Chapter Ten

Can Belief Be Justified Through Coherence Alone?

Catherine Elgin and James Van Cleve both answer this question negatively. But whereas Van Cleve advocates a moderate version of foundationalism, Elgin defends a broadly coherentist view. According to her, justification is primarily a matter of explanatory coherence. The justification an individual belief enjoys is derived from the coherence of the overall system. In his essay, Van Cleve argues that, although coherence is indeed a source of justification, it cannot by itself render a belief completely justified. According to Van Cleve, no belief could be justified unless it were possible for some beliefs to acquire complete justification without receiving support from any other beliefs. In their respective responses, Elgin and Van Cleve continue the dispute, focusing on issues such as conjunction closure, corroboration by independent witnesses, empirical generalizations, revisability, and the skeptical threat of being deluded.

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

Contemporary Debates in Epistemology, Second Edition. Edited by Matthias Steup, John Turri, and Ernest Sosa.

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 with each other 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 focusing 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 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 nineteenth-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 nineteenth-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 nineteenth-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 hand, 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-à-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, pp. 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 unreliabilty 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 near-sighted 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 forgo the strained analogy and simply characterize such devices as detectors. Then an understanding of 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 meshing 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' 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 by 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.

Epistemic 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 as follows. (i) 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. (ii) 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. (iii) The standards of assessment are themselves the fruits of epistemic activity, and can change in response to feedback (Goodman, 1984, p. 69). (iv) 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.

References

Blanshard, B. (1939) *The Nature of Thought*. London: George Allen & Unwin.
BonJour, L. (1985) *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.
Elgin, C. (1996) *Considered Judgment*. Princeton, NJ: Princeton University Press.
Goodman, N. (1984) *Fact, Fiction, and Forecast*. Cambridge, MA: Harvard University Press.
Lehrer, K. (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, J. (1971) *A Theory of Justice*. Cambridge, MA: Harvard University Press.
Rescher, N. (1973) *The Coherence Theory of Truth*. Oxford: Clarendon Press.

Further Reading

Adler, J. (1986) Knowing, betting, and cohering. *Philosophical Topics* 14: 243–257.
Bender, J., ed. (1989) *The Current State of the Coherence Theory*. Dordrecht: Kluwer.
Goodman, N. and Elgin, C. (1988) *Reconceptions*. Indianapolis: Hackett.
Harman, G. (1973) *Thought*. Princeton, NJ: Princeton University Press.
Lehrer, K. (1974) *Knowledge*. Oxford: Clarendon Press.
Sellars, W. (1963) *Science, Perception and Reality*. London: Routledge & Kegan Paul.
Sellars, W. (1968) *Science and Metaphysics*. London: Routledge & Kegan Paul.
Sosa, E. (1980) The raft and the pyramid. *Midwest Studies in Philosophy* 5: 3–26.
Sosa, E. (1985) The coherence of virtue and the virtue of coherence. *Synthese* 64: 3–28.
Williams, M. (1980) Coherence, justification, and truth. *Review of Metaphysics* 37: 243–272.

Why Coherence Is Not Enough: A Defense of Moderate Foundationalism

James Van Cleve

I

Foundationalism has been characterized as the view that "the knowledge which a person has at any time is a structure or edifice, many parts and stages of which help to support each other, but which as a whole is supported by its own foundation" (Chisholm, 1964). To unpack the metaphor, we may say that the foundation consists of

basic beliefs – beliefs that a subject is justified in holding even in the absence of any justifying reason for them – and all other justified beliefs derive their justification at least in part from such basic beliefs.

The classical argument for foundationalism is an infinite regress argument going back to Aristotle: unless there were basic beliefs, every justified belief would rest on further justified beliefs, and so on without end. An infinite regress is not in fact the only alternative to basic beliefs, but the other alternatives are equally problematic, or so foundationalists maintain. To canvass all the options, let us set down four propositions that jointly imply the existence of an infinite regress of justified beliefs:

1 Some beliefs are justified.
2 No belief is justified unless some other belief serves as a reason for it.
3 One belief cannot serve as a reason justifying another unless the first is itself justified.
4 If A serves as a reason justifying B, then B cannot serve (directly or indirectly) as a reason justifying A.

These four propositions jointly entail the existence of an infinite regress of justified propositions. We are therefore faced with five alternatives: accept the infinite regress or reject one of the four assumptions.

Skeptics (of the universal ilk) deny 1, maintaining that no beliefs whatever are justified. *Foundationalists* deny 2, maintaining that some beliefs are justified in the absence of reasons. *Positists* (not to be confused with positivists) deny 3, maintaining that chains of justifying reasons can terminate in reasons that are not justified themselves, but are simply individual or societal posits.[1] *Coherentists* deny 4, maintaining that beliefs can be justified in virtue of relations of mutual support. *Infinitists* accept all four assumptions and the resulting infinite regress.

All five options have their takers. Skepticism and infinitism are both defended elsewhere in this volume (see the essays by Fumerton and Klein). Positism finds advocacy in Wittgenstein's *On Certainty* and in assorted postmodern thinkers (Wittgenstein, 1969). But I think it is fair to say that the leading contenders among the five options are foundationalism and coherentism.

Coherentism is sometimes characterized as a view that sanctions circular reasoning, but that is an oversimplified construal of it. Coherentists do not typically endorse simple loops in which A justifies B, B justifies C, and C justifies A; rather, they envision vast webs of belief in which everything is supported by some significant portion of the remaining beliefs: A by B and C, B by A, D, J, and K, and so on.

For their part, foundationalists do not typically deny the power of coherence to contribute to the overall epistemic status of a body of belief. They simply insist that coherence cannot do all the work on its own – there must be at least a modicum of intrinsic credibility or non-inferential warrant possessed by basic beliefs before coherence can have its amplifying effect.

Laurence BonJour distinguishes three grades of foundationalism (BonJour, 1985, pp. 26–30). According to *strong* foundationalism, basic beliefs are "not just adequately justified, but also *infallible*, *certain*, *indubitable*, or *incorrigible*" (BonJour, 1985, pp. 26–27). According to *moderate* foundationalism, the non-inferential warrant possessed by basic beliefs need not amount to absolute certainty or any of the other

privileged statuses just mentioned, but it must be "sufficient by itself to satisfy the adequate-justification condition for knowledge" (BonJour, 1985, p. 26). Finally, according to *weak* foundationalism,

> basic beliefs possess only a very low degree of epistemic justification on their own, a degree of justification insufficient by itself either to satisfy the adequate-justification condition for knowledge or to qualify them as acceptable justifying premises for further beliefs. Such beliefs are only "initially credible," rather than fully justified. (BonJour, 1985, p. 28)

We must rely on coherence among such initially credible beliefs to amplify their level of warrant up to the point where it is adequate for knowledge.

As BonJour notes, weak foundationalism could be regarded as a hybrid view, mixing together foundational and coherentist elements. In fact, Susan Haack prefers to call it "foundherentism," which she illustrates with the example of a crossword puzzle. Experience corresponds to the clues, which give an initial presumption in favor of certain beliefs or entries in the puzzle; the initial beliefs are then confirmed by the way in which they interlock with other entries (or sometimes discarded because they do not fit in). Thus does coherence amplify (or its absence erode) the initial warrant possessed by basic beliefs (Haack, 1993).[2]

But if coherence can elevate the epistemic status of a set of beliefs in this way, what prevents it from generating warrant entirely on its own, without any need for basic beliefs? This is a question that has been asked by several authors, including BonJour:

> The basic idea is that an initially low degree of justification can somehow be magnified or amplified by coherence, to a degree adequate for knowledge. But how is this magnification or amplification supposed to work? How can coherence, not itself an independent source of justification on a foundationalist view, justify the rejection of some initially credible beliefs and enhance the justification of others? (BonJour, 1985, p. 29)

The implied suggestion is that if coherence can do what the weak foundationalist allows, it can also do what the thoroughgoing coherentist says it can do.

In the next section, I address this challenge. I argue that there is indeed a good rationale for the weak foundationalist's insistence that before coherence can do its work, there must be initially credible inputs for it to work upon. But I also argue that the level of initial credibility cannot be as low as that envisioned by weak foundationalists, who should therefore upgrade their position to moderate foundationalism. Finally, in section III, I offer some critical reflections on the coherentist alternative defended by Catherine Elgin.

II

An excellent example of a weak foundationalist theory is provided by C.I. Lewis's theory of memory knowledge. It has two elements:

> First; whatever is remembered, whether as explicit recollection or merely in the form of our sense of the past, is *prima facie* credible because so remembered. And second; when the whole range of empirical beliefs is taken into account, all of them more or less dependent

> upon memorial knowledge, we find that those which are most credible can be assured by their mutual support, or as we shall put it, by their *congruence*. (Lewis, 1946, p. 334)[3]

Lewis defines a congruent set as one in which any member is more probable given the rest than it is on its own:

> A set of statements, or a set of supposed facts asserted, will be said to be congruent if and only if they are so related that the antecedent probability of any one of them will be increased if the remainder of the set can be assumed as premises. (AKV, p. 338)

A point on which Lewis repeatedly insists is that congruence alone cannot generate probability or warrant. Rather, some of the statements must have initial credibility, which congruence can then amplify:

> The feature of such corroboration through congruence that should impress us, is the requirement that the items exhibiting these congruent relationships must – some of them at least – be *independently given facts* or have a probability which is antecedent. (AKV, p. 352)

How *much* probability must the congruent items have? Only a "slight" amount, Lewis tells us, illustrating his point with the example of individually unreliable witnesses who tell the same story:

> Our previous example [AKV, p. 239] of the relatively unreliable witnesses who independently tell the same circumstantial story, is another illustration of the logic of congruence; and one which is more closely typical of the importance of relations of congruence for determination of empirical truth in general. For any of these reports, taken singly, the extent to which it confirms what is reported may be slight. And antecedently, the probability of what is reported may also be small. But congruence of the reports establishes a high probability of what they agree upon, by principles of probability determination which are familiar: on any other hypothesis than that of truth telling, this agreement is highly unlikely. (AKV, p. 346)

But how much probability is a "slight" amount? In one place Lewis says, "Anything sensed as past is just a little more probable than that which is incompatible with what is remembered and that with respect to which memory is blank" (AKV, p. 358). He thereby implies that a remembered piece of information has a probability greater than 0.5, given that it is remembered. (This is because it is an axiom of the probability calculus that $P(\sim h,e) = 1 - P(h,e)$; thus a proposition has greater probability than its negation if and only if it has probability greater than 0.5.) Thus according to Lewis's version of weak foundationalism, the congruence of a set of remembered items can raise their level of justification arbitrarily high, but only if the items have initial credibility amounting to a probability (given that we seem to remember them) greater than 0.5.

To this latter aspect of Lewis's theory BonJour has raised an objection. There is no need, he says, for Lewis's requirement that memory reports or other cognitive deliverances have initial credibility:

> What Lewis does not see, however, is that his own example shows quite convincingly that no antecedent degree of warrant or credibility is required. For as long as we are confident

> that the reports of the various witnesses are genuinely independent of each other, a high enough degree of coherence among them will eventually dictate the hypothesis of truth telling as the only available explanation of their agreement – even, indeed, if those individually reports initially have a high degree of *negative* credibility, that is, are much more likely to be false than true (for example, in the case where all of the witnesses are known to be habitual liars). (BonJour, 1985, pp. 147–148)

We are now presented with a clear-cut issue to investigate: in order for the congruence of a set of items to raise their credibility to near 1, what level of antecedent credibility is required? Must it be greater than that of their negations and thus greater than 0.5, as Lewis maintains? Or may it be less than 0.5, as BonJour implies when he says that the reports of the witnesses (or of our memory) may be more likely false than true?

For light on this question, we may look to discussion of a traditional topic in probability theory: how to assess the probability that independent witnesses who agree in their testimony are telling the truth. This is a problem to which a number of classical authors have proposed answers. One of the standard answers is due to George Boole (1952, p. 364):[4]

> Let p be the general probability that A speaks the truth, q the general probability that B speaks the truth; it is required to find the probability that, if they agree in a statement, they both speak the truth. Now agreement in the same statement implies that they either both speak truth, the probability of which beforehand is pq, or that they both speak falsehood, the probability of which beforehand is $(1 - p)(1 - q)$. Hence the probability beforehand that they will agree is $pq\ (1 - p)(1 - q)$ and the probability that if they agree, they will agree in speaking the truth is accordingly expressed by the formula

$$w\left[= P\left(\text{A and B speak truly, they agree}\right)\right] = \frac{pq}{pq + (1-p)(1-q)}$$

For an explanation of the rationale behind Boole's formula, I must refer the reader to what I have said elsewhere.[5] Here there is space only to note the bearing of his formula on the issue separating Lewis and BonJour. Suppose that A and B each tell the truth 60 percent of the time; that is, suppose that p and q are each equal to 0.6. The reader may verify that in this case, w = 0.69. That is, setting A's credibility and B's each equal to 0.6, the probability that X is true given that they each testify to it is 0.69. More generally, if we plug in any numbers greater than 0.5 as p and q, w will be greater than the mean of p and q. So far we have an illustration of the point, common ground for Lewis and BonJour, that congruence can boost credibility.

But what if we plug in values of p and q equal to or less than 0.5? If 0.5 goes in, 0.5 comes out; and if p and q are each less than 0.5, the output value will be less than their mean. For example, if p and q are each 0.1, w is approximately 0.01. In other words, if the witnesses have what BonJour calls "negative credibility" (credibility less than 0.5), the probability of a statement given that they both testify to it is not enhanced but diminished! So if Boole's formula is correct, Lewis is vindicated and BonJour refuted: with initial credibilities less than 0.5, coherence makes things worse rather than better.

There is reason to be suspicious of Boole's formula, however. In deriving it, he tacitly assumes that there are only two possible answers to the questions put to the witnesses – true or false. That is why he can equate agreement with "both speak truly

or both speak falsely." In a more realistic scenario, the witnesses would be asked multiple-choice questions, and if they agreed in giving the same answer out of ten possible choices (let us say), their agreement would be much more impressive.

A formula allowing for an arbitrary number of possible answers to questions has been devised by Michael Huemer (1997).[6] For simplicity's sake, Huemer assumes that the witnesses have the same level of credibility, so p = q. If n is the number of possible answers and X is the answer on which independent witnesses A and B agree, then Huemer's formula may be written as follows (omitting details of the derivation):

$$P\big(X, A\,\text{says}\,X\ \&\ B\,\text{says}\,X\big)[= W] = \frac{np^2 - p^2}{np^2 - 2p + 1}$$

This formula agrees with Boole's in the special case where n = 2, but gives dramatically different results in the cases not covered by Boole's. Specifically, it enables coherence to have its amplifying affect even when the credibility level of the witnesses is below 0.5, just so long as it is greater than the chance or random guessing level of 1/n. For example, if p is only 0.3, but there are ten possible answers (say, ten possible last digits in a glimpsed license plate number), then w = 0.62. If p = 0.3 and n = 100 , then w = 0.86. For any value of p, just so long as it is greater than zero, we can bring the final probability of X (i.e., its probability given that the witnesses agree on it) as close as we like to certainty by making n high enough. With a more complex version of the formula, we can also make w higher by increasing the number of witnesses.

If Huemer's formula is correct, then, BonJour is vindicated and Lewis refuted. Initial credibilities need not be greater than 0.5; they need only be greater than the chance level of 1/n (where n = the number of possible answers). If enough witnesses agree without collusion in giving the same answer from among a large enough number of choices, it becomes overwhelmingly likely that they are correct, even if their initial level of credibility was scarcely above zero.

Is BonJour right, then? Can coherence alone be a source of warrant without need of inputs with initial credibility? An answer of yes would be too hasty, for there is another requirement of initial credibility we have yet to consider. The requirement at which we have so far demurred is the requirement that what is reported must be more probable than not (and thus have probability greater than 0.5) given that a witness (or an ostensible memory) attests to it. If several individually unreliable reporters agree without collusion, then the fact to which they bear common witness may have high probability in the end. But in attaching a high final probability to the fact attested, we are of course taking for granted that the various witnesses *do* testify to it. If we had reason to think that the courtroom and all its proceedings were happening only in a dream or a novel, the fact that the ostensible reports hang together would count for little. And so it is with the reports of memory, the senses, and cognitive systems more generally: coherence among them lends high final credibility only on the assumption that the reports genuinely occur.

How, then, do we know *these* things: that witness A does say X, that I do ostensibly remember Y, that I do seem to see Z? Many foundationalists would say that these are the grounds on which the rest of our knowledge rests, and that they must themselves be matters of basic knowledge.[7] Lewis himself famously maintained that nothing can be probable unless something is certain, and among the certainties he

placed the facts that I do have this or that presentation of sense or memory. His insistence on certainty is controversial,[8] but it seems to me that a good case can be made that there must at least be high intrinsic credibility – perhaps high enough to constitute knowledge – attaching to the facts that such-and-such cognitive states (be they experiences, ostensible memories, or beliefs at large) are actually taking place. If this is right, we must not only abjure pure coherentism: we must also adopt a moderate rather than a weak foundationalism.

I see only one plausible alternative to an assumption of high initial credibility or knowledge-sufficient warrant at the foundational level, and that is the view that the promptings of sense or memory function as *external* conditions of knowledge. An external condition of knowledge is a condition that makes knowledge possible regardless of whether it is itself known to obtain. For example, in Goldman's reliability theory, if a subject comes to believe p as the result of a reliable process, his belief is knowledge regardless of whether the subject knows anything about the reliability of the process (Goldman, 1979). Perhaps the facts that I have such-and-such ostensible perceptions or memories could function in this external way, contributing to my knowledge even if not themselves known. The idea would be that my ostensible perceivings and rememberings are not pieces of evidence on which I conditionalize when their epistemic status is high enough;[9] instead, they are facts whose mere obtaining confers credibility on their contents.

I turn now to another point at which I think a theory of knowledge that invokes coherence must make a concession to either foundationalism or externalism.

The example of the witnesses who agree is in one respect a drastically oversimplified case of coherence. The agreement of the witnesses is literal identity, or at least logical equivalence, of content: witness 1 says X and so does witness 2. But the coherence that figures in epistemology is typically a much looser sort of hanging together. The coherence of ostensible memories is not their all being memories that p, for the same p or something logically equivalent. Nor is the coherence of beliefs or cognitions generally like that. Rather, it is a type of coherence that is exemplified by the following items:

I seem to remember seeing a skunk last night;
I seem to remember smelling a skunk last night;
I seem to remember that the lid was on the garbage can when I went to bed;
I now see that the can has been knocked over and trash strewn about;
There was a skunk here last night;

and so on. In other words, it is not identity or even equivalence of content, but rather something like the relation Lewis calls congruence: a matter of each item being more probable given the rest than it is on its own.

What are these coherence-constituting relations of probability founded upon, and how do we know that they obtain?

One answer has been given by Russell: "It is only by assuming laws that one fact can make another probable or improbable" (Russell, 1948, p. 188). Perhaps Russell goes too far in requiring strict laws in order for one fact to make another probable, but it is plausible that we at least require rough empirical generalizations. Where do these generalizations come from? Presumably, they are inferred inductively from particular facts gathered by memory. And now the following difficulty emerges: ostensible

memories give rise to knowledge only with the help of coherence; coherence depends on laws or empirical generalizations; and such generalizations can be known only with the help of memory. In short, we cannot get coherence without the help of laws, and if memory does not suffice on its own to give knowledge of particular facts from which the laws are inferred, we cannot get laws without the help of coherence. It appears to follow that we cannot have any knowledge from memory unless the occurrence of ostensible memories is prima facie sufficient for knowledge. Such was Russell's own conclusion: 'memory is a premise of knowledge. ... When I say memory is a premise, I mean that among the facts upon which scientific laws are based, some are admitted solely because they are remembered' (Russell, 1948, pp. 188–189). Note the word "solely." Russell is saying that individual memories must be capable of giving rise to knowledge on their own, without benefit of coherence. This is compatible, of course, with allowing that the warrant provided by memory is defeasible, as Russell did allow. But the resulting view is nonetheless a foundationalism of memory knowledge stronger than that of Lewis, who required only an initial "slight" presumption in favor of the truth of any ostensible memory. Russell's view accords to memory greater epistemic powers than that: ostensibly remembering that p is a source of prima facie warrant that, if undefeated (and if p is true) is strong enough for knowing that p. In BonJour's terms, we have again advanced from weak to moderate foundationalism, this time as regards the contents of ostensible memories rather than the occurrences of them as mental events.

Russell's argument assumes that coherence and the laws that underlie it contribute to our knowledge only if they are themselves known. As in the case of the occurrences of our ostensible memories, one could challenge this assumption by going external, holding that coherence does its work regardless of whether the subject knows it obtains. This is the second point at which I believe coherentism can avoid a concession to foundationalism only by making a concession to externalism. The fact that p, q, and r do cohere with one another, as well as the facts that they are deliverances of our cognitive systems to begin with, are facts that must either function externally or be known foundationally.

III

As the debate between foundationalists and coherentists has progressed, each side has moved in the direction of the other. Contemporary foundationalists are seldom foundationalists of the strong Cartesian variety: they do not insist that basic beliefs be absolutely certain. They also typically allow that the elements in a system of belief can acquire enhanced justification through their coherence. On the other side, many coherentists admit that coherence alone is not the sole source of justification – there must be some initially credible inputs before coherence can work its wonders. Is there anything more to disagree about, or do foundationalists and coherentists now meet in the middle?

There are indeed still points of difference. To highlight several of them, I shall discuss the broadly coherentist views of Catherine Elgin as developed in her book *Considered Judgment* (1996; cited as CJ hereafter) and in her contribution to this volume.

Elgin characterizes herself as a proponent of reflective equilibrium.[10] As she conceives of it, reflective equilibrium has two chief requirements: "The components of a system in reflective equilibrium must be reasonable in light of one another, and the system as a whole reasonable in light of our initially tenable commitments." (CJ, p. 107; see also pp. ix, 13, and 127–128). It is by the second requirement that Elgin distinguishes her view from a pure coherentism: the components of a system in equilibrium must be answerable not just to one another, but also to our initially tenable commitments.[11]

The second requirement puts a "tether" on permissible systems (CJ, pp. 10, 107, 128), thereby enabling Elgin to avoid some of the objections to pure coherentism. For example, one of the standard objections to coherentism is that the contents of a consistent fairy tale would be a body of warranted propositions (see Schlick, 1973, p. 419). Not so for Elgin, since the propositions in the story may be reasonable in light of each other without being reasonable in light of our initially tenable commitments.

Of the various things we believe, which have initial tenability? In Elgin's view, they *all* do, if we genuinely believe them: anything actually held has some initial tenability or presumption in its favor (CJ, pp. 101–102). The presumption may only be slight and it may be lost in the end, but it is there in the beginning.

What about the various principles of logic, evidence, and method whereby some things are reasonable in light of others? In Elgin's view, these have the same status as everything else: they are initially tenable if held, and they may gain or lose in tenability depending on how they fit in with everything else (CJ, p. 104).

Why is Elgin's view as so far set forth not simply a form of weak foundationalism, in which initially tenable claims function as basic beliefs? She cites two differences: unlike the justification that attaches to foundationalism's basic claims, initial tenability can be lost; it can also be augmented through coherence (CJ, p. 110).

It is not clear to me, however, that either of these features should be regarded as a prerogative of coherentists alone. In a typology of possible foundationalisms, Roderick Firth has suggested that the minimal tenet of foundationalism is simply this: basic beliefs have some measure of initial warrant that is not derived from coherence. This level of warrant may be increased by coherence with other statements or diminished, even to the vanishing point, by lack of coherence with other statements.[12] Firth's minimal view thus incorporates both of the features Elgin sees as antithetical to foundationalism.[13]

Nonetheless, I see two other questions on which foundationalists are apt to disagree with Elgin. First, are all commitments initially tenable, or only those in some specially marked out class? Second, are all commitments likewise revisable, or are some immune from subsequent rejection? On each question, Elgin takes the more egalitarian stand (CJ, pp. 101–102 and 121).

Using terminology from Michael Huemer, we may say that the first issue is the issue of phenomenal conservatism versus a more general doxastic conservatism (Huemer, 2001, pp. 99–115). Phenomenal conservatism is the view that if anything seems to be the case, one is prima facie justified in accepting it. The seemings can include perceptual seemings (it looks to me as if there is a red object over there), memorial seemings (I seem to remember being chased by a dog one day on my way to kindergarten), and intellectual seemings (it strikes me as self-evident that the relation of equality is transitive). Doxastic conservatism is the more sweepingly democratic view that anything the subject

believes has some presumption in its favor – the products of wishful thinking and superstition no less than the deliverances of perception and memory. Foundationalists are typically phenomenal conservatives, while Elgin (in her book) is a doxastic conservative.

In her contribution to this volume, Elgin no longer espouses doxastic conservatism, or at any rate holds that coherentists need not be committed to it. She says that perceptual deliverances may be assigned special weight, provided they do not have it a priori. They have it only in virtue of coherence considerations: by accepting the deliverances of perception, we get systems that remain coherent over time (that is, they continue to jibe with new deliverances of perception). But might it not likewise be true that by accepting fantasies, we get systems that remain coherent with future contents of fantasy – especially if the fantasizer has a one-track mind? So it seems to me that there is a privileging of perception presupposed and not explained by this coherentist underpinning for it.

I turn now to the second point: are all commitments on a par as regards the possibility of revision? Elgin holds, with Quine, that in the quest for reflective equilibrium, anything may be revised. There is no commitment that may not be sacrificed in order to maximize the tenability of the entire system. Here I would like to protest that there are certain principles of logic, at least, that cannot be given up, because they are framework principles without which coherence could scarcely be defined. What would happen if we gave up the law of non-contradiction? It is not clear that there could any longer be such a thing as what Quine calls a recalcitrant experience, forcing changes elsewhere in the system. If a new deliverance stood in contradiction to things we already accepted, we could simply accept it with a "What – me worry?" shrug.

Here is another difficulty for a coherentism that holds everything revisable, at least if we understand this as "anything could be justifiably rejected," symbolized as "(p) possibly J ~ p." Suppose there is some proposition q (the law of non-contradiction, perhaps?) whose truth is necessary for anything to belong to a coherent system and therefore necessary for anything to be justified. Since ~q entails that nothing is justified, we now have *possibly J (nothing is justified)*, which is absurd.[14]

A related question about the status of the rules and principles of logic and evidence is whether they have force only because a subject is committed to them. I gather Elgin would say yes, but I say no. Consider a system of beliefs containing the elements p, q, and $p \rightarrow \sim q$. Suppose it is not a sheer fact of logic, independent of anything the subject believes, that the system is inconsistent and in need of revision. Could we make the system intolerably inconsistent by adding the principle: if any two of {p, q, and $p \rightarrow \sim q$} are true, then the third must be false? No, for anyone who could live comfortably with the original system could live comfortably with the expanded system. Such is the lesson the tortoise taught Achilles (Carroll, 1895).

I wish to raise one more question about the role of logical and evidential relations in Elgin's coherentism. Elgin defines both coherence and equilibrium in terms of the relation "p is reasonable in light of q, r, and s." What is the required epistemic status of this relation (or of the logical and other relations on which it supervenes)? Must such relations be known to hold among the propositions in a system before the propositions are warranted for the subject? And if so, how does such knowledge arise?

I see three possibilities to consider in answer to these questions. First, the holding of coherence-constituting relations might be regarded as an external condition of

knowledge, making knowledge possible regardless of whether the subject knows that such relations hold. Although this seems to me a good way for a coherentist to go, I gather that Elgin would not find it congenial. On more than one occasion, she expresses her dissatisfaction with externalist stratagems in epistemology (CJ, pp. 22, 46, 51; this volume, p. 251). Second, it might be held that logical relations and other relations of support are known to hold because they are necessary relations, apprehendable a priori. But this, too, is an option Elgin would reject. In the first place, it would be a concession to foundationalism; in the second place, she has Quinean qualms about there being any propositions at all that are true necessarily and known a priori (CJ, pp. 40–46 and 53–57). Third, it might be held that coherence-constituting relations are known to hold because the propositions saying that they hold are part of the best overall coherent system that is reasonable in light of our antecedent commitments.[15] But that way lies an infinite regress. A proposition p affirming a relation of coherence would be justified only because the subject is justified in believing that p belongs to a coherent system. That belief in turn would be justified only because the subject is justified in believing that *p belongs to a coherent system* belongs to a coherent system, and so on. Even if it is the same system every step of the way, we still get a regress in which ever more complicated propositions must be believed with justification to belong to the system. The untenability of such a regress suggests to me that we should go instead with one of the first two options, agreeing with the externalist that coherence propositions need not be known at all or with the foundationalist that they are known because they are either basic propositions or propositions inferrable from basic propositions.

IV

I have not lived up to the title of this essay, for I have not offered a complete defense of moderate foundationalism. I hope nonetheless to have shown that an internalist coherentism cannot be a satisfactory theory of justification. We must be either externalists or foundationalists, and if we are foundationalists, our foundationalism must be of the moderate rather than the weak variety.

Notes

1 The distinction between positism and foundationalism is lost on those who cannot hear the word "justified" as anything but a past participle, implying that some act or relation of justifying has occurred whereby a belief is justified by something else that serves as a reason for it. For foundationalists, "justified" simply connotes a favorable epistemic status, which a belief may have even though the subject has no reason for it. In this connection, another term, such as "evident" or "credible," might be less misleading than "justified."

2 Haack summarizes her theory in Haack (2000). She does not classify foundherentism as weak foundationalism, believing it essential to foundationalism that the foundations not receive support from other elements in the structure.

3 Hereafter I shall cite this work in the text as AKV.

4 This paper was first published in 1857.

5 James Van Cleve, "Can coherence generate warrant?" (2011).

6 I discuss Huemer's work at some length in "Can coherence generate warrant?"
7 For foundationalists who hold that propositions about the physical world are always derived, not basic, that A says X would not be basic after all. It would rest, however, on deliverances that are basic.
8 For further discussion of Lewis's view, see Van Cleve (1977). For a method of assigning probabilities in relation to evidence without assuming that the evidence is certain (in the sense of having probability 1), see Jeffrey (1983).
9 By "conditionalizing" I mean drawing the inference "h has probability n given evidence e, and e is certainly true; therefore, h has probability n."
10 The term was coined by John Rawls (1971, p. 20f.). It refers to a state of affairs in which specific judgments and general principles have been brought into agreement with each other through a process of mutual adjustment.
11 Elsewhere Elgin adds other requirements. One is that there must not be a competing system that is more tenable overall (Elgin, 1996, p. 107). Another is that the best explanation of the system's coherence must be that it is at least roughly true ("Non-foundationalist Epistemology: Holism, Coherence, and Tenability," this volume).
12 See Firth (1964, pp. 466–467) for his progressively weaker foundationalisms.
13 Even so staunch a foundationalist as Roderick Chisholm incorporates Elgin's two allegedly distinctive features to some extent. Although the warrant belonging to Chisholmian "self-presenting" propositions can be neither increased nor diminished by their relations to other propositions, the same is not true of the warrant possessed by "indirectly evident" propositions. Their epistemic status is defeasible, and it may be raised by concurrence (Chisholm's term for coherence). See Chisholm (1977, chapters 2 and 4).
14 In fairness to Elgin, I note that she may give special status to the law of non-contradiction. She says that tenable systems must at least be logically consistent, or else be all-entailing (CJ, p. 136). Against this rationale, however, I note that in revisionary relevance logic, contradictions do not entail everything.
15 Perhaps this is what coherentists should say in response to Russell's question about how the laws underlying coherence are known.

References

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

Boole, G. (1952) On the application of the theory of probabilities to the question of the combination of testimonies or judgments. In *Studies in Logic and Probability*. London: C.A. Watts.

Carroll, L. (1895) What the tortoise said to Achilles. *Mind* 4: 275–280.

Chisholm, R.M. (1964) The myth of the given. In *Philosophy*, ed. R.M. Chisholm *et al*. Englewood Cliffs, NJ: Prentice Hall.

Chisholm, R.M. (1977) *Theory of Knowledge*, 2nd edn. Englewood Cliffs, NJ: Prentice Hall.

Elgin, C. (1996) *Considered Judgment*. Princeton, NJ: Princeton University Press.

Firth, R. (1964) Coherence, certainty, and epistemic priority. *Journal of Philosophy* 61: 545–557. Reprinted in *Empirical Knowledge*, ed. R. Chisholm and R.J. Swartz. Englewood Cliffs, NJ: Prentice Hall, 1973.

Goldman, A. (1979) What is justified belief? In *Justification and Knowledge*, ed. G. Pappas. Dordrecht: Reidel.

Haack, S. (1993) *Evidence and Inquiry*. Oxford: Blackwell.

Haack, S. (2000) A foundherentist theory of empirical justification. In *Epistemology*, ed. E. Sosa and J. Kim. Oxford: Blackwell.

Huemer, M. (1997) Probability and coherence justification. *Southern Journal of Philosophy* 35: 463–472.
Huemer, M. (2001) *Skepticism and the Veil of Perception*. Lanham, MD: Rowman & Littlefield.
Jeffrey, R. (1983) *The Logic of Decision*, 2nd edn. Chicago: University of Chicago Press.
Lewis, C.I. (1946) *An Analysis of Knowledge and Valuation*. LaSalle, IL: Open Court.
Rawls, J. (1971) *A Theory of Justice*. Cambridge, MA: Harvard University Press.
Russell, B. (1948) *Human Knowledge: Its Scope and Limits*. New York: Simon & Schuster.
Schlick, M. (1973) The foundation of knowledge. In *Empirical Knowledge*, ed. R. Chisholm and R.J. Swartz. Englewood Cliffs, NJ: Prentice Hall.
Van Cleve, J. (1977) Probability and certainty: a re-examination of the Lewis–Reichenbach debate. *Philosophical Studies* 32: 323–334.
Van Cleve, J. (2011) Can coherence generate warrant *ex nihilo*? Probability and the logic of concurring witnesses. *Philosophy and Phenomenological Research* 82: 337–380.
Wittgenstein, L. (1969) *On Certainty*. Oxford: Blackwell.

Reply to Van Cleve

Catherine Z. Elgin

Foundationalists consider some beliefs basic. A subject is justified in holding them "even in the absence of any justifying reason for them" (p. 256).[1] Other beliefs are justified by being suitably related to basic beliefs. Ostensible memories, beliefs about one's current experiences, and the fruits of introspection standardly count as basic. Strong foundationalists consider basic beliefs certain. Moderate foundationalists hold that "the non-inferential warrant possessed by basic beliefs need not amount to absolute certainty ... but it must be 'sufficient by itself to satisfy the adequate-justification condition for knowledge' (BonJour, 1985, p. 26)' (p. 257). Weak foundationalists set the level of credibility considerably lower. Despite its title, James Van Cleve's paper is not so much a defense of moderate foundationalism as an attack on what he takes to be its chief rivals – coherentism and weak foundationalism. Playing offense rather than defense is a reasonable argumentative strategy; but his failure to defend the position may suggest that moderate foundationalism is epistemologically unproblematic. I will argue that this is not so, and defend coherentism and/or weak foundationalism against Van Cleve's objections.

Each component in a coherent system supports and is supported by many others. Foundationalists acknowledge that coherence amplifies credibility, but insist that amplification can occur only where credibility is there to be amplified. If they are right, some measure of initial credibility is needed. Their only question is: How much? An unamplified sound retains its original volume. If coherence functions as a credibility amplifier, then a belief that fails to cohere with other beliefs should retain whatever level of credibility it began with. But failure to cohere deprives claims of credibility. Perhaps "amplification" is a bad metaphor. Still, one wants to know why failure to cohere discredits a claim.

Beliefs satisfying the adequate justification condition can have a probability of less than 1. If knowledge is closed under obvious implication, a subject who knows that *p*

and knows that q is evidentially in a position to know that $p\&q$. But if p and q are evidentially independent and each has a probability of less than 1, the probability of the conjunction is less than the probability of each conjunct. The more we conjoin, the lower the probability, eventually falling below the threshold for knowledge. Moderate foundationalism thus has reason to disvalue comprehensiveness. Granted, more complicated relations among propositions can enhance probability. But the closure principle for conjunction, which seems obvious and benign, is threatened.[2]

Moderate foundationalism might preserve closure by assigning probability 1 to the contents of all beliefs that satisfy the adequate justification standard. Then there is no doubt to aggregate. Adler (2002, pp. 250–255) argues that such a move is acceptable if we distinguish between credibility and confidence, and recognize that we can be less than fully confident in a belief to which we assign probability 1. Another alternative is to deny that the probability calculus provides the metric for credence. Whatever solution moderate foundationalists choose, they should acknowledge that their position has untoward consequences.

Weak foundationalism maintains that deliverances – items that present themselves as candidates for belief – have a slight measure of initial credibility, and that credibility is enhanced through their incorporation into a coherent system of beliefs. Coherentism denies that initial credibility is needed. I construed my position as a very weak form of foundationalism, since I was marking out the epistemological difference between deliverances and propositions a subject has no inclination to believe. On Van Cleve's characterization, my position is coherentist. Coherentists, he says, maintain "that beliefs can be justified in virtue of relations of mutual support" (p. 256). If we restrict the considerations we are evaluating to beliefs (or deliverances), we need not introduce initial credibility. Within the realm of beliefs and deliverances, coherence is enough. Inasmuch as Van Cleve argues against both coherentism and weak foundationalism, this is a point of clarification, not an objection, to what he says. But it is worth noting that the difference between coherentism and weak foundationalism turns on where we start.

Corroboration by independent witnesses can raise probability above the threshold for knowledge, even if the initial credibility of their testimony is slight, so long as there are enough witnesses and they have enough choices. This favors weak foundationalism. But, Van Cleve notes, Huemer's (1997) proof applies only to agreement on a single proposition. The sort of coherence we are interested in involves looser relations of support among distinct beliefs. So weak foundationalists cannot assume that Huemer has settled the debate in our favor. Perhaps an analogous theorem could be proved given a suitably rigorous description of the web of beliefs. Perhaps not. But if Huemer's result is relevant to the larger debate, it provides the weak foundationalism with grounds for optimism.

Van Cleve demurs. Even if corroboration vindicates testimony with low initial credibility, he says, we still have to know that the testimony actually takes place. We might after all be dreaming. "A good case can be made that there must be at least high intrinsic credibility – perhaps high enough to constitute knowledge – attaching to the facts that such-and-such cognitive states (be they experiences, ostensible memories, or beliefs at large) are actually taking place" (p. 261). I am not convinced. If the dream possibility is a skeptical challenge, it applies to Van Cleve's position as much as to mine. If not, then, as Descartes argued in Meditation V, we have reasons, albeit less than certain reasons, to believe that we are not dreaming.[3] The reasons need not be

more than initially credible; for coherence leverages weakly supported deliverances when the best explanation of their occurrence is that things are at least roughly as they seem. Even putatively basic beliefs may be far from sufficient for knowledge. I may be unsure whether I seem to hear a distant train whistle or only imagine that I hear it. If no one else on the platform reacts, no train arrives, and I recall my propensity for wishful thinking, I conclude that I imagined the sound; I did not even seem to hear it.

Van Cleve objects that coherence depends on empirical generalizations. Coherence cannot be what generates knowledge, since we need knowledge of particulars to justify those generalizations. But the tenability of a generalization does not rest on its etiology. A generalization is initially credible if it is a deliverance. If it is true and is suitably interwoven into a comprehensive web of beliefs, it is known. But it need not be already known to be a candidate for incorporation, and in the process of devising a system its credibility may both influence and be influenced by other things the subject is inclined to believe.

Although Quine contends that any commitment can be revised, he does not maintain that all are equally good candidates for revision. He advocates a principle of minimal mutilation: in revising, we should retain as many of our central beliefs as possible. But even central beliefs may require revision. "Inanimate objects are identical when their parts are identical" is a fundamental metaphysical principle. But if it holds universally, then "F=ma" does not. Molecules in a viscous fluid move at different rates. In prototypical applications of "F=ma," forces act on objects like billiard balls that have a well-defined boundary. In viscous fluids, the "forces" on the "object" are effects on the momentum of molecules moving in and out of that "object." Retaining "F=ma" requires continually redefining what constitutes a single particle, letting different molecules comprise it at different times. Fluid mechanics preserves the law. Rather than insisting that all component molecules of a particle be the same from one instant to the next, they let the individual molecules come and go, but keep the average enclosed mass constant (Wilson, 2006, pp. 158–159). "F=ma," evidently, is so central a law of physics that scientists are willing to radically revise the criteria for the identity of a fluid particle over time in order to preserve it. Quine's position explains why this is reasonable.

What about the law of non-contradiction? Could it ever be rational to revise it? Paraconsistent logicians say "yes" (Priest, 2002). They contend that the benefits of rejecting the law of non-contradiction outweigh the costs. I disagree, not because I think the law is unrevisable, but because I think that other revisions better accommodate the problem cases.

Even if the law of non-contradiction lies at the very center of the web of belief, so that it is always preferable to make revisions elsewhere in the system, the same does not hold of all the principles of classical logic, much less of the principles of evidence. The history of logic is rife with debates about which principles ought to be accepted. The law of excluded middle remains the subject of controversy. Evidential principles are even less secure. We endorse visual deliverances and reject premonitions not because the former are basic and the latter are not, but because visual deliverances have, and premonitions lack, impressive track record, and we have an acceptable theory of vision, but no clue how premonitions could connect to their subject matter. The support by an overarching, tenable system vindicates vision and discredits premonition.

Contra Van Cleve, fantasy poses no threat to coherentism. So long as the fantasizer realizes that he is fantasizing, the thoughts he entertains are not deliverances, hence

lack initial credibility. Confabulation is a more serious worry. A confabulator composes a coherent narrative by unconsciously ignoring, bracketing, or downplaying considerations that detract from the account he seeks to construct, and accepting unwarranted considerations that support it. Clearly a confabulated account has no claim to epistemic respectability. But self-deception undermines coherence (Adler, 2002, pp. 74–101). Confabulators achieve local coherence by sacrificing coherence across a broader range of beliefs. Suppose a father deceives himself into thinking that his daughter's dismal grades result from her teachers' failure to recognize her quirky brilliance. To sustain his belief, he overlooks factors he otherwise considers relevant to student performance – terrible study habits, unfinished homework, ignorance of the most basic facts about the subject matter. He violates his otherwise accepted methodological beliefs about how to judge such matters. He would not invoke such considerations to account for his paperboy's poor grades. He thus weakens the overall coherence of his belief system by carving out exceptions for a special case. Because self-deception weakens coherence, it does not undermine the contention that epistemic justification is grounded in coherence.

Delusions might seem to pose a stronger challenge, for they are more coherent than self-deceptive beliefs. A mental patient who believes he is Napoleon in exile interprets all his experiences in terms of his delusion. He takes nurses, doctors, and aides to be lackeys, courtiers, and guards; visitors to be loyal subjects; those who refuse to do his bidding to be traitors; those who tell him he is ill or mistaken in his beliefs to be part of the plot to prevent his retaking the throne. Unlike the self-deceptive father, he does not carve out special exceptions. He explains all his experiences in terms of the delusion. Whenever a tension occurs, he rejects the deliverance that conflicts with his identification with Napoleon. Although his beliefs are more coherent than the self-deceptive father's, his delusion would have to be extensive to achieve any reasonable level of coherence. He would, for example, have to believe that he and everyone around him was speaking French rather than English. He would have to believe that oil lamps or candles rather than electric lights were illuminating the area. He would have to believe that the vehicles he sees are horse-drawn carriages rather than cars and trucks. Still, there would be many everyday experiences for which he had no explanation – riding an elevator, watching television, or using a telephone. Such gaps deprive his worldview of the level of coherence that normal epistemic agents regularly achieve. Self-deception and delusion might seem to undermine coherentism. People, it seems, can manipulate themselves into believing nearly anything. But they cannot easily incorporate self-deceptive or delusive beliefs into a comprehensive, coherent belief system. Coherence, which might look like an unduly weak constraint, turns out to be a very strong one.[4]

I conclude that Van Cleve's objections to weak foundationalism and coherentism are not compelling. There is good reason to believe that through systematizing, we weave initially tenable deliverances into epistemically tenable webs of belief.

Acknowledgments

I am indebted to Jonathan Adler and Samuel Elgin for useful comments on early drafts of this paper.

Notes

1 Page numbers in parentheses refer to James Van Cleve's essay, "Why Coherence Is not Enough: A Defense of Moderate Foundationalism," this volume.
2 In discussing the lottery paradox, Kyburg (1997) argues that because of the aggregation of doubt, we should reject closure under conjunction. So perhaps the principle is not so benign as it looks.
3 I am grateful to Jonathan Adler for strengthening my argument here.
4 I am grateful to Jonathan Adler for helping me to articulate this point.

References

Adler, J. (2002) *Belief's Own Ethics*. Cambridge. MA: MIT Press.
BonJour, L. (1985) *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.
Huemer, M. (1997) Probability and coherence justification. *Southern Journal of Philosophy* 35: 463–472.
Kyburg, H. (1997) The rule of adjunction and reasonable inference. *Journal of Philosophy* 94: 109–125.
Priest, G. (2002) Paraconsistent logic. In *Handbook of Philosophical Logic*, 2nd edn, Vol. 6, ed. D. Gabbay and F. Guenthner, pp. 287–393. Dordrecht: Kluwer.
Wilson, M. (2006) *Wandering Significance*. Oxford: Clarendon Press.

Reply to Elgin

James Van Cleve

I select five points for further debate.

Should adequate justification be conjunctive? Moderate foundationalists hold that basic beliefs need not be certain; they need only have a level of justification adequate for knowledge. Elgin notes as a potential problem for moderate foundationalism that it may have to renounce the conjunctivity of adequate justification. If we do epistemology within the framework of the probability calculus, to say that adequately justified beliefs need not be certain is to say that they need not have probability 1. It follows that a conjunction of adequately justified independent beliefs may fail to be adequately justified, since the probability of the conjunction will be given by multiplying several numbers less than 1, and the product may be less than whatever threshold is set for adequacy.

I agree with Elgin that adequacy should be conjunctive. She notes two ways for avoiding the result to the contrary: assign probability 1 to adequate justification after all, or deny that the probability calculus provides the metric for justification. I favor the latter alternative, for one and perhaps two reasons. First, one might construe adequate justification as a qualitative notion, not measurable by any number.[1] Second, even if one did assign a number to the level of justification adequate for knowledge,

one could very well deny that the laws of probability are the laws of rational belief. I cite two points of discrepancy. (i) In the probability calculus, any necessary truth (no matter how recondite and remote from our knowledge) gets assigned probability 1. By the same token, whatever is entailed by a proposition has a probability at least as high as what entails it. As often noted, these are utterly unrealistic as laws governing actual degrees of rational belief. Necessary truths may be far from certain in any epistemic sense, and the necessary consequences of the obvious may not be obvious. (ii) It is a theorem of the probability calculus that if e1 and e2 are pieces of evidence with probability 1, then for any proposition h, P(h)=P(h, e1)=P(h, e2). Let e1=I am sensing redly, e2=I am hearing F-ly, and h=there is a red object in front of me. Surely e1 lends a higher degree of rational support to h than e2 does, yet the conditional probability of h on either piece of evidence is the same.[2]

Is everything in principle subject to revision? Elgin holds with Quine that the answer is yes. I had suggested that the law of non-contradiction is an exception, being essential if recalcitrant experiences are to force revisions in the web of belief. Elgin cites the work of Priest, who has advocated giving up the law of non-contradiction. Yes, he is a *dialetheist*, someone who holds that some contradictions are true. But he is not a *trivialist*, someone who holds that everything, including all contradictions, is true. In paraconsistent logic, one may accept some contradictions without accepting them all (Priest and Tanaka, 2009). I suggest that it could never be rational to reject the statement *not all contradictions are true*.

Must the facts of deliverance have high initial credibility? Elgin (whom I should have classified as a weak foundationalist) maintains that deliverances need have only a small amount of initial credibility, much less than would be required for knowledge. I take it she is referring to the *contents* of deliverances – if I have an ostensible memory that *p* or seem to hear testimony that *p*, *p* thereby acquires a slight amount of credibility for me. I had maintained that the *facts* of deliverance – that I *do* ostensibly remember that *p*, for instance – must have *high* credibility, perhaps even enough for knowledge. If there is not adequate justification for believing the deliverances are actually occurring, their coherence counts for little. Elgin replies that if I am raising the possibility of dreaming as a skeptical challenge, it threatens my own position. I say not so, for deliverances as I conceive of them are seemings to see, presentations of *p* as past, and so on, which (as Descartes notes in Meditation 2) can be known to be occurring dream or no dream.

Are particulars epistemically prior to generals? In developing a point of Russell's, I had presupposed a kind of epistemic priority: contingent general truths are known or adequately justified only if one first has adequate justification for particular propositions from which they may be inferred. Elgin does away with all such priority. Once particulars and generals are all there in the web of belief (however they got there), each has its modicum of initial credibility, and they all rise together to the level of knowledge when they are suitably coherent. Their coherence is explicated in part as their each being reasonable in light of the others. How is it known that a generalization is reasonable in light of its instances, or the instances reasonable in light of the generalization? Is this perhaps an a priori matter? Or is it an external condition of knowledge, making knowledge possible regardless of whether the subject has any justification for it? If neither of the foregoing, it is presumably a proposition that must already have been in the web alongside the particulars and the generals, getting its positive status

just as they do, some of it just by being delivered and more by cohering with the rest. But now we may inquire whence derives the reasonableness of it given its companions, and a regress looms.

Are epistemic principles on a par with everything else in the web? An epistemic principle is a general principle specifying the conditions under which various propositions have various levels of justification for a subject. Elgin says "yes" to our question, but I think the considerations of the previous paragraph show that some epistemic principles must have a different status, either being external conditions of knowledge or, if requiring to be known, being known a priori. Consider two of Elgin's own epistemic principles – that any deliverance has a slight initial amount of credibility, and that a belief has a high amount of credibility if it is a member of a coherent set whose coherence is best explained by supposing its members true. If these are neither a priori nor external, they must have credibility for the subject that they derive by being members of a coherent set – a set whose members are reasonable in light of each other. If D is reasonable in light of A, B, and C, that is presumably because there is an epistemic principle to the effect that if S believes (with justification?) each of A, B, and C, then S is justified in believing D. If this principle is neither a priori nor external, the subject must be justified in believing it. Whence that justification? As before, a regress looms.

Notes

1 This is the approach taken in Chisholm (1977). Of his five epistemic categories, only two are representable in the probability calculus, the top category (certainty) implying probability 1 and the bottom category (having some presumption in its favor) implying probability greater than 0.5. Being evident, the level of justification required for knowledge, is assigned no number.

2 So is the degree of confirmation of h on either piece of evidence, defining this either as the ratio P(h,e)/P(h) or the difference P(h,e) – P(h).

References

Chisholm, R.M. (1977) *Theory of Knowledge*, 2nd edn. Englewood Cliffs, NJ: Prentice Hall.

Priest, G. and Tanaka, K. (2009) Paraconsistent logic. In *The Stanford Encyclopedia of Philosophy* (Summer 2009 edn), ed. E.N. Zalta. http://plato.stanford.edu/archives/sum2009/entries/logic-paraconsistent/ (last accessed March 14, 2013).