

Coherentism via Graphs

Selim Berker
Harvard University
sberker@fas.harvard.edu

[Penultimate draft of a paper that was published in
Philosophical Issues 25 (2015): 322–52; please cite that version.]

1. Introduction

Some of our beliefs are justified, if indeed they are justified, only because they stand in inferential relations to other beliefs. And some of those latter beliefs are justified, if indeed they are justified, only because they stand in inferential relations to yet other beliefs. Where does this regress end? As the traditional story goes, there are four main options. According to *foundationalists*, the regress terminates in a set of beliefs each of whose justification does not depend on its inferential relations to any other beliefs. According to *infinitists*, the regress goes on without end. According to *coherentists*, the regress eventually loops back on itself. And according to *skeptics* of a certain sort, because none of these alternatives is satisfactory, all of our beliefs are unjustified.

It is customary to use *directed graphs* to depict these different takes on the justificatory structure of a subject's beliefs at a given time. For instance, we might draw a directed graph in which the nodes represent our subject's beliefs and the arrows represent inferential relations between those beliefs, like so:

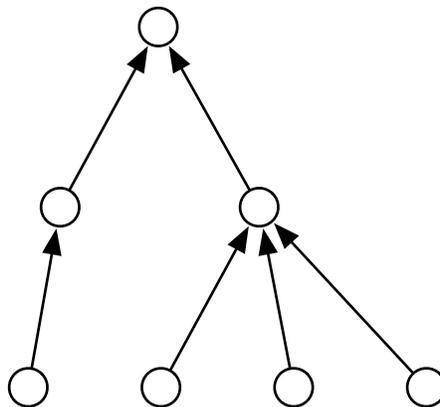


Figure 1

Directed graphs of this sort are an extremely useful—indeed, almost indispensable—tool for thinking about the structure of justification. But despite the ubiquity of directed graphs in discussions of

foundationalism, infinitism, and coherentism, the full power of this tool is rarely harnessed. A number of important distinctions are often overlooked or ignored in an effort to keep things simple.

For example, although the lines of support for actual beliefs held by actual subjects are usually the sort of ramifying mess found in Figure 1, authors frequently clean up this mess in order to make the regress of justification easier to state. As I will put it, these authors replace *the ramifying regress*, in which some beliefs are supported by multiple beliefs, with *the serial regress*, in which each belief is supported by at most one belief. The standard way of doing this is to assume that whenever a belief in P is supported by a belief in Q_1 and a belief in Q_2 , this is justificatorily equivalent to a case in which a belief in P is supported by a single belief in the conjunctive proposition $Q_1 \& Q_2$.¹

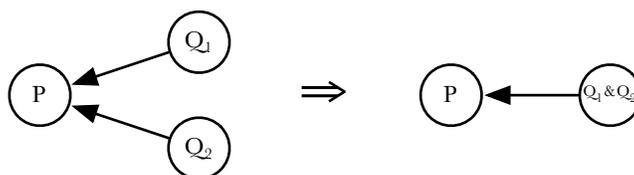


Figure 2

With this assumption in place, the basic issue dividing foundationalists, infinitists, and coherentists can be stated all too easily: which of the following justificatory structures does one prefer?

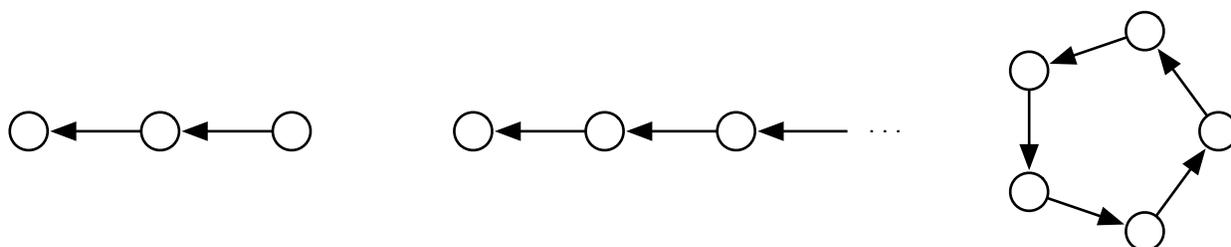


Figure 3

However, this simplification of the issues is not without its costs. For once the non-skeptical alternatives have been portrayed in this manner, the coherentist option looks, to put it bluntly, stupid—a wholly unattractive position that only a truly desperate philosopher would endorse. How can a mere circle of support, of whatever length, make a belief appropriate to hold? It is no coincidence that coherentists tend to prefer the metaphor of *a web* to that of *a circle*. For most coherentists, it is crucial to keep in the structure

that gets washed out when we assimilate a ramifying regress into a serial regress.

What follows shall essentially be a more detailed elaboration of this basic point. I will first introduce three distinctions that it is important to keep track of when using directed graphs to represent justificatory structure. Then I will demonstrate the importance of these distinctions by showing how they can help us solve a variety of puzzles that arise when attempts are made to formulate a viable form of coherentism. In particular, I think these distinctions give us the tools both (a) to respond to several important objections to coherentism and (b) to provide that most elusive of elements in a coherence theory of justification, namely a substantive, non-metaphorical characterization of what coherence amounts to.

So really this essay has two independent goals. First, it is a defense of the importance of paying attention to three distinctions when using directed graphs to depict justificatory structure. This is a topic that should be of interest to all epistemologists—foundationalists, infinitists, coherentists, and skeptics alike. (Although the debate among these camps is not as fashionable as it once was, this is not an issue that is going away: every epistemological theory must take a stand on the regress of justification, even if only implicitly.) Second, this paper is a defense of what we might call a graph-theoretic approach to the notion of coherence at the heart of coherentism. This approach is in marked contrast to the two approaches that dominate contemporary discussions of coherentism. According to the first, coherence is formulated in terms of probabilistic relations among the contents of one's beliefs.² According to the second, coherence is determined by a weighted sum of various factors that come in degrees, such as the degree of connectedness among one's total system of belief or the degree to which that system does not contain anomalous beliefs.³ Although I will only briefly touch in the current essay on why I think so, in my opinion both of these approaches are, without supplementation, inherently limited in what they can tell us about the nature of coherence. A much more promising approach, I will suggest, is to take coherence to be a graph-theoretic property of the structure of support for a given belief.

One last comment before I proceed. There are two main ways of representing a directed graph. The first is pictorially, using the sorts of drawings found in Figures 1–3. The second way, preferred by mathematicians, is to characterize a directed graph in terms of an ordered pair whose first component is a

set of objects—the graph’s vertices (or nodes)—and whose second component is a set of ordered pairs of those objects—the graph’s directed edges (or arcs). For example, $\langle \{v_1, v_2, v_3\}, \{ \langle v_2, v_1 \rangle, \langle v_3, v_2 \rangle \} \rangle$ is a way of representing the leftmost graph in Figure 3 via this second method. Since the current essay is intended as a non-technical introduction to the importance of observing certain distinctions when thinking about the justificatory regress, as much as possible I restrict myself to pictorial means of characterizing directed graphs. I make full use of the mathematical theory of directed graphs, in all its abstract glory, in a more technical companion piece to this article.⁴

2. Preliminaries

Let me start with some clarificatory remarks.

First, the main normative notion that I will be focusing on in this essay is *justified belief*. Many of the issues to be discussed here arise for other normative categories in epistemology, such as *reason for belief*, *evidence*, and *knowledge*. Some justified beliefs are inferential in nature; but so, too, are some reasons for belief, bits of evidence, and pieces of knowledge based in inferential relations. The former gives rise to a regress of justification; but so, too, do the latter give rise to analogous regresses of reasons, of evidence, and of knowledge. I have fixed on justification as my normative category of choice in this essay mainly for familiarity’s sake: almost all of the most important discussions of epistemic regresses in the post-Gettier era have been framed in terms of “justification”-talk. Variants of every point I will be making hold for the other normative categories as well.

Second, by “justification” I mean “sufficient overall epistemic justification.” (I almost always omit these qualifiers for brevity’s sake.) In particular, I will not be considering the conditions under which one might have a *practical* rather than an *epistemic* justification for believing something. Similarly, I will concentrate my discussion on whether one has *sufficient overall* justification to believe a certain thing, not on whether one has *pro tanto* (or *prima facie*) justification to believe that thing. In many contexts this second distinction is of the utmost importance, but suppressing it for the purposes of the current essay will bring the structural distinctions I want to make into better focus.

Third, in this essay I will only be considering the justificatory structure of the beliefs held by *a*

single person at a single point in time. For other purposes, it might be useful to consider directed graphs in which there are nodes for the beliefs of more than one person. Moreover, according to some epistemological theories, it is inferential relations between one's beliefs at different times that are relevant to the justificatory regress. However, variants of the structural points I want to make also arise for these social and diachronic approaches to the structure of justification, so it will be harmless to frame my discussion within the confines of an individualistic, synchronic approach to justification and its grounds.

Fourth, I will be helping myself to notions of deductive, inductive, and abductive inferential support between propositions without pausing much to ask what these things come to. There are difficult questions here that need answering. What exactly is it for one proposition to deductively entail another? What is it for a set of propositions to provide adequate inductive support for another? What differentiates abductive from inductive support, if indeed they are distinct varieties of support? Since my main concern is elsewhere, I will simply assume that satisfactory answers to these questions can be provided.

Fifth, in addition to talk of deductive, inductive, and abductive support relations between *propositions*, it will be useful to have a way of referring to deductive, inductive, and abductive support among *beliefs*. When (i) subject S believes P, Q_1, \dots, Q_n , and (ii) P is deductively, inductively, or abductively supported by Q_1, \dots, Q_n , I will say that (iii) S's belief in P *is supported by* S's belief in Q_1, \dots, Q_n .⁵ This terminology is fairly standard in the coherentism literature. Consider, for example, the common claim that coherent beliefs must be mutually supporting: it is difficult to make sense of this idea without defining support among beliefs roughly as I have.⁶

My final clarificatory remark is perhaps the most important: I will only be considering varieties of positive support in this essay. This is a gross oversimplification, since just as beliefs can support each other via deductive, inductive, and abductive links, so too can beliefs be in tension with one another. (Think, for example, of contradictory beliefs.) However, matters will already get complex enough when we restrict ourselves to an investigation of positive support relations, so I leave it as a task for another occasion to extend the framework developed here so as to take into account negative structural relations among beliefs.

3. Three Distinctions

With these preliminaries in place, I can now explain the three distinctions that will be the main focus of our discussion. Whenever one draws a directed graph for philosophical purposes, it is absolutely essential to ask, “What exactly do the nodes represent? And what exactly do the arrows represent?” The answer to the first of these questions might seem easy in the case at hand. Surely, it might be insisted, the nodes in our graphs of justificatory structure represent our subject’s beliefs. More generally, we might take the nodes to represent any of the following doxastic states or entities:

- beliefs (either occurrent or standing);
- suspensions of judgment;
- credences;
- dispositions to be in one of the above doxastic states;
- propositions accessible to one of the above doxastic states.

I agree that it is important to include doxastic nodes of these sorts in our directed graphs. However, we need not stop there. For there is nothing to bar us from incorporating nodes that represent other sorts of justificatorily-relevant entities into our graphs. A prime candidate: many theorists hold that a subject’s perceptual experiences are justificatorily relevant (in these sense that they either partially or entirely make it the case that the subject is justified in believing something), and if this is so, then it will be useful to have a way of including perceptual experiences in our diagrams of justificatory structure. Similarly, some theorists hold that (i) intuitions are not a species of belief, and (ii) intuitions are justificatorily relevant. And so on: for every kind of mental entity which a theorist takes to ground justification, we can include a distinctive sort of node for that kind of entity. Such non-doxastic nodes could include any of the following:

- perceptual experiences;
- intuitions;
- graspings of meaning;
- memory impressions;
- emotions.

This leads to the first distinction on which I want to focus:

distinction #1 (doxastic nodes vs. non-doxastic nodes): In addition to nodes representing doxastic states (such as beliefs), we can include in our directed graphs nodes representing non-doxastic states (such as perceptual experiences).

In what follows, I will take standing beliefs to be my primary example of a doxastic node, and I will take perceptual experiences to be my primary example of a non-doxastic node. (Hereafter I drop the “standing” and “perceptual” qualifiers.) I will use a convention whereby a belief node is designated by a circle and an experience node is designated by a star. I take no stand on whether experiences have propositional or conceptual content, content of a non-propositional or non-conceptual sort, or no content at all.

Eventually I will be showing how distinction #1 is of use to coherentists. But an immediate application of this distinction arises in the case of foundationalism. As has often been remarked, there are three main accounts that foundationalists give of *basic beliefs* (the non-inferentially or immediately justified beliefs upon which, according to foundationalists, all other justified beliefs rest): some foundationalists hold that basic beliefs are self-justifying; some foundationalists hold that basic beliefs are non-justified⁷ justifiers (i.e. they themselves are not in need of justification even though they can confer justification on other beliefs); and some foundationalists hold that basic beliefs are justified in virtue of standing in a certain sort of relation to some state such as an experience which acts as a non-justified justifier. The difference between these three varieties of foundationalism, at least for an individual line of support, can be neatly summarized by the following three directed graphs:

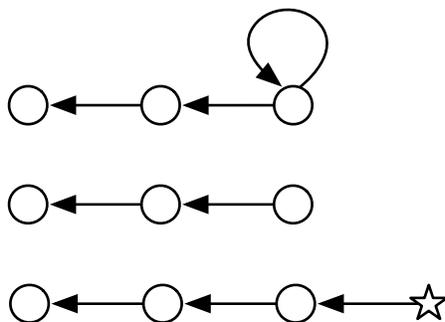


Figure 4

Without our first distinction, it would be impossible to graphically represent the difference between the second and third of these forms of foundationalism.

Earlier I said that whenever one employs directed graphs for philosophical purposes, it is crucial to ask two questions: “What exactly do the nodes represent? And what exactly do the arrows represent?” Let us turn now to the second of these questions. It might seem that the answer to this question is also easy—in fact, it might seem that we have already answered it. Don’t the arrows just represent relations of either deductive or inductive or abductive support (or some combination thereof)? One reason that this isn’t quite right has to do with the distinction just mentioned. In the bottom graph in Figure 4, the arrow between the experience node and the rightmost belief node isn’t deductive, inductive, or abductive; instead it is some sort of *sui generis* relation of perceptual support, as we might call it. So we should be more expansive in our conception of the varieties of support that an arrow can represent: these might include relations of deductive support, inductive support, abductive support, perceptual support, and perhaps other varieties of support as well (such as intuitive support and introspective support, if there be such), in addition to combinations of these support relations.⁸

However, there is a deeper reason why we have not yet fully answered our second question. Sometimes, in directed graphs of justificatory structure, the arrows are used to represent *relations of support*, and sometimes they are used to represent *the flow or transmission of justification*; in fact, they are usually taken to represent both at the same time. For example, suppose the following is one part of a given directed graph of justificatory structure:

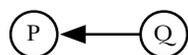


Figure 5

(When needed, I include in my graphs a name for the proposition believed inside each belief node.) This part of a graph is frequently taken to represent the fact that our subject S’s belief in proposition P is deductively, inductively, or abductively supported by S’s belief in proposition Q. But just as frequently, and often at the same time, this part of a graph is taken to represent the fact that, because S’s belief in Q is justified and because S’s belief in Q inferentially supports S’s belief in P,⁹ S’s belief in P is justified:

justification flows from one belief to the other.

This dual role for a single feature in a graph is unfortunate, especially since distinct claims are being made here: a relation of support is not the same as a flow of justification, even if the latter always obtains (partially) in virtue of the former. For example, suppose $Q = \langle \text{I will live forever} \rangle$ and $P = \langle \text{I will live for two hundred more years} \rangle$. Q deductively supports P , but it is easy enough to imagine a case in which, although I believe both Q and P , I am—alas—not justified in believing either. But if my beliefs in P and in Q are not justified, there cannot be a transmittance of justification between them: justification juice cannot flow from one to the other when it is not present in either. So this is a case in which although Figure 5 would be part of a complete diagram of the justificatory structure of my beliefs were the arrow interpreted *as a relation of support*, it is not part of that diagram when the arrow is interpreted *as a flow of justification*. Hence the second distinction to which I think it is important to pay attention:

distinction #2 (relations of support vs. flows of justification): Arrows representing relations of support between nodes should be distinguished from arrows representing flows of justification.

To avoid any ambiguity on this front, I will be using a convention whereby (i) arrows in my graphs always designate relations of support; (ii) a shadow under a node signifies that it is justified; and (iii) a shadow under an arrow signifies that justification is flowing or being transmitted along that line of support. (Note that a lack of a shadow in a diagram does not necessarily signify a lack of justification or a lack of justificatory flow: sometimes I will not be concerned with marking the presence or absence of justification in a diagram.) Thus the following graph:



Figure 6

is a compact way of representing all of the following: S believes P (the left circle); S believes Q (the right circle); S's belief in Q supports her belief in P (the arrow); S's belief in P is justified (the shadow under the left circle); S's belief in Q is justified (the shadow under the right circle); and S's belief in P is justified *because* S's belief in Q is justified and supports her belief in P (the shadow under the arrow).

Distinction #1 concerns the interpretation of nodes in directed graphs depicting justificatory

structure, and distinction #2 concerns the interpretation of arrows in such graphs. The last distinction I want to stress concerns both nodes and arrows at once. When we speak of two belief nodes supporting a third belief node (whether deductively, inductively, or abductively), this sort of talk is ambiguous. Sometimes we mean that each of the two belief nodes *independently* supports the third, as when a belief in P and a belief in Q each on their own deductively support a belief in the disjunction $P \vee Q$. And sometimes we mean that the two belief nodes *jointly* support the third, as when a belief in P and a belief in $P \supset Q$ together deductively support a belief in Q. These are different structural relations between nodes, and it is crucial that we be able to represent this difference in our directed graphs, like so:



Figure 7

The left of these I call “v-support,” and the right I call “y-support.” (To see why, tilt your head to the right.) I will also use “v-support” and “y-support” to refer to generalizations of these support structures in which three or more nodes each independently support a node (in the case of v-support) or together support a node (in the case of y-support). This gives us

distinction #3 (v-support vs. y-support): We should have a way of depicting the difference between several nodes each independently supporting some node (v-support) and several nodes together supporting some node (y-support).

In many ways, this third and final distinction is the crux of my graph-theoretic approach to formulating and defending coherentism.

Four comments about the v- vs. y-support distinction are in order. First, once we include instances of y-support in our graphs of justificatory structure, then strictly speaking we are no longer drawing “directed graphs,” as mathematicians use that term. Rather, we are now drawing what mathematicians call “directed hypergraphs”: a generalization of directed graphs in which the directed edges can have multiple tails or multiple heads. Second, in addition to the distinction between relations of v- and y-support, there is an analogous distinction between v- and y-flows of justification. A case in which the

following two facts obtain:

[S's belief in P is justified *because* (S's belief in Q₁ is justified *and* S's belief in P is supported by S's belief in Q₁)],

[S's belief in P is justified *because* (S's belief in Q₂ is justified *and* S's belief in P is supported by S's belief in Q₂)]

is a typical example of a v-flow. By contrast, a case in which the following single fact obtains:

[S's belief in P is justified *because* (S's belief in Q₁ is justified *and* S's belief in Q₂ is justified *and* S's belief in P is y-supported by S's beliefs in Q₁ and in Q₂)]

is a typical example of a y-flow. Third, a helpful way of understanding the v- vs. y- distinction (whether for support or for flow) is to think in terms of overdetermination: v-flows are instances of justificatory overdetermination, whereas y-flows are not, and similarly relations of v-support involve overdetermined support, whereas relations of y-support do not. Fourth, the v- vs. y- distinction is perhaps the clearest way of seeing the dangers of assimilating a ramifying regress into a serial regress. It is easy to appreciate the importance of the v- vs. y- distinction, but there is no way to draw that distinction after we collapse a ramifying regress into a serial regress.

4. Three Puzzles for Coherentists

Each of our three distinctions is, on its own, fairly simple. However, when we put these three distinctions together, they have some powerful effects on how we think about justificatory structure. In particular, they give coherentists the resources to address a number of pressing puzzles that arise for their position.

Puzzle #1: It has become standard in the literature to distinguish between *linear* and *holistic* forms of coherentism. As their names suggest, linear coherentists accept a familiar linear conception of justification, whereas holistic coherentists take a more holistic approach to justification's source. But what exactly does this distinction come to? Here is a typical explication of the difference, by Laurence Bonjour:

The only apparent hope for a coherentist response to this objection [that chains of justification cannot loop back on themselves], stemming originally from [Bernard] Bosanquet, is the suggestion that the objection depends on the plausible but ultimately mistaken idea that relations of justification fundamentally involve a *linear*, asymmetrical order of epistemic dependence among the beliefs in question. The contrary suggestion is that justification, when properly understood, is ultimately *nonlinear* and *holistic* in character, with all of the beliefs in the relevant system of beliefs standing in relations of *mutual* support, but none being epistemically prior to the others. (Bonjour 2003, 44–45)

BonJour's general idea is clear. But it is difficult to extract from this passage a crisp definition of linear or holistic coherentism. Hence my first puzzle for coherentists: how precisely do we formulate this oft-made but somewhat-obscure distinction between linear and holistic varieties of coherentism?

As I see it, my distinction #2, between relations of support and flows of justification, is the key to the distinction between linear and holistic coherentism. Linear coherentists acknowledge both relations of support and flows of justification. Holistic coherentists, meanwhile, are fine with relations of support, but they deny the existence of flows of justification. According to holistic coherentists, a belief is never justified *in virtue of* standing in a relation of support with another belief that is justified: justification does not *transfer* or *transmit* from one belief to another via inferential relations. This difference between linear and holistic coherentists manifests in their attitudes toward directed hypergraphs of justificatory structure. For linear coherentists, the directed edges in those graphs can be interpreted *either* to designate relations of support *or* to designate flows of justification. (In my conventions: these graphs can include arrows as well as shadows under those arrows.) Holistic coherentists, on the other hand, will insist that the directed edges in those graphs be understood *only* as relations of support, *not* as flows of justification. (In my conventions: the graphs will include arrows but no shadows under those arrows.)

Thus I am proposing an interpretation on which the holistic coherentist, qua holist, is committed to no more and no less than the following: justification does not flow or transmit between beliefs.¹⁰ This way of understanding the commitments of holistic coherentism allows us to clear up some common confusions about the position, including the following:

- It is sometimes said that holistic coherentists deny a presupposition of the regress argument. This is only half correct. We need to distinguish *the regress of support*, which is the regress we get when items in the regress are linked via relations of support, from *the regress of justification*, which is the regress we get when items in the regress are linked via flows of justification. It is only the latter regress which holistic coherentists deny a presupposition of.
- Similarly, it is sometimes said that holistic coherentism is a form of foundationalism.¹¹ Again, this is only half correct. There is a sense in which holistic coherentism is a foundationalist response to

the regress of justification—“one-step foundationalism,” as Klein (2005, 278) nicely puts it, in which every justified belief is justified in virtue of some property concerning its coherence with other beliefs, thereby obviating the need to continue the regress further than a single step. However, this characterization of holistic coherentists is not entirely fair, since holistic coherentists don’t think that the regress of justification—with its appeal to the transmission of justification as what binds together successive stages of that regress—is a legitimate one. Moreover, when we turn our attention to *the regress of support*, we find that holistic coherentism represents a decidedly non-foundationalist response to that regress.

- Finally, it is sometimes said that holistic coherentism makes the justification of a single belief depend on the coherence of one’s *entire system* of beliefs (and other justificatorily-relevant mental states). But holistic coherentism *as such* is committed to nothing of the sort. All it is committed to is a denial of flows of justification. This leaves it open that the justification of an individual belief might be determined by the coherence of some subsystem of one’s beliefs (and other justificatorily-relevant mental states).

In what follows, I restrict myself to consideration of holistic forms of coherentism, since—like most authors—I consider holistic coherentism more defensible than linear coherentism.¹² Thus in all future diagrams of justificatory structure, I will only draw arrows of support and will not be concerned with representing the flow of justification (since holistic coherentists deny the existence of such flow).

Puzzle #2: Coherence, it is often assumed, is exclusively a relation among beliefs (or propositions accessible to belief, or other doxastic entities of some sort—I will continue to focus on the case of belief). This leads to one of the most commonly cited objections to coherentism, in both its holistic and linear varieties: the so-called *isolation objection* or *input problem*. In brief, the worry is this: if justification depends only on the internal coherence of a system of beliefs, then justification does not require any *contact* or *input* from the world outside that system of beliefs—and, in particular, does not require one to take into account one’s *experiences* when forming and maintaining beliefs. But surely, the thought goes, this is implausible. Part of what it is to have justified beliefs is to respond appropriately to one’s experiences. If coherentism

cannot account for this truism, then so much the worse for coherentism.

The obvious response to this objection is to move to a view on which the bearers of coherence are not sets of beliefs alone, but rather sets of beliefs plus experiences. Such a form of coherentism is known as “non-doxastic coherentism.”¹³ Non-doxastic coherentism has much to recommend it, especially since it offers a nice, tidy response to what is often considered a knockdown objection to coherentism. But non-doxastic coherentism is also a deeply puzzling position, which perhaps explains why it has not been widely embraced by those with coherentist sympathies. In particular, it is not clear how coherentists can include experiences among a coherent set of mental states without lapsing into foundationalism. After all, the items in a coherent set are supposed to be mutually supporting, and experiences (it is standardly assumed) are not the sort of mental state that can be supported by other mental states. So if experiences are to feature in a network of mutually supporting mental items, they must do so by providing support to other items in that network—and, in particular, by providing support to beliefs. However, if we allow that experiences are non-supported supporters of beliefs, haven’t we thereby committed ourselves to one of the central tenets of foundationalism? Thus it appears that coherentists are caught in a dilemma: if they deny experiences a justificatory role, they must face the isolation objection, but if they allow experiences to be justificatorily relevant, they risk becoming foundationalists. This is my second puzzle.

In order to see our way out of this puzzle, let us look more closely at the dilemma’s second horn. As we have already mentioned, one form of foundationalism—perhaps its most commonly defended form, nowadays—holds that there is a non-inferential variety of support that exists between an experience and a belief, in virtue of which basic beliefs are justified. Let us, somewhat anachronistically, refer to this kind of support as *the categorical given*.¹⁴ We have been pictorially representing it like so:



Figure 8

Note that an experience which provides support in this manner acts as a *regress stopper*: it offers support but is not itself in need of support, so our regress comes to an end. Moreover, I have been individuating foundationalism, infinitism, and coherentism—as is customary—in terms of their distinctive responses to

the regress problem. So if this is the only way in which an experience can support a belief, we are skewed by the second horn of our dilemma: taking experiences to be the sort of thing that can feature in a network of mutually supporting (and hence coherent) mental states commits us to foundationalism.

However, by combining my distinctions #1 and #3, we can see that there is another way for experiences to play a supporting role for beliefs. According to distinction #1, we should include both doxastic and experiential nodes in our diagrams of the structure of support for a belief. According to distinction #3, it can be the case that several nodes in such a diagram jointly support some other node. Put these ideas together, and we are left with the possibility that one (or more) experience nodes and one (or more) belief nodes might *together* support a given belief node. In graphic terms:

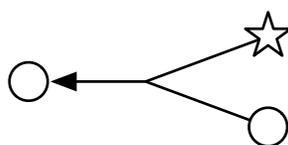


Figure 9

Adapting some terminology of Anil Gupta's, let us call this kind of support *the hypothetical given*.¹⁵ Candidates for (the content of) the belief on the right in Figure 9 could include any of the following: <Lighting conditions are normal>, <I haven't ingested any hallucinogenic drugs recently>, <My visual faculties are generally reliable>, etc. But regardless of which belief—or beliefs—together with an experience—or experiences—act as a y-supporter for another belief, the crucial point is that a support structure of this kind does not act as a regress stopper. Presumably y-support of the sort depicted in Figure 9 can only be part of what makes it the case that a belief is justified if the belief node on the right is itself supported by other nodes in our system. Thus coherentists who countenance the hypothetical given can allow experiences to play an epistemic role in the regress of support (and in their account of coherence) without taking experiences to be regress stoppers: experiences would provide epistemic friction in the system, but they would not generate support or justification *ex nihilo*.¹⁶ This solves our second puzzle.¹⁷

Puzzle #3: My final puzzle starts with the following tempting line of thought: “Isn't a more accurate name for coherentism ‘circularism’? And isn't circularism incredibly implausible? Why would

anyone hold that a belief can be justified in virtue of its having a circular chain of support?" In a way I have already raised this problem for coherentism in this essay's opening pages. There I said that one important reason to formulate the regress problem in terms of the ramifying rather than the serial regress is that, within the confines of the serial regress, worries about circularity constitute a fatal objection to coherentism. However, it might seem puzzling why moving to the ramifying regress is any help here. After all, the coherentist response to the ramifying regress still involves an appeal to circular lines of support, albeit many such circular lines of support—one additional one for each ramification in our structure. But if a single circle is epistemically pernicious, shouldn't multiple circles be just as pernicious, if not more?

It is here that distinction #3 (between v- and y-support) becomes crucial. It is not ramifications as such that help answer the anti-circularism challenge, but rather ramifications involving y-support. I agree with the hypothetical objector that *circles of full support* are not justification-making. However, I think that *circles of partial support* can be relevant to whether a belief is justified. Compare the following two (vastly oversimplified) toy models of what it might take for a set of beliefs to be coherent. First, let a *circular triad* be a set of three beliefs $\{b', b'', b'''\}$ such that b' on its own supports b'' , b'' on its own supports b''' , and b''' on its own supports b' . Pictorially:

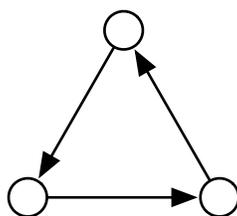


Figure 10

Second, let an *entwined triad* be a set of three beliefs $\{b_1, b_2, b_3\}$ such that every two of them y-supports the third: b_1 and b_2 together support b_3 ; b_1 and b_3 together support b_2 ; and b_2 and b_3 together support b_1 .

Pictorially:

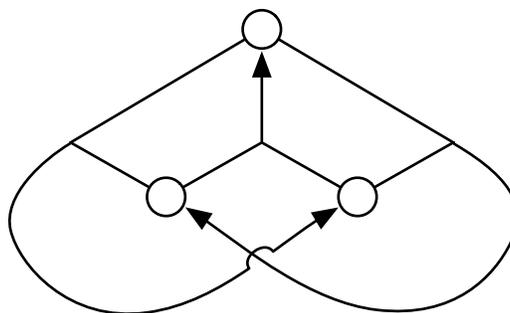


Figure 11

Our objector is assuming, I believe, that a circular triad is the coherentist’s model for the sort of support structure that confers justification on a belief. But once we appreciate that coherentists usually take an entwined triad to be a better—albeit still imperfect—model for a justification-conferring support structure, the force of the objector’s challenge dissipates.¹⁸

It is important that we be clear here about the exact nature of our objector’s challenge. On one way of pressing that challenge, it involves a brute appeal to intuition. The objector has an intuition that circular support structures can’t underwrite a belief’s status as justified, and she takes this intuition to refute coherentism. She looks at Figure 10 and thinks, “Yuck! That can’t be what makes a belief justified.” I share the objector’s “yuck”-reaction to Figure 10. But when I consider entwined triads such as Figure 11, I am less certain. It is far from clear to me that Figure 11 is hopeless as a model for the sort of support structure that can serve to make a belief justified, in the way in which Figure 10 seem hopeless. In short, the intuitive version of the anti-circularism objection is much less compelling when we consider structures in which the only circles of support are circles of *partial* support.

Moreover, there is a way in which the intuitive version of our objection is disappointing. Since (holistic) coherentism is defined in terms of its distinctive response to the regress of support, (holistic) coherentism is, by definition, a view according to which a circular support structure is required in order for a belief to be justified. So on the intuitive version of the anti-circularism challenge, our objector is doing little more than appealing to an intuition that (holistic) coherentism is false. That, in the end, is not the most insightful objection. If circular support structures are not justification-affording, we would like to know *why* they are not—to know what it is about such structures that makes them so problematic.

Thus the anti-circularism objection is more powerful when it does not rest content with a bare intuition that circular support structures can't ground justification, and instead fills in the worry in a way that gives it more teeth. Here is one way of doing so. Let us assume that inferential support among propositions obeys contraposition, so that if P supports Q (whether deductively, inductively, or abductively), then $\neg Q$ supports $\neg P$ (whether deductively, inductively, or abductively). It follows that if some subject's beliefs in P , Q , and R together constitute a circular triad, then there is another possible subject whose beliefs in $\neg P$, $\neg Q$, and $\neg R$ together constitute a circular triad in the reverse direction: the first subject's belief in P supports her belief in Q , which supports her belief in R , which in turn supports her belief in P , whereas the second subject's belief in $\neg Q$ supports her belief in $\neg P$, which supports her belief in $\neg R$, which in turn supports her belief in $\neg Q$ (see Figure 12).

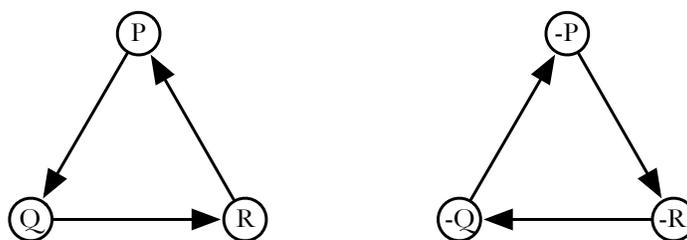


Figure 12

So if circular triads are justification-affording, it follows that if the support relations among the first subject's beliefs in P , Q , and R make each belief justified, the support relations among the second subject's beliefs in $\neg P$, $\neg Q$, and $\neg R$ make each of *these* beliefs justified. But, it might be insisted, this is an absurd result: the sorts of considerations that make belief in a set of propositions justified must be able to distinguish between those propositions and their negations. Hence, our objector concludes, circular triads cannot be justification-affording.¹⁹

This more specific version of the anti-circularism objection is, I believe, a forceful challenge to the view that *circular triads* can underwrite justification. But it is important to realize that this version of the objection is completely powerless against the view that *entwined triads* can underwrite justification. Suppose a subject's beliefs in P , Q , and R together make up an entwined triad. How do we construct a negated analogue of this entwined triad that is itself an entwined triad? From the fact that P and Q together

support R , it doesn't follow that $\neg R$ by itself supports $\neg P$ and $\neg Q$ together—indeed, it's not even clear what this claim means.²⁰ So, when we try to construct an analogue of Figure 12 for entwined triads, we are left with the following:

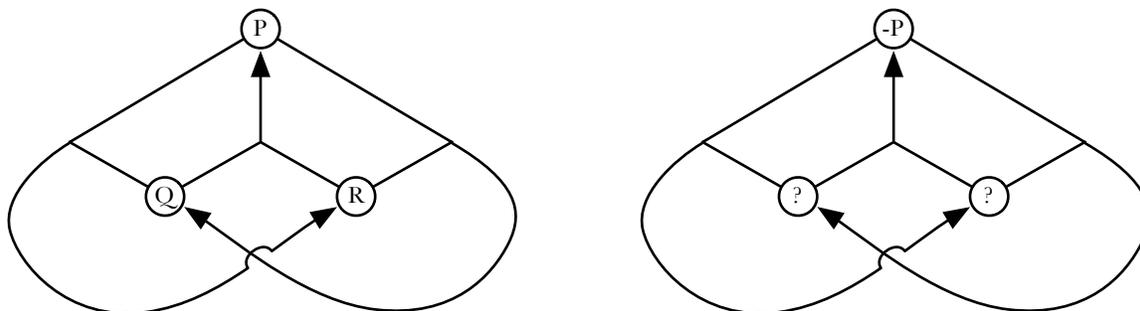


Figure 13

Thus I conclude that both the intuitive and the more specific version of the anti-circularism charge lose much of their force when applied to proposals on which it is circles of partial support, rather than circles of full support, which undergird coherentist justification. Moreover, matters look even better for the coherentist when we allow experiences into our support structures via an appeal to the hypothetical given. For example, consider what I will call an *entwined sextet*: three beliefs and three experiences $\{b_1, b_2, b_3, e_1, e_2, e_3\}$ such that $\{b_1, b_2, e_3\}$ supports b_3 , $\{b_1, b_3, e_2\}$ supports b_2 , and $\{b_2, b_3, e_1\}$ supports b_1 . As a graph:

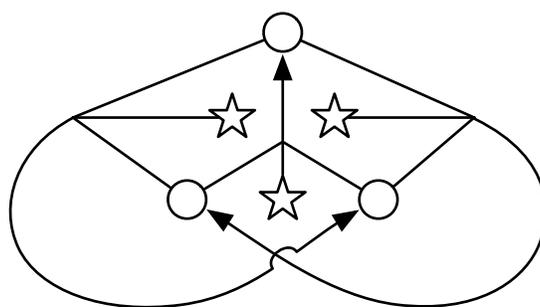


Figure 14

I doubt that there actually exist three beliefs and three experiences which bear these particular support relations to one another. But if such a sextet *were* to exist, I don't think the presence of circles of partial support in the overall system would give us any reason at all to deny that such a structure confers

justification on the beliefs contained therein. So when we take an entwined sextet to be our model for the sort of structure that gives rise to coherentist justification, the intuitive version of the anti-circularism objection completely lapses. The more specific version of that objection which we have been considering also lapses: it is even more difficult to construct a negated analogue of a given entwined sextet than it is to construct a negated analogue of a given entwined triad. The lesson, once again, is that worries about circularity only seem crippling to coherentists when we neglect the possibility of γ -support.

5. A New Way of Characterizing Coherence

The distinction between circular triads and entwined triads/sextets is not just a way of making “circularist” responses to the regress of support look less problematic; it also points us toward something more important: a way of providing a substantive characterization of what coherence comes to.

The basic idea here is simple. The differences between a circular triad, an entwined triad, and an entwined sextet are the kinds of differences we can express in graph-theoretic terms—that is, in terms of the properties and relations encoded by the sorts of directed hypergraphs we have been considering. So if an entwined triad or an entwined sextet is closer to the type of structure that constitutes a coherent set of mental states than a circular triad is, we have good evidence that coherence itself is expressible in graph-theoretic terms.

More boldly, we may argue as follows. Coherence, we have been assuming, is a matter of the structure of support among a subject’s beliefs, experiences, and other justificatorily-relevant mental states at a given time. But we can use directed hypergraphs to represent all of those support relations. That is, we can use directed hypergraphs to represent all of the relations that have a bearing on coherence. Moreover, ways of combining graph-theoretic notions are also graph-theoretic notions. So it follows that coherence is itself expressible as a graph-theoretic property of our directed hypergraphs.

Which graph-theoretic property, though? That, of course, is the tricky part. Here is an example of the sort of proposal I have in mind. Suppose we are working with the hypothesis that coherence is a graph-theoretic property possessed by an entwined triad and an entwined sextet, but not possessed by a circular triad. Then there is a natural way of extending this hypothesis into an account of coherence and

of coherentist justification. Let a *thread of support* from node x to node y be a *path of partial support* that starts at x and ends at y . Let a *web of support* from (non-empty) set of nodes N to node y be a *path of full support* that starts at all of the members of N and ends at y .²¹ Now, because deduction is monotonic, extra webs in a given structure are easy to come by: if set of nodes N deductively supports node y , and if N^* is an arbitrary set of nodes, it follows that from the monotonicity of deduction that $N \cup N^*$ deductively support y , and hence that there is a web from $N \cup N^*$ to y . So let us say that a web from N to y is *superseded* when (and only when) there exists a web from a proper subset of N to y . Then the proposal I have in mind makes the following two central claims:

- (C₁) Set of nodes N is *coherent* if and only if, and because, for every two belief nodes b_1 and b_2 in N , there is a non-superseded web of support from b_1 and at least one other member of N to b_2 .
- (C₂) Belief b is *justified* if and only if, and because, there exists a set of nodes, N , such that (i) N is coherent, and (ii) b is a member of N .

(C₁) tells us what it takes for a set of beliefs, experiences, and other justificatorily-relevant mental states to be coherent. (C₂) tells us what it takes for an individual belief to be justified. Together, they constitute a graph-theoretic form of coherentism about epistemic justification.

I don't mean to be endorsing this specific theory; in fact, I think there are fatal problems with it, which I shall get to momentarily. Rather, I bring up (C₁) and (C₂) in order to illustrate a way of characterizing coherence—and of formulating coherentism—that is very different from the approaches one finds in the contemporary literature. That literature is dominated by two general approaches. According to *the probabilistic approach*, coherence is taken to be a property of sets of propositions that bear certain probabilistic relations to one another—the debate being over which specific probabilistic relations are the right ones. According to *the weighted-sum approach*, we first provide a list of various *coherence-increasing factors* (such as the degree to which the beliefs in a set are connected to one another) and various *coherence-decreasing factors* (such as the fact that the beliefs in a set are inconsistent with one another). Coherence is then taken to be determined by a weighted sum of the coherence-increasing factors minus a weighted sum of the coherence-decreasing factors.

Much important work has been done in both of these traditions. But I also think each tradition

has its shortcomings. One major problem with the probabilistic approach is that it cannot allow experiences to play a role in determining coherence. Suppose $\{P_1, \dots, P_n\}$ is a set of n propositions which bear the right probabilistic relations to one another to make for coherence, according to a given version of the probabilistic approach. Then if I believe each of $\{P_1, \dots, P_n\}$, the corresponding set of n beliefs counts as coherent, regardless of their relation to my experiences. For this reason advocates of the probabilistic approach are forced to endorse doxastic coherentism, and hence must face the isolation objection.²² The weighted-sum approach, on the other hand, tends to be developed in a way that is unsatisfyingly vague: usually almost nothing is said either about how the individual coherence-increasing and -decreasing factors are precisely measured or about how they are weighed against each other in the determination of overall coherence. Similarly, we are rarely given a reason to think the proposed list of coherence-increasing and -decreasing factors is exhaustive, or to think that all of the items in the list are independent of one another. Although these considerations do not constitute a decisive objection to the weighted-sum tradition (perhaps it is a mistake to look for too much precision in a theory of coherence), they give us good reason to search for a more systematic way of conducting our inquiries into the nature of coherence.²³

My graph-theoretic approach is, I believe, precisely such a more systematic approach. By using the tools of graph theory, we can theorize about the factors that make for coherence and incoherence in a rigorous manner. The approach is also more versatile than the probabilistic approach, since it easily allows for the bearers of coherence to include experiences and other non-doxastic states. But this is not to say that the graph-theoretic approach is entirely independent of the two mainstream approaches to coherence. Many of the coherence-increasing and -decreasing factors identified by advocates of the weighted-sum approach will be ones which the proponent of the graph-theoretic approach will want to account for in graph-theoretic terms, and which she can incorporate into her graph-theoretic analysis of coherence. Similarly, many of the proposals that have been made in the probabilistic literature on coherence will have graph-theoretic analogues, so the lessons learned in that literature can have bearing on how we think about graph coherence.

Let us return to our example of a graph-theoretic account of coherence and justification, namely (C₁) and (C₂). (C₁) is a natural candidate for the graph-theoretic property that distinguishes structures like an entwined triad or sextet from structures like a circular triad: we want all of the beliefs in the structure to support each other by way of threads that are part of a larger non-superseded web, not by way of threads that are themselves webs. So if we are searching for a graph-theoretic property that is possessed by entwined triads and sextets and not possessed by circular triads, (C₁) is a good start. Moreover, (C₂) is a natural way of moving from an account of coherence to a full theory of justified belief.²⁴

Nevertheless, (C₁) and (C₂) are inadequate as they stand. There are clear counterexamples to the criterion for justified belief provided by (C₁) and (C₂) in both directions: failures both of necessity and of sufficiency. Let us consider problems of necessity first. It is very plausible to think that justified belief is closed under inferential support, in the following sense:

single-premise closure: If S believes P, S believes P⁺, S's belief in P is justified, and S's belief in P deductively (or inductively, or abductively) supports S's belief in P⁺, then S's belief in P⁺ is justified.²⁵

Single-premise closure is difficult to square with (C₁) and (C₂). Suppose the principle's antecedent is satisfied. In order to satisfy its consequent, we need there to exist a set of nodes N containing S's belief in P⁺ which is coherent in the way specified by (C₁). And in order for N to be coherent in the way specified by (C₁), it must at least be the case that there is a thread of support from S's belief in P⁺ to some other belief of S's. But the conditions in the antecedent of our principle give us no assurance that S's belief in P⁺ partially supports any other belief which S holds: P⁺ might be a *very* abstruse consequence of P, and her belief in P⁺ might be the only belief she currently holds that is at all relevant to this abstruse issue. (For example, let P⁺ be the disjunction of P and some utterly inane proposition about some completely trivial matter.)²⁶

Luckily, there is an obvious way to patch up our proposal so that it satisfies single-premise closure.

All we need to do is replace (C₂) with

(C₂^{*}) Belief *b* is *justified* if and only if, and because, there exists a set of nodes N such that (i) N is coherent, and (ii) there is a non-superseded web of support from some subset of N to *b*.

It is easily shown that (C₁) plus (C₂^{*}) together entail single-premise closure. (Just take the non-superseded

web of support which makes S's belief in P justified and append to it the link between S's belief in P and S's belief in P⁺ to create a non-superseded web of support that makes S's belief in P⁺ justified.) And when we adjust our account of justified belief in this way, the threat to its necessity is removed.^{27, 28}

Threats to the account's sufficiency, however, are not so easily removed. There are at least two problems here. First, a major limitation of (C₁) and (C₂^{*}) is that although they *allow* experiences to play a role in justifying beliefs, they do not *require* one to take heed of one's experiences when forming beliefs. This is a consequence of the fact that, from (C₁)'s perspective, an entwined triad is just as good a candidate for a coherent set of mental states as an entwined sextet is: in both structures, there is a path of partial support from each belief in the structure to every other belief in the structure, and whether any of those paths of partial support requires the assistance of an experience node in order to constitute a path of full support is irrelevant, so far as (C₁) is concerned. Thus one could count as having a justified belief in virtue of holding an entwined triad of beliefs even if those beliefs are wildly at odds with one's experiences. For this reason, we need a stronger connection between justification and experience in order to avoid the isolation objection.

The second problem with the combination of (C₁) and (C₂^{*}) is more straightforward: there are clear counterexamples to the proposal's sufficiency. Consider the following trio of propositions: {P & Q, Q & R, R & P}, for arbitrary P, Q, and R. Any two of these propositions deductively y-supports the third. So it follows that a belief in P & Q, a belief in Q & R, and a belief in R & P are an entwined triad. Thus if we pick our P, Q, and R such that no one of them supports any of the others on its own, we are saddled with the result that all it takes, according to (C₁) and (C₂^{*}), to have a justified belief in one of these three conjunctive propositions is to believe the other two. But that's absurd. I can't justifiably believe <I will live forever, and there are infinitely many primes> just by also believing <There are infinitely many primes, and my son's favorite color is red> and <My son's favorite color is red, and I will live forever>.

The lesson from both of these problems is that our initial hypothesis that coherence is a property possessed by entwined triads and entwined sextets but not by circular triads was mistaken. In particular, it is not true that every instance of an entwined triad constitutes a coherent set of mental states: more

structure is needed. Now there are a number of ways in which we might revise (C₁) and (C₂^{*}) in order to provide that extra structure. In my opinion, the best way of doing so is to relax an assumption we made at the outset of this essay, and to consider negative structural relations in our hypergraphs in addition to positive relations of support. However, that is not a claim I will defend here. After all, my purpose in this section has been to *illustrate* a certain way of thinking about coherence and coherentism that is very different from the main approaches one encounters in the literature. (C₁) and (C₂^{*}) serve that purpose, even if there are serious objections to the claim that they constitute a fully adequate coherentist theory of justification. I think it should be clear that the two objections we have been considering are only objections to this specific proposal, not objections to my entire graph-theoretic approach.

6. Is Graph-Theoretic Coherence Truth-Conducive?

Still, there are other objections that might be taken to refute the entire graph-theoretic approach. I want to end by considering one objection of that sort.

A large portion of the literature on probabilistic approaches to coherence is devoted to the following question: is coherence, when construed probabilistically, truth-conducive? The reason this question has loomed large is twofold. First, it is commonly assumed that a “no”-answer to the question is a decisive objection to coherentism: epistemic justification must be truth-conducive, the thought goes, so if coherence is not truth-conducive, coherentist accounts of epistemic justification are false. Second, the question “Is probabilistically-construed coherence truth-conducive?” has loomed large because a number of authors have argued either—more modestly—that many specific accounts of coherence in probabilistic terms yield a notion of coherence that is not truth-conducive, or—more boldly—that *any* account of coherence in probabilistic terms must have this unfortunate result.²⁹

We can imagine a similar objection being levied against graph-theoretic understandings of coherence. Graphical coherence, it might be claimed, is not truth-conducive: that sort of a structural relation among one’s mental states does not in general promote the acquisition of true beliefs and the avoidance of false ones. And, the objection continues, this fact gives us good reason to reject the idea that epistemic justification is grounded in graphical coherence.

The problem with this objection, I believe, is with its background assumption that epistemic justification must be truth-conducive. In other work, I have argued that although the slogan “Truth-conduciveness is the mark of epistemic justification” is widely endorsed by contemporary epistemologists, that slogan represents a distinctively *consequentialist* way of understanding epistemic normativity which should be rejected.³⁰ In my view, it is simply not true that epistemic justification is a matter of what best promotes, either as an objective matter of fact or from the subject’s own perspective, the overall balance of epistemically good outcomes (such as the obtaining of true beliefs) minus epistemically bad outcomes (such as the obtaining of false beliefs). So in response to the question, “Is graph-theoretic coherence truth-conducive?” I reply, “Who cares? Truth-conduciveness is the wrong thing to be looking for when theorizing about epistemic justification.”

The two most prominent coherentists of the previous generation, Laurence Bonjour and Keith Lehrer, were also truth-conducivists, albeit truth-conducivists of a subjective consequentialist sort, and their commitment to truth-conducivism profoundly influenced the way they developed and defended coherentism.³¹ Because of this, Bonjour and Lehrer were not, at the most fundamental level, coherentists. For both of them, the most fundamental epistemic property—that which ultimately grounds all other epistemic properties—was truth-conduciveness-from-the-subject’s-perspective; having coherent beliefs was, for them, merely a way—perhaps the only way—of acquiring beliefs with that desired property. This is not true coherentism. It is like a utilitarian who says that the Principle of Utility describes our most fundamental moral duty, but the best way to conform to that duty is to internalize Kant’s Categorical Imperative. That is not true Kantianism.

Thus in order to truly defend coherentism, we must renounce the idea that there is a more fundamental epistemic standard—conducive toward the true and away from the false—to which the coherentist epistemic standard must answer.³² In this way, I view the current essay as an exploration of non-consequentialist epistemology. It is not the only way of developing a non-consequentialist epistemic theory—there also exist non-consequentialist versions of foundationalism and infinitism, as well as various hybrid views. But it is an approach I consider particularly fruitful.

Moreover, let me remind the reader of a point with which I began this essay. Although I have been emphasizing the ways in which the graph-theoretic resources I have been developing here are of use to coherentists (whether of a consequentialist or non-consequentialist sort), those resources should also be of interest to infinitists and foundationalists (again, whether of a consequentialist or non-consequentialist sort). For nearly every objection or puzzle that arises for coherentism, there is a parallel objection or puzzle that arises for infinitism: infinitists also make a vague but ubiquitous distinction between linear and holistic forms of their view, infinitists also must explain how they can reply to the isolation objection without secretly becoming foundationalists, etc. So my graph-theoretic strategies for replying to these objections and puzzles can be employed by infinitists as well. And, in the case of foundationalism, although much of what I say here is not of concern to foundationalists when developing their account of immediate justification, foundationalists may want to appeal to some of my graph-theoretic machinery when providing an account of how mediately justified beliefs depend on immediately justified ones (since many traditional foundationalists appeal to a notion of coherence at this stage in their story). Any plausible epistemic theory must recognize that there can be complicated and subtle structural relations among our beliefs that have a bearing on their epistemic standing; so any plausible epistemic theory should aspire to theorize about the nature of epistemic structure in as rigorous a manner as possible.³³

References:

- Aikin, Scott F. 2011. *Epistemology and the Regress Problem*. New York: Routledge.
- Alston, William P. 1976. "Two Types of Foundationalism." *Journal of Philosophy* 73: 165–85. Reprinted in Alston 1989, 19–38. Page references are to the 1989 reprint.
- . 1989. *Epistemic Justification: Essays in the Theory of Knowledge*. Ithaca, NY: Cornell University Press.
- Bernecker, Sven, and Duncan Pritchard, eds. 2011. *The Routledge Companion to Epistemology*. New York: Routledge.
- Berker, Selim. 2011. "Gupta's Gambit." *Philosophical Studies* 152: 17–39.
- . 2013a. "Epistemic Teleology and the Separateness of Propositions." *Philosophical Review* 122: 337–93.
- . 2013b. "The Rejection of Epistemic Consequentialism." *Philosophical Issues* 23: 363–87.
- . 2015. "Reply to Goldman: Cutting Up the One to Save the Five in Epistemology." *Episteme*.
- . MS-a. "A Graph-Theoretic Account of Epistemic Structure." In progress.

- . MS-b. “Three Approaches to Coherence: Probabilistic, Pluralistic, and Graph-Theoretic.” In progress.
- . MS-c. “The Unity of Grounding.” In progress.
- BonJour, Laurence. 1985. *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.
- Cling, Andrew. 2008. “The Epistemic Regress Problem.” *Philosophical Studies* 140: 401–21.
- . 2009. “Reasons, Regresses, and Tragedy: The Epistemic Regress Problem and the Problem of the Criterion.” *American Philosophical Quarterly* 46: 333–46.
- . 2014. “The Epistemic Regress Problem, the Problem of the Criterion, and the Value of Reasons.” *Metaphilosophy* 45: 161–71.
- DePaul, Michael R. 2011. “Foundationalism.” In Bernecker and Pritchard 2011, 235–44.
- Elgin, Catherine Z. 1996. *Considered Judgment*. Princeton, NJ: Princeton University Press.
- Feldman, Richard. 2003. *Epistemology*. Upper Saddle River, NJ: Prentice Hall.
- Fine, Kit. 2012. “Guide to Ground.” In *Metaphysical Grounding: Understanding the Structure of Reality*, edited by Fabrice Correia and Benjamin Schnieder, 37–80. Cambridge: Cambridge University Press.
- . 2012b. “The Pure Logic of Ground.” *Review of Symbolic Logic* 5: 1–25.
- Goldman, Alvin. 1992. “Epistemic Folkways and Scientific Epistemology.” In *Liaisons: Philosophy Meets the Cognitive and Social Sciences*, 155–75. Cambridge, MA: MIT Press.
- Gupta, Anil. 2006. *Empiricism and Experience*. Oxford: Oxford University Press.
- Klein, Peter D. 2005. “Infinitism Is the Solution to the Regress Problem.” In Steup and Sosa 2005, 131–40. Reprinted in Steup, Turri, and Sosa 2014, 274–83. Page references are to the 2014 reprint.
- . 2011. “Infinitism.” In Bernecker and Pritchard 2011, 245–56.
- Klein, Peter D., and Ted A. Warfield. 1994. “What Price Coherence?” *Analysis* 54: 129–32.
- Kvanvig, Jonathan L., and Wayne D. Riggs. 1992. “Can a Coherence Theory Appeal to Appearance States?” *Philosophical Studies* 67: 197–217.
- Lehrer, Keith. 1974. *Knowledge*. Oxford: Oxford University Press.
- Lycan, William G. 2012. “Explanationist Rebuttals (Coherentism Defended Again).” *Southern Journal of Philosophy* 50: 5–20.
- Lyons, Jack C. 2009. *Perception and Basic Beliefs: Zombies, Modules, and the Problem of the External World*. Oxford: Oxford University Press.
- Mackie, J. L. 1974. *The Cement of the Universe: A Study of Causation*. Oxford: Oxford University Press.
- Oakley, I. T. 1976. “An Argument for Scepticism concerning Justified Beliefs.” *American Philosophical Quarterly* 13: 221–28.
- Olsson, Erik J. 2005. *Against Coherence: Truth, Probability, and Justification*. Oxford: Oxford University Press.
- . 2008. “The Place of Coherence in Epistemology.” In *New Waves in Epistemology*, edited by Vincent F. Hendricks and Duncan Pritchard, 192–214. New York: Palgrave Macmillan.
- Plantinga, Alvin. 1993. *Warrant: The Current Debate*. Oxford: Oxford University Press.
- Pollock, John L. 2001. “Nondoxastic Foundationalism.” In *Resurrecting Old-Fashioned Foundationalism*, edited by Michael R. DePaul, 41–57. Lanham, MD: Rowman and Littlefield.
- Poston, Ted. 2014. *Reason and Explanation: A Defense of Explanatory Coherentism*. New York: Palgrave Macmillan.

- Pryor, James. 2005. "There Is Immediate Justification." In Steup and Sosa 2005, 181–202. Reprinted in Steup, Turri, and Sosa 2014, 202–22. Page references are to the 2014 reprint.
- . 2012. "When Warrant Transmits." In *Mind, Meaning, and Knowledge: Themes from the Philosophy of Crispin Wright*, edited by Annalisa Coliva, 269–303. Oxford: Oxford University Press.
- Roche, William A. 2012. "A Reply to Cling's 'The Epistemic Regress Problem.'" *Philosophical Studies* 159: 263–76.
- . 2013. "Coherence and Probability: A Probabilistic Account of Coherence." In *Coherence: Insights from Philosophy, Jurisprudence, and Artificial Intelligence*, edited by Michał Araszewicz and Jaromír Šavelka, 59–91. Dordrecht: Springer.
- Rosen, Gideon. 2010. "Metaphysical Dependence: Grounding and Reduction." In *Modality: Metaphysics, Logic, and Epistemology*, edited by Bob Hale and Aviv Hoffmann, 109–35. Oxford: Oxford University Press.
- Sayre-McCord, Geoffrey. 1996. "Coherentist Epistemology and Moral Theory." In *Moral Knowledge?: New Readings in Moral Epistemology*, edited by Walter Sinnott-Armstrong and Mark Timmons, 137–89. Oxford: Oxford University Press.
- Schaffer, Jonathan. 2009. "On What Grounds What." In *Metametaphysics*, edited by David Chalmers, David Manley, and Ryan Wasserman, 347–83. Oxford: Oxford University Press.
- Schechter, Joshua. 2013. "Rational Self-Doubt and the Failure of Closure." *Philosophical Studies* 163: 429–52.
- Sosa, Ernest. 1980a. "The Foundations of Foundationalism." *Noûs* 14: 547–65. Reprinted in Sosa 1991. Page references are to the 1991 reprint.
- . 1980b. "The Raft and the Pyramid: Coherence versus Foundationalism in the Theory of Knowledge." *Midwest Studies in Philosophy* 5: 3–25. Reprinted in Sosa 1991. Page references are to the 1991 reprint.
- . 1991. *Knowledge in Perspective: Selected Essays in Epistemology*. Cambridge: Cambridge University Press.
- Steup, Matthias, and Ernest Sosa, eds. 2005. *Contemporary Debates in Epistemology*. 1st edition. Malden, MA: Blackwell.
- Steup, Matthias, John Turri, and Ernest Sosa, eds. 2014. *Contemporary Debates in Epistemology*. 2nd edition. Malden, MA: Blackwell.

¹ See, for example, Alston 1976, 30; Klein 2005, 276; and Cling 2008, 403n3.

² As, for instance, in the massive literature spawned by Klein and Warfield 1994. See Olsson 2008 for overview.

³ Prominent examples of this approach include BonJour 1985 and Sayre-McCord 1996.

⁴ See Berker MS-a.

⁵ We might need to add to (i) and (ii) a condition according to which S believes P *on the basis of* her belief in Q_1, \dots, Q_n and their support of P. I will be setting aside complications of this sort relating to the basing relation.

⁶ Although my terminology here is, I believe, widespread, it is not universal. For instance, in a recent series of articles Andrew Cling (2008, 2009, 2014) uses the word "support" differently from how I do: for Cling, the expressions "supports" and "is a reason for" are synonymous. That is *not* how I will be using the word "support."

All "support"-talk in epistemology is metaphorical, and there is no single correct way to interpret that metaphor. Nevertheless, here is one reason to prefer my way of using the word "support" over Cling's: given Cling's terminology, coherentists cannot accept the common claim just mentioned in the text, namely that coherence is a matter of mutual support among one's beliefs. If mutual support is what makes it the case that we have coherence,

and if coherence is what makes it the case that we have reasons, and if the reason relation just is the support relation, then we are caught in an explanatory circle. (I assume here that circles of ground are vicious.)

⁷ One frequently encounters the phrase “unjustified justifiers” in the literature, but it is more accurate to refer to these beliefs as “non-justified justifiers,” since they are neither justified nor unjustified (they stand outside the realm of justification). I take the term “non-justified” from Goldman 1992, 159.

⁸ For certain purposes it might be useful to draw different arrows for different types of support (for example, a solid arrow for deductive support, a dotted arrow for inductive support, and so on). I will be ignoring such complexities in this essay.

⁹ And maybe because various other conditions are satisfied as well, if more is needed for the transmittance of justification.

¹⁰ Or, at least, this is what a *pure* holistic coherentist is committed to. There is also the possibility of a *mixed* holistic/linear coherentist view, according to which some beliefs are justified on holistic coherentist grounds, without any appeal to the flow of justification, but those beliefs are able to transfer justification along inferential lines to other beliefs that are not justified on holistic coherentist grounds. On the possibility of such a view, see Plantinga 1993, 79, and Elgin 1996, 17n5. In the text, whenever I speak of “holistic coherentism,” I mean “pure holistic coherentism.”

¹¹ A claim most famously due to Ernest Sosa (1980a, 154; 1980b, 180), but also echoed by Alvin Plantinga (1993, 78–80), Peter Klein (2005, 278; 2011, 250–51), Scott Aikin (2011, 2, 110), and Michael DePaul (2011, 237), among others.

¹² Why? Because linear coherentists are committed to justified belief requiring a circle of partial ground (a series of facts such that each partially grounds the next, together with the last partially grounding the first). Under the assumption that some of our beliefs are justified, this commitment goes against orthodoxy in metaphysics, where most theorists (Schaffer 2009; Rosen 2010; Fine 2012a, 2012b) hold that the grounding relation is both transitive and irreflexive, two properties which together entail that there cannot exist circles of partial ground. Even though I am open to the possibility that this orthodoxy among metaphysicians might be mistaken, I don’t wish to rest anything in the current essay on that possibility. For my own views on the nature of grounding, see Berker MS-c.

¹³ More generally, non-doxastic coherentism refers to any coherentist theory on which the bearers of coherence are sets that include both doxastic and non-doxastic items. See Pollock 2001, 41–42; Lyons 2009, 21; Roche 2012, 274; and Poston 2014, 10. James Pryor (2005, 207) calls this view “impure coherentism,” but I dislike that name, since it suggests that it is somehow more in the nature of coherentism to restrict coherence to a relation among only beliefs. For some forceful arguments against this suggestion, see Kvanvig and Riggs 1992.

¹⁴ I say “somewhat anachronistically” because the traditional doctrine of the given has more built into it than simply a commitment to experiences being non-supported supporters of beliefs.

¹⁵ See Gupta 2006, §4A. My use of the term “hypothetical given” deviates from Gupta’s in two ways. First, Gupta is not always clear on whether the hypothetical given is defined in terms of y-support or defined in terms of the y-flow of justification: much of his language suggests the latter (see Berker 2011, 21n8), but the analogy he draws between the hypothetical given and logical entailment suggests the former (see Gupta 2006, 80–82). (In Berker 2011, I interpreted Gupta as holding that the hypothetical given is a type of y-flow, because that is what best fits the textual evidence, but I now think Gupta is better off taking the hypothetical given to be a type of y-support and revising his language accordingly.) I, on the other hand, use the term “the hypothetical given” exclusively to pick out a certain sort of y-support relation. The second difference between my terminology and Gupta’s is that, for Gupta, it is not a *belief* (or *set of beliefs*) which together with an experience justifies/supports a given belief, but rather what he calls a *view*: a set of concepts, conceptions, and beliefs. I have not included concepts and conceptions among the items which together with an experience provide support because I do not see how concepts and conceptions on their own, when not deployed in specific beliefs, can help make it the case that a support relation obtains.

Another terminology note: my distinction between “categorical” and “hypothetical” forms of the given (and hence of perceptual support) is not the same as Pryor’s distinction between “categorical” and “hypothetical” forms of inferential support in his 2012, 285. Pryor’s distinction there is closer to my earlier distinction between flows of justification and relations of support.

¹⁶ A similar story can be told about other non-doxastic nodes. For example, if intuitions are non-doxastic states that, together with certain beliefs, are able to support other beliefs, then coherentists can take coherence to be a property of sets of beliefs, experiences, and intuitions without lapsing into foundationalism. (In fact, I myself am

most interested in applying the form of coherentism developed here to the epistemology of intuitions about normative matters, but I will continue to focus on experiences as my non-doxastic node of choice in the text.)

¹⁷ My proposal here is not entirely new, but I think presenting it in terms of support structures like that found in Figure 9 is an improvement over existing proposals in the vicinity. For example, in an important article, Jonathan Kvanvig and Wayne Riggs (1992) argue that coherentists can appeal to “appearance states” (i.e. non-factive experiences) if they take a given appearance state to be an INUS condition for justified belief, in J. L. Mackie’s (1974) sense: an *Insufficient* but *Non-redundant* condition for [S is justified in believing P] which is part of a larger condition that is itself *Unnecessary* but *Sufficient* for the obtaining of that fact. There are two problems here. First, by appealing to Mackie’s notion of an INUS condition, Kvanvig and Riggs open themselves to all of counterintuitive consequences that come with that notion (and which have led to the general abandonment of Mackie’s INUS analysis of causation). Second, what Kvanvig and Riggs are suggesting is in effect the analogue of Figure 9 where instead of y-support we have a y-flow of justification, so their proposal is of no use to *holistic* coherentists who deny that justification flows or transmits from one node in our structure to the next.

Similarly, here is Pryor’s (2005, 207) description of how non-doxastic coherentism is possible: “Impure [i.e. non-doxastic] coherentists are willing to give some non-beliefs, such as perceptual experiences, a justifying role. They will just deny that those states are able to justify a belief *all by themselves*. They can only do so in cooperation with other justified beliefs.” But, again, this proposal is formulated in terms that assume a linear conception of justification, and thus is not the sort of thing that can be endorsed by holistic coherentists.

¹⁸ After the main ideas in this paper were developed, I discovered that I. T. Oakley (1976, 225) and William Roche (2012, 268) also mention the sort of structure I call an entwined triad when discussing coherentism, although not to make the point I make here.

¹⁹ I am inspired here by Cling 2009, 341. Note that this version of the anti-circularism objection makes it into a version of another oft-cited worry for coherentism, namely the so-called *alternative-systems objection*. However, there is an important difference between this version of the alternative-systems objection and the version that is usually discussed. According to the standard version of the alternative-systems objection, all forms of coherentism are committed to the following: if a given subject has a justified belief in P in virtue of that belief being part of a coherent system of beliefs, then there is another possible subject who has a justified belief in $\neg P$ in virtue of that belief being part of a *different, equally coherent system of beliefs*. I agree with Richard Feldman (2003, 67) that this result should not be particularly worrisome for coherentists. But the result we are considering here is more troubling. For now coherentists appear to be committed to the following: if a given subject has a justified belief in P in virtue of that belief being part of a certain system of beliefs, then there is another possible subject who has a justified belief in $\neg P$ in virtue of that belief being part of a *system of beliefs whose propositional contents are the negations of the propositional contents of the beliefs in the first system*.

²⁰ It is true that, in the case of deductive support, if P and Q together deductively support R, then $\neg R$ and Q together deductively support $\neg P$. However: (i) it is doubtful that an analogous claim holds for inductive or abductive support, and (ii) this claim doesn’t help us show that if belief in P, Q, and R constitutes an entwined triad, then belief in $\neg P$, Q, and $\neg R$ can constitute an entwined triad, since it is not clear how we can derive the result that $\neg P$ and $\neg R$ together support Q.

²¹ Here I am relying on an intuitive understanding of what a path of partial or full support amounts to. For more formal definitions of these notions, see Berker MS-a.

²² Can partisans of the probabilistic approach resist this argument by (i) assuming that experiences have propositional content and (ii) insisting that an experience with a given propositional content plays a justificatory role if its propositional content together with various other propositions the subject believes constitute a coherent set of propositions? No, they cannot. Consider one natural proposal along these lines:

- (*) Set of mental states {experience with content P_1, \dots , experience with content P_m , belief with content P_{m+1}, \dots , belief with content P_n } is coherent if and only if set of propositions $\{P_1, \dots, P_m, P_{m+1} \dots, P_n\}$ is coherent on probabilistic grounds.

Now suppose $\{P_1, \dots, P_n\}$ is a coherent set of propositions, and compare the following:

- case 1: I have an experience with content P_1 , I believe P_2, \dots , I believe P_n .
 case 2: I believe P_1 , I believe P_2, \dots , I believe P_n .

It is not true that my mental life is equally coherent in these two cases, but (*) commits us to this unfortunate result.

What about a slightly different proposal, in which we make the fact of one's having a given experience, rather than the content of that experience, be what coheres with the other propositions one believes? For example, consider

- (**) Set of mental states {experience with content P_1 , . . . , experience with content P_m , belief with content P_{m+1} , . . . , belief with content P_n } is coherent if and only if set of propositions {<I have an experience with content P_1 >, . . . , <I have an experience with content P_m >, P_{m+1} , . . . , P_n } is coherent on probabilistic grounds.

But again we have trouble. Suppose {<I have an experience with content P_1 >, P_2 , . . . , P_n } is a coherent set of propositions. Then (**) is not able to account for the obvious difference in coherence between the following two cases:

case 3: I have an experience with content P_1 , I believe P_2 , . . . , I believe P_n .

case 4: I believe <I have an experience with content P_1 >, I believe P_2 , . . . , I believe P_n .

In Berker MS-b I argue that similar problems face all attempts to extend the probabilistic approach to non-doxastic forms of coherentism.

²³ There is also a deeper problem with the weighted-sum approach. Why do the various coherence-increasing and -decreasing factors combine together into one notion which only then grounds justification, rather than each being its own source of some *pro tanto* good- or bad-making feature of belief? For further elaboration of this point, see Berker MS-b.

²⁴ Here is another attractive feature of the combination of (C₁) and (C₂). They obey the following plausible principle:

counterclosure: If S believes P, S believes Q, S's belief in P is justified, and S's belief in P is justified partially in virtue of being deductively (or inductively, or abductively) supported by S's belief in Q (either on its own or with the help of other beliefs), then S's belief in Q is justified.

The proof of this is straightforward. Suppose the antecedent of the above principle is satisfied. Then S's belief in P is justified in virtue of being a member of some coherent set of nodes N that includes in it a belief in Q which deductively (or inductively, or abductively) supports S's belief in P. So S's belief in Q is a member of a coherent set of nodes, and hence is justified.

It is important to distinguish counterclosure from a closely related principle that only advocates of a linear conception of justification will find plausible:

linear counterclosure: If S believes P, S believes Q, S's belief in P is justified, and S's belief in P is justified partially in virtue of being deductively (or inductively, or abductively) supported by S's belief in Q (either on its own or with the help of other beliefs), then: S's belief in Q is justified *and S's belief in P is justified partially in virtue of S's belief in Q being justified*.

The combination of (C₁) and (C₂), being a form of holistic coherentism, does not entail linear counterclosure. Given (C₁) and (C₂), it is only justified beliefs which can make it the case that some belief they inferentially support is justified, but it is not *in virtue of* those beliefs being justified that they play this role. Rather, what we have is a common-ground structure: that which makes it the case that the supporting beliefs are able to help ground the supported belief's status as justified also makes it the case that the supporting beliefs are themselves justified.

²⁵ We might need to add a clause in the antecedent specifying that S believes P⁺ *on the basis of* her belief in P and its inferentially support of P⁺ (or maybe we don't, if we have already built a basing condition into our account of support among beliefs: see n. 5). I will be ignoring this complication, since it doesn't affect my overall point.

²⁶ There are those who doubt that single-premise closure holds with full generality; see, for example, Schechter 2013. However, even those who doubt the general truth of single-premise closure usually concede that it holds *in most cases*, and that is all I need for the argument in this paragraph. Many individual cases in which it is extremely plausible that single-premise closure holds would not be ones in which it holds if (C₁) and (C₂) are true.

²⁷ The combination of (C₁) and (C₂^{*}) is a natural candidate for the sort of mixed holistic/linear coherentist theory mentioned in n. 10, if we add to it the claim that any belief which is justified according to (C₂^{*}) but not according to (C₂) is justified in virtue of having a non-superseded web of support which originates in a set of justified beliefs. (This is compatible with that belief *also* being justified in virtue of having a non-superseded web of support

which originates in a set of nodes that are part of a coherent set: what we have here is a case in which X grounds Y, Y grounds Z, and X also grounds Z.)

²⁸ It is easy enough to complicate (C₂*) to make it compatible with *multi*-premise closure, if we are so inclined.

²⁹ The literature on this topic is vast. For a helpful overview, see Olsson 2008. For a book-length argument that probabilistic coherence is not truth-conducive, see Olsson 2005. For a catalogue of nearly all the contributions to the debate, see Roche 2013, 59n2.

³⁰ See Berker 2013a, 2013b, and 2015.

³¹ As should be apparent to anyone who reads Bonjour 1985 and Lehrer 1974, the classic statements of their respective coherentist views. For further evidence that Bonjour is a (subjective) truth-conducivist, see Berker 2013a, 352–53.

³² William Lycan (2012, 18) makes a similar point.

³³ This essay supersedes an earlier manuscript with the title “Graphic Coherence.” For comments on that manuscript, I thank Tyler Doggett, Tom Donaldson, Hilary Kornblith, Lisa Miracchi, Ram Neta, Jim Pryor, Gustavo Ribeiro, Susanna Siegel, Jonathan Vogel, and Jonathan Weisberg. For discussion of this material, I also thank Nazım Adaklı, Scott Aikin, David Atkinson, David Christensen, Jamie Dreier, Kate Elgin, Nina Emery, Jeremy Fantl, Jeremy David Fix, Bryan Frances, Tobias Fuchs, Georgi Gardiner, Stephen Grimm, Paul Guyer, Richard Heck, Chris Hill, Xingming Hu, Mike Huemer, İlhan İnan, Matthew Jernberg, Peter Klein, Amanda MacAskill, Elizabeth Miller, Jeanne Peijnenburg, Ted Poston, Michael Rescorla, Geoffrey Sayre-McCord, Joshua Schechter, Andrew Sciarretta, Jada Twedt Strabbing, Stephen Wilf, Michael Williams, Reed Winegar, and audiences at Brown University, Fordham University, Harvard University, Sabancı University, and the Workshop on Infinite Regresses at Vanderbilt University. Special thanks go to Jeff Overall; many of the ideas in this essay were inspired by conversations we had while he was writing his senior thesis under my direction in 2010–11. Work on this essay was generously supported by fellowships from the American Council of Learned Societies and the Radcliffe Institute for Advanced Study.