always a simple yes/no matter. We may find that although an evidence statement is unacceptable or unsupportive as it stands, with suitable modifications, it would be. And we may find that the modifications themselves are acceptable. Coherence remains crucial. Sometimes it is achieved directly, sometimes by discrediting or disarming threats.

The coherence that affords epistemic justification is not just coherence among object-level deliverances. We have higher-order commitments about what sorts of object-level deliverances are trustworthy, about how much credibility to accord them, about how they ought to mesh, and about what to do when commitments clash. These higher-order commitments supply reasons to revise or reject some deliverances but not others when conflicts occur. The coherence that constitutes epistemic justification is something we achieve, not something that simply falls out of the relations in which our object-level deliverances happen to stand to one another.

The second worry is that coherence can readily be achieved through epistemically illicit means. A good 19th century novel is highly coherent, but not credible on that account. Even though *Middlemarch* is far more coherent than our regrettably fragmentary and disjointed views about the book theft, the best explanation of its coherence lies in the novelist's craft, not in the truth (or approximate truth) of the story. The coherence of the story affords virtually no reason to think it is true. This is surely right. But rather than taking this objection to completely discredit the contention that coherence conduces to epistemic acceptability, I suggest that it indicates something different: Coherence conduces to epistemic acceptability only when the best explanation of the coherence of a constellation of claims is that they are (at least roughly) true.

Although epistemology generally focuses on the beliefs of a single individual, I began with a public case because the otherwise unlikely agreement of independent witnesses clearly shows how the best explanation of the coherence of a given body of claims may be that they are

(at least roughly) true. The case of a single individual can be trickier. Sometimes people confabulate. They compose a coherent narrative by ignoring, bracketing or overlooking factors that detract from the story they seek to construct. The process may be unconscious. Obviously, when a subject is confabulating, the coherence of her beliefs is not explained by their truth. If it is hard to tell whether she is confabulating, it is hard to tell whether coherence confers epistemic standing on her beliefs. But to understand how, why, and when coherence engenders credibility, it is best to put this complication aside. Then we see that the story I have told could be told of a single epistemic agent as well. If the best explanation of the coherence of an agent's system of thought is that it is at least roughly true, and she has no overriding reason to think otherwise, she is justified. Anne is aware of what she thinks she saw, and what she thinks the other witnesses report. She is privy to the relevant background information about apparent sight lines and the like. Since her various relevant cognitive commitments mesh and the best explanation of their meshing is that they are at least roughly true, according to epistemological holists, she is justified in accepting them.

One might argue that even the best 19th century novel does not pose as great a threat as we sometimes suppose. No matter how deeply immersed I am in the story, a single glance up from the page is enough to convince me that I am not in a drawing room in 19th century England. The story, though internally coherent, manifestly fails to mesh with the rest of my experience. This is true, but the question is what to make of it. On the one hand, too restricted a cluster of mutually supportive claims seems inadequate to engender credibility. We can't make the story credible simply by ignoring everything else we believe. On the other, insisting that all our commitments need to cohere seems unduly demanding. If acceptability requires coherence with everything we accept (or with everything we accept for cognitive purposes [Lehrer, 1986]), it is but a short step to skepticism. One wayward belief, however remote from current concerns, could discredit an entire constellation of beliefs. Theories that ground justification in coherence then face a problem of scope.

Worries about scope, however, seem not to do justice to the problem that confronts us

here. Faced with a clash between the deliverances of the novel and those of my glance, it is obvious which I should accept. There is no temptation to resolve the tension by dismissing perceptual deliverances or taking *them* to be the fiction. They seem to possess an epistemic privilege that prevents considerations of coherence from overriding them. The capacity of perceptual deliverances to trump the claims of a tightly knit novel may seem conclusively to demonstrate that epistemological justification cannot consist in coherence.

The matter deserves further consideration though. Until the source of perception's epistemic privilege is clear, it is premature to rule coherence out. A variety of reasons have been offered. Foundationalists argue on a priori grounds that knowledge requires that there be some independently credible beliefs. They hold that perceptual deliverances are among the independently credible beliefs because perceptual deliverances derive at least some of their warrant from the circumstances in which they occur, not their relation to other deliverances. Exactly how credible they are is a matter of dispute [BonJour, 1985, pp. 26-30]. But they must, foundationalists contend, have some measure of credibility that does not derive from their accord with other convictions. Reliabilists argue that a deliverance is epistemically acceptable if produced by a reliable mechanism. Some perceptual mechanisms are reliable, hence some perceptual deliverances are acceptable. Since the reliability of perceptual mechanisms is independent of the relations of their deliverances to other deliverances, perceptual deliverances are independently credible.

There are at least two separate insights here. The reliabilist argument targets the need for a link to the world. The reason for crediting the casual glance while dismissing the deliverances of the novel is that we take it that perception provides the link. The way the world is constrains our perceptual deliverances more immediately and directly than it does our other beliefs. Insofar as the contents of knowledge claims concern the way the world is, it makes sense that the constraints the world supplies should override other considerations. The foundationalist position underscores the idea that some deliverances -- in particular, those of perception -- seem at least *prima facie* credible independently of their connections to other beliefs.

What the objections show is that if perception is to provide the sort of check on theorizing that we think it should, egalitarianism vis ... vis object-level deliverances will not do. An egalitarian theory would hold that each deliverance has an equal claim on our epistemic allegiance. On the principle of one man, one vote, there is no basis for privileging some deliverances over others. If a perceptual deliverance fails to cohere with an otherwise coherent theory, the perceptual deliverance ought to be rejected then, since the claims of the many outweigh the claims of the one. But no matter how comprehensive and integrated an empirical account is, no matter how many other beliefs the account manages to incorporate, observations should have the capacity to discredit it. They have that capacity only if the epistemic claims of perceptual deliverances at least sometimes outweigh those of theory. But it does not follow that perceptual deliverances must be utterly immune to revision or rejection on the basis of considerations of coherence. Nor does it follow that the epistemic privilege granted to perceptual deliverances is independent of coherence considerations.

If we think about our situation when we glance away from the novel, we recognize that we draw on more than the sentences comprising the novel and our current perceptual deliverances. We tacitly rely on a fairly extensive and epistemologically informed understanding of novels and perception. We know enough about underlying mechanisms to have reason to credit some perceptual deliverances. We know enough about literature to realize that novels are typically literally false. That constitutes sufficient reason for even casual perceptual deliverances to override the claims of the novel.

Juxtaposing the novel with perception might seem to make the problem too easy, though. Regardless of what we think about perception, if we recognize that a novel is a work of fiction, we have reason to discount any direct claims it may seem to make on our epistemic allegiance. (I say direct claims because I believe that novels play a significant, albeit indirect role in the advancement of understanding. But how they do so is not germane to this discussion [Elgin, 1996, 183-200].) The serious challenge comes from a coherent factual account that conflicts with perceptual deliverances. If holism holds that such an account always overrides perceptual

deliverances, it seems plainly unacceptable. However tightly woven an empirical account may be, we would be epistemically irresponsible to ignore recalcitrant evidence. Foundationalists take this latter point to be decisive: If observation can show a theory to be unjustified, then coherence cannot be the locus of justification.

This would be so, if observation worked in isolation. For then, owing to its epistemic privilege, one perceptual deliverance would have the capacity to discredit an entire system of thought. But this is a myth. Only observations we have reason to trust have the power to unseat theories. So it is not an observation in isolation, but an observation backed by reasons that actually discredits the theory.

The holist response to the challenge presented by observation is this: A priori, perceptual deliverances have no special weight. They are just deliverances jockeying for inclusion in coherent bodies of thought. But over time, as we attend to the fates of our various deliverances, we learn that the incorporation of some, but not others, yields accounts which are borne out by further experience, hence which retain their coherence over time. This gives us grounds for discrimination. We realize that the deliverances we take to be perceptual are more likely to be confirmed than spontaneous deliverances that just leap to mind. So we assign greater weight to perceptual deliverances than to passing thoughts. Moreover, we learn that not all perceptual deliverances are on a par. Those that are credible tend to come in mutually reinforcing streams, so isolated perceptual deliverances count for little. We begin to draw distinctions among perceptual deliverances. For example, we discover that peripheral vision is less trustworthy than central vision. So we have reason to discount what we see out of the corner of the eye. This is not to say that we dismiss the deliverances of peripheral vision out of hand, but that we demand more in the way of corroboration. Some of us discover that we are color blind or tone deaf or myopic. That is, we learn that our perceptions of colors, tones or the dimensions of distant objects are not to be trusted. And so on. We come to assign different weights to perceptual deliverances depending on how well they accord with other things we take ourselves to have reason to credit -- other appearances of the same object, the reports of other observers, the

implications of our best theories about the visible properties of items of the kind in question, and so forth.

The issue is not simply how well a given content meshes with other things we believe, but how well a given content from a given source in given circumstances does. The weight we attach to perceptual deliverances derives from our understanding of the world and our access to it. Initially, perhaps, this is just a matter of track records. Some perceptual deliverances seem to integrate better into acceptable systems of thought than spontaneous thoughts that just leap to mind. Later, as we develop physiological and psychological accounts of ourselves, which explain our perceptual mechanisms, we gain additional reasons to take some perceptual deliverances to be credible. The epistemic privilege that some perceptual deliverances enjoy then derives from an understanding of ourselves as perceiving organisms. That is, the reason for assigning those deliverances significant epistemic weight derives from the coherent account of perception that backs the assignment. Contrary to what foundationalists contend, the justification for privileging perception derives from the relation of perceptual judgments to the rest of our theory of ourselves as cognitive agents interacting with a mind-independent world.

The reliabilist account seems to fare slightly better. What justifies assigning my visual inputs significant epistemic weight seems to be that vision is a reliable perceptual mechanism. What justifies dismissing my forebodings is that premonition is not. This is not quite right though. It is not the brute reliability or unreliability of a source that supplies the justification, but an understanding of that reliability or unreliability. Even if my forebodings are accurate, so long as we have no reason to trust them, they bear little weight.

This argument explains both why some perceptual deliverances have the capacity to unsettle theory, and why those deliverances are not intrinsically privileged. They owe their epistemic status to their place in our evolving understanding of the world and our modes of access to it. This has two welcome consequences. The first is that the privilege they enjoy is revocable. When I learn that I am color blind, I need to revise my views about which of my visual deliverances are acceptable. The second is that non-perceptual deliverances can in

principle be equally weighty. This is an advantage in accounting for the epistemic status of scientific evidence and of testimony.

A look at modern science shows that it is not just (or perhaps even mainly) bare perceptual deliverances that have the capacity to discredit theory. The outputs of measuring devices do too. In an effort to retain a tie to classical empiricism, some philosophers of science argue that measuring devices are simply extensions of our senses. Just as eyeglasses enable nearsighted people to see what otherwise they could not, telescopes and microscopes enable everyone to see what otherwise we could not. So if seeing something in suitable circumstances has sufficient weight to undermine a coherent cluster of claims, seeing something through a telescope or microscope should be able to do so too. This idea is not unreasonable so long as we restrict ourselves to devices like optical telescopes and microscopes. But it stretches the bounds of plausibility to contend that radio telescopes, electron microscopes, MRIs and the like are also mere extenders of the sense of sight. It seems better to forego the strained analogy and simply characterize such devices as detectors. Then an understanding what they detect, how they detect, and why they should be trusted supplies reason to accord their outputs considerable weight. Even without the strained analogy, the argument for crediting the outputs of scientific instruments thus parallels the argument for crediting perceptual deliverances. For although they are not perceptual mechanisms, the devices are among our modes of access to the world.

Testimony poses a similar problem. We acquire many of our beliefs from the testimony of others, and consider those beliefs justified. Some philosophers say that the justification for accepting testimony is a priori. *Ceteris paribus*, we are justified in accepting what people tell us. Others say it is inductive. We should believe only those who have shown themselves to be relevantly reliable in the past. The former seem to endorse gullibility, the latter to unduly limit acceptability. Something more sensitive is wanted. Evidently the question is not whether testimony *per se* is or is not *prima facie* acceptable. Some testimony is frankly incredible; some requires a good deal of corroboration; some is straightforwardly acceptable. The acceptability of a bit of testimony depends on how well its content coheres with other relevant deliverances, how

well the belief that the testifier is competent with respect to her allegation coheres, and how well the belief that she is sincere coheres. Because of its mesh with our background beliefs, straightforwardly acceptable testimony scores high on all of these measures. Just as different perceptual deliverances are accorded different weights, so are different testimonial deliverances. Testimony with sufficiently strong backing can discredit a hitherto coherent cluster of beliefs.

Even though the deliverances of perception, testimony, and instrumental readings have no special standing a priori, in light of our developing theories of the world and our modes of access to it, some of them turn out to have considerable epistemic weight. This satisfies the demand that acceptable beliefs be appropriately constrained by the way the world is. It also reveals that holism has the resources to recognize that deliverances can differ in weight, some being more credible than others. The claims of the few can in suitable circumstances outweigh the claims of the many.

Achieving coherence is not just a matter of excluding untoward deliverances though. In the interests of systematicity, we may incorporate considerations we have no antecedent reason to believe. For example, although there is no direct evidence of positrons, symmetry considerations show that a physical theory that eschewed them would be significantly less coherent than one that acknowledged them. So physics's commitment to positrons is epistemically appropriate. Considerations we have no independent reason to believe can acquire tenability then because they strengthen the coherence of the systems they belong to.

The issue of scope remains. The totality of a person's beliefs and/or deliverances is not particularly coherent. Not only are there outliers and inconsistencies among beliefs, there are also clusters of beliefs that are relatively isolated from one another. Meg's cluster of beliefs about the pituitary gland, the evidence that bears on the acceptability of these beliefs, the trustworthiness of bits of testimony on the subject, and the proper methods for assessing such things has few and loose connections to her cluster of views about parliamentary procedure, the evidence that bears on these views, the trustworthiness of testimony about the subject, and the proper methods for assessing them. It seems that she could easily be badly wrong about the

former without her error having any significant effect on the tenability of her views about the latter. Outliers and inconsistencies among beliefs are in principle relatively unproblematic. According to a holism, outliers lack justification. Because they lack suitable connections to other things we believe, we have no reason to credit them. Inconsistencies among beliefs conclusively demonstrate that some of the beliefs are false. But it is not obvious that mutual indifference of belief clusters is objectionable. It is not clear that we should consider Meg epistemically defective because of the lack of close ties between the two clusters. On the other hand, if the clusters of beliefs are too small and too numerous, complacency over their mutual indifference seems problematic. We do not want to license ignoring inconvenient tensions among beliefs by consigning them to mutually irrelevant clusters.

The problem neither has nor needs an a priori resolution. Our evolving theories of the world and our access to it provide us with an appreciation of the relations in which our various clusters of beliefs should stand to one another and the requirements they should satisfy. Such a laissez-faire attitude might seem to allow for the acceptability of crazy constellations of views. If we leave it to our evolving theories to decide what range of considerations acceptable accounts must answer to, we may be forced to endorse isolated islands of claptrap. The worry is more apparent than real. We have theories about theories, which enable us to assess the reasons, methods, standards and evidence that our various object-level theories appeal to. Some requirements, such as logical consistency, apply globally. Regardless of how far apart Meg's views about politics and endocrinology are, unless they can be conjoined without contradiction they are not all acceptable. Other requirements, like the need to respect judicial precedents or to accord with biochemical findings, are more limited in range. But even these do not enable us to isolate belief clusters entirely. Even if Meg's views about endocrinology and politics have few points of contact, her views about endocrinology and hematology have many.

Consistency requirements do more than rule out express contradictions. The requirement that like cases be treated alike demands that if a consideration has weight in one area but not in another, there be an acceptable reason for the difference. In order to be tenable, a system of

mutually reinforcing claims must either answer to the logical and evidential standards to which other theories are subject or be backed by a tenable account of why those standards do not apply. Some theories have such backing. There are, for example, cogent reasons why mathematics is not subject to empirical testing. So infinitary mathematics is not threatened the absence of empirical evidence for its findings. In epistemically objectionable cases, no such reasons are available. The claims of astrology, although mutually reinforcing, are epistemically unacceptable because they yield predictions that are either too vague to be tested or are not borne out when tested. Since astrology makes empirical claims, there are considerations to which it ought to be responsive which it fails to accommodate. To say that something cannot be ruled out a priori is not to say that it cannot be definitely and decisively ruled out.

Epistemological positions that construe knowledge as justified true belief generally treat being justified, being true, and being believed as three separate features of a propositional content. The standard objection to coherentism is that coherence among propositional contents is so easily achieved that it affords no reason to believe that the contents are true, hence no justification for them. This overlooks the fact that the contents in question are not just any propositional contents, they are belief contents or deliverance contents. That is, they are contents that present themselves as true. This makes a difference. For the fact that they present themselves as true gives us some slight reason to think that they are true. The word 'slight' is crucial. I do not contend that we have sufficient reason to credit such contents. But at least two considerations speak in favor of granting them a slight measure of credibility. Beliefs form the basis for action, so the success of our actions affords evidence of the truth of the corresponding beliefs. Moreover, we learn from experience. Once we come to recognize that premonitions tend not to be borne out, we cease to credit them. We may continue to experience feelings of foreboding, for example, but they cease to qualify as deliverances.

Manifestly these considerations are far too weak to demonstrate that beliefs or deliverances are epistemically justified. They do, however, give us reason to think that beliefs and deliverances have some claim on our epistemic allegiance. They have an epistemic edge.

We have better reason to incorporate them into our systems of thought than to incorporate contents we are neutral about. Beliefs and deliverances are, I suggest, initially tenable. But initial tenability is a weak and precarious epistemic status. Considerations of overall coherence often require revision or rejection of initially tenable commitments. Initially tenable commitments can conflict. They may be mutually incompatible or non-cotenable. Or they may be sufficiently isolated that they are incapable of giving support to or gaining support from other things we believe. Then they cannot be incorporated into an epistemically acceptable system.

Epistemically acceptability, I contend, requires reflective equilibrium [Elgin, 1996; Rawls, 1971]. A system of thought is in equilibrium if its elements are reasonable in light of one another. This is a matter of coherence. An equilibrium is reflective if the system is as reasonable as any available alternative in light of our initially tenable commitments. Such a system is not required to incorporate as many initially tenable commitments as possible. As we have seen, there are weighting factors that favor some incorporations over others. Moreover, rather than incorporating commitments, a system may show why we were misled into accepting them, or may include modifications of them.

The standards of reasonableness are second-order commitments, and are subject to the same sorts of considerations as our first-order deliverances. The fact that we accept them indicates that they are prima facie acceptable. But they can conflict, or fail to yield verdicts in cases where they should, or yield verdicts that we find unacceptable. Then they too are subject to revision or rejection in order to yield a comprehensive system of first- and second-order commitments that is on reflection something we can endorse.

Whether the sort of holism that results is a coherence theory is not clear. Using Bonjour's [1985] categories, it might be classified as a very weak foundationalism or as a coherence theory. Deliverances derive their initial tenability from their status as deliverances. That suggests that something other than coherence is involved. But initially tenable commitments display at least two features that are not characteristic of standard foundational beliefs. First, there are no intrinsically privileged kinds of deliverances. The account does not insist that there is something

epistemically special about perception or introspection or analyticity. It simply says that the fact that a consideration presents itself as true gives it a modest measure of tenability. Second, even that small measure of tenability is easily lost. Tenable theories are justified in part by reference to initially tenable deliverances, but they need not incorporate the deliverances by reference to which they are justified.

Whether we call such an epistemology a coherence theory does not in the end matter. The virtues of the theory are these: (1) It does not privilege any sorts of beliefs or representations a priori. What beliefs and representations are worthy of acceptance is something we learn by developing increasingly comprehensive, coherent accounts of the world and our access to it. (2) It enables us to start from whatever deliverances we happen to have. But because it insists that we subject those deliverances to rigorous assessment, such a starting point is not question begging. (3) The standards of assessment are themselves the fruits of epistemic activity, and can change in response to feedback [Goodman, 1984, p. 69]. (4) Hence, everything is subject to revision. A system of thought that we can on reflection accept today may be one that we cannot on reflection accept tomorrow. But so long as a system is in reflective equilibrium and the best of explanation of its being so is that it is at least roughly true, it and its components are justified. What results is neither certainty nor skepticism but a fallible, provisional, but reasonable epistemological stance.

CATHERINE Z. ELGIN

REFERENCES

- Blanshard, Brand (1939). The Nature of thought, London: George Allen & Unwin.
- BonJour, Laurence (1985). The Structure of empirical knowledge, Cambridge, MA: Harvard.
- Elgin, Catherine (1996). Considered judgment, Princeton, NJ: Princeton.
- Goodman, Nelson (1984). Fact, fiction, and forecast. Cambridge, MA: Harvard.
- Lehrer, Keith (1986). 'The Coherence Theory of Knowledge,' Philosophical topics, 14, 5-25.
- Lewis, C. I. (1946). An Analysis of knowledge and valuation, La Salle, IL: Open Court.
- Rawls, John (1971). A Theory of justice, Cambridge, MA: Harvard.
- Rescher, Nicholas (1973). The Coherence theory of truth, Oxford: Clarendon.

FURTHER READING

- Adler, Jonathan. (1986). 'Knowing, betting, and cohering.' Philosophical Topics, 14, 243-257.
- Bender, John (ed.). (1989). The Current state of the coherence theory. Dordrecht: Kluwer.
- Goodman, Nelson and Elgin, Catherine. (1988). Reconceptions. Indianapolis, IN: Hackett.
- Harman, Gilbert. (1973). Thought. Princeton, NJ: Princeton.
- Lehrer, Keith. (1974). Knowledge. Oxford: Clarendon.
- Sellars, Wilfrid. (1968). Science and metaphysics. London: Routledge & Kegan Paul
- Sellars, Wilfrid. (1963). Science, perception and reality. London: Routledge & Kegan Paul.
- Sosa, Ernest. (1985). 'The Coherence of virtue and the virtue of coherence.' Synthese, 64, 3-28.
- Sosa, Ernest. (1980). 'The Raft and the Pyramid.' Midwest studies in philosophy, 5, 3-26.
- Williams, Michael. (1980) 'Coherence, Justification, and Truth.' Review of Metaphysics, 37, 243-272.

Biographical Note

Catherine Z. Elgin is the author of *Considered Judgment* (Princeton, 1996), *Between the Absolute and the Arbitrary* (Cornell, 1997), *With Reference to Reference* (Hackett, 1983), and co-author, with Nelson Goodman of *Reconceptions in Philosophy and Other Arts and Sciences* (Hackett, 1988). She is professor of the philosophy of education at Harvard Graduate School of Education.