

Coherentism

Coherentism is the primary alternative to foundationalism in the history of epistemology. Both approaches are positions concerning the structure of knowledge or justification, and both arise primarily in response to the regress argument for skepticism. I will begin by characterizing the regress argument and the way in which coherentism arises as a response to it. After doing so, I will focus on the primary objections to coherentism and the primary tasks facing the view, beyond that of responding to objections.

I. Coherentism and Regress

The regress argument begins by insisting that what we know and have good reason to believe must have a basis of some sort that accounts for this positive status, and the standard form of the argument requires that such a basis involve some chain of supporting reasons. Once this beginning point is granted, the question arises concerning the nature of this chain of reasons. Here, the skeptic notes that there are three possibilities: either the chain is infinitely long, or it stops at some point, or it turns back on itself, proceeding in a circle. The skeptic then attempts to show that all three options are problematic, arguing that an infinitely long chain involves a vicious regress of reasons that cannot justify anything, that any stopping point would be arbitrary and unjustified thereby ultimately leaving the original belief with only arbitrary support, and that circular reasoning cannot justify anything. The positions targeted by these arguments are infinitism, foundationalism, and coherentism. Defenders of each approach argue that the skeptical stance taken on their favored approach is not correct.

Coherentists, in particular, argue that it is a mistake to think of their position as involving the fallacy of circular reasoning. Though an attempt to avoid the regress argument for skepticism by endorsing circular reasoning would count as a version of coherentism, coherentists have not taken this approach. Instead, coherentists are best understood as objecting to the regress argument at an earlier step, maintaining that the initial premise of the regress argument is mistaken. That step claims that knowledge and justification require some chain of supporting reasons, and coherentists typically insist that such a requirement is mistaken. They claim that such a requirement involves a linear conception of reasons, whereas positive epistemic status involves a holistic conception of reasons on which the support for any given belief arises out of its holistic relationship to an entire system of information.

Thus, whereas foundationalists prefer the metaphor of a building which rests on a suitable foundation of some sort, coherentists prefer the metaphors of a web of belief, with all parts of the web hanging or falling together, or a ship at sea having to be rebuilt enroute, regardless of where the problems are found. These metaphors suggest a picture of positive status that depends on the structural integrity of the entire entity in question, and thus that such status is best conceived of as a holistic characteristic of the system in question.

In short, then, coherentists agree with foundationalists and skeptics that the infinitist option involves a vicious regress, and side with the skeptic that there is no adequate non-arbitrary stopping point to a chain of reasons of the sort that foundationalism requires. Though it is possible to adopt these stances by endorsing circular chains of reasons, typical coherentists favor a holistic conception of reasons so that justification involves a relationship between a target

belief and an entire system of information.

II. From Schematic to Substantive Coherentism

Such a characterization of coherentism does little to identify the substance of the view, however. Instead, it identifies the region of logical space in which coherentism is found, leaving the larger project of clarifying the central elements of coherentism untouched. One way to summarize this point is to say that the response to the regress argument above identifies coherentism in schematic form only, whereas a full defense of the view will require moving from schematic to substantive coherentism.

To do so, coherentists need to discharge two primary obligations. They must first say what the relation of coherence involves, and they must identify which system of information is the relevant system for implying positive epistemic status for its members.

The standard answer among coherentists to the second question is that the system of information in question is the system of beliefs of the individual whose beliefs are being evaluated. Thus, a standard slogan among coherentists is that nothing can justify a belief except another belief, or that there is no exit from the circle of beliefs.¹

There is, however, another line of thought among coherentists that suggests a less restrictive viewpoint. Coherentists often object to foundationalism by claiming that there is no given or uninterpreted element in sensory experience, and by claiming that mere experience by itself cannot justify anything because experience is always in need of interpretation.² One way to honor these points is to insist that only beliefs can justify other beliefs, but a weaker way to do so focusing on this idea that nothing less than interpreted experience can justify anything. If this latter route is taken, the important question to answer is the question of what interpreted experience involves. On this point, coherentists might insist on a distinction between mere awareness and perceptual experience. That is, coherentists might deny that perceptual experience can be reduced to sensory awareness, and that perceptual experience itself needs to be understood in conceptual terms so that every perceptual experience has propositional content and that the subject of such an experience must have a concept for each representational item of the proposition in question.³ On such a view, the relevant system of information includes not only beliefs of the individual in question, but also experience and its contents as well.

The primary concern regarding this option of including experiences in the relevant system of information is that it may commit a theorist to foundationalism, in which case it cannot be endorsed by coherentists. In particular, the worry is that by appealing to experience in any form, a theorist has to endorse the central foundationalist element that some beliefs have intrinsic warrant or self-warrant. Foundationalists maintain that if it appears to me that it is raining, this appearance generates some degree of intrinsic or self-warrant for my belief that it is raining, whereas no such self-warrant or intrinsic warrant arises if I infer that it is raining from believing that the weatherman on TV is telling the truth about current conditions.

Such a concern, though plausible, is mistaken. Central to foundationalism is the view that some beliefs have intrinsic warrant or self-warrant, and typical foundationalists claim that such warrant arises out of the relationship between experience and foundational beliefs. Coherentists may maintain, not that this relationship between experience and belief is sufficient for such beliefs to have some degree of warrant, but rather that it is necessary. In one way, this point is already implicit in the distinction between experience and awareness noted above, for an

any interpretation of the sort involved in an experience having propositional content may be appropriate or inappropriate, and the notion of appropriateness here is an epistemic notion of propriety that coherentists typically use to insist that a legitimate stopping point for foundationalism has not yet been found. Moreover, the natural coherentist story to tell about which interpretations of experience are appropriate is a holistic one, concerning background knowledge in terms of which the interpretation arises. So the mere fact that coherentists might appeal to interpreted experience in their explanation of the relevant system of information does not turn their view into a form of foundationalism.

The other major issue for coherentists is the nature of the coherence relation itself. On this issue, coherentists have not been as helpful as one could wish. Early proposals included the bizarre suggestion that a system of information coheres just in case every member of the system entails every other member. Such a proposal severely restricts the range of justified beliefs, since contingent propositions are not entailed by necessary ones, and yet it is obvious that some of our contingent beliefs are justified.

A slightly weaker requirement interprets coherence in terms of entailment by the rest of the system of information in question. This proposal is an advance over the prior one, since it allows for the possibility of justified contingent beliefs. Even so, it isn't adequate to the defeasible justification involved in learning. I can learn that my wife is at the store from overhearing my daughter tell her to get milk and eggs in addition to the items on her list of grocery supplies, but none of my information here or elsewhere in my belief system entails that she is at the store. Nonetheless, I am justified in believing that she is, in spite of this lack of entailment.

More sophisticated accounts of coherence aimed at recognizing the defeasible and probabilistic nature of much of our justification came into the literature on coherentism in the work of C.I. Lewis.⁴ Lewis claimed that coherence among independent witnesses enhanced the probability of truth (when each source has some degree of initial credibility), and thus thought of coherence as involved probability relations as well as logical ones, with both probability and coherence relevant to the question of which of our beliefs are justified. Lewis was no coherentist, of course, and the point about the need for initial credibility of the sources is essential to understanding why he was a foundationalist rather than a coherentist, but the point to note is the introduction of the language of probability into discussions concerning the nature of coherence, an influence on the literature that is still present in current discussion.

After this introduction of the language of probability into discussions of coherence, one direction of influence led to versions of Bayesian coherentism, which clarify coherence solely in probabilistic terms. According to such an approach, a system of (partial) beliefs is coherent if it is probabilistically consistent. Such an approach involves replacing the psychological notion of belief with that of a degree of belief. For those who follow the more traditional path involving the notion of belief itself, one quite popular approach is in terms of coupling the language of coherence with that of explanation, and the most detailed account of coherence along this path is developed by Laurence Bonjour, and it incorporates explanatory connections as one element among several.⁵ In particular, Bonjour's account involves the following five features:

1. logical consistency;
2. probabilistic consistency;
3. the range of inferential connections between beliefs;
4. the inverse of the degree to which the system is divided into unrelated, unconnected

subsystems of belief; and

5. the inverse of the degree to which the system of belief contains unexplained anomalies.
(pp. 95,98)

As BonJour admits, this account is far from complete, but even if we had more detail on how the items on the list function together, this account of coherence is at the systemic level only. One important fact about the justification of our beliefs is the variability among our beliefs regarding their degree of justification. Some are justified while others are not, and within each class, some are more justified than others. An account of coherence that applies only at the systemic level cannot explain this variability.

On this score, Bayesian coherentism offers hope. If justification for a degree of belief involves probabilistic coherence within a system of (partial) beliefs, then we can distinguish different levels of justification for a given proposition in terms of what degree of belief is probabilistically consistent within that system. So suppose p and q are propositions and S is a person, and, given S 's system of partial beliefs, S can only remain probabilistically consistent by believing p to degree .8 (on a scale from 0 to 1 inclusive) and q to degree .9. In such a case, q is more justified for S than is p .

While this approach has an advantage over a merely systemic account of coherence, it has a cost as well. The notion of probabilistic coherence it employs restricts the relevant system of information to the system of partial beliefs, thus eliminating interpreted experience from relevant system of information relative to which justification arises.

A full defense of coherentism thus would need to make advances on the current options available in terms of the relevant system of information as well as the nature of the coherence relation itself. Problems for coherentism are not limited to this lack of development, however, and we turn in the next section to the primary arguments used by those who reject coherentism.

III. Arguments Against Coherentism

The primary arguments against coherentism can be grouped into three categories. There are objections deriving from the nature of coherentism as a holistic view, arguing that coherentists cannot properly explain what it is to base a belief on other beliefs or experiences. There is also the longstanding objection that coherentists cut off justification from the world, sometimes called the input problem and sometimes called the isolation problem. And there are problems concerning the truth connection, one well-known example of which is the alternative systems objection. I will discuss each of these in turn.

A. The Basing Relation Difficulty

If justification is conceived of in terms of linear chains of reasons, then it may be fairly easy to explain the difference between someone who bases their belief on the reasons that justify the belief and someone who believes the same thing, has the reasons in question for that belief, but believes on some other basis. The fundamental difference between the two is this: to properly base the belief in question, it must be based on the item preceding it in the chain of reasons that justifies the belief in question. The issue of exactly what is involved in basing a belief on a reason might still be a difficult question to answer, but identifying which particular item this relation must relate the target belief to is not difficult.

Once we adopt a holistic account of justification, however, this simple answer is no longer available. If we tried the same approach, we would require the basing relation to hold between the target belief and that which justifies that target belief. But, by hypothesis, that which justifies the belief is the entire system of information in question, typically interpreted to involve at least all of the beliefs of the person in question. The result is a strange and awkward view: justified beliefs must be based on everything in one's system of beliefs. Since the notion of basing is an explanatory or causal notion concerning the psychology of the person in question, this result has all the appearance of a *reductio*: how can the entire system of beliefs play a causal or explanatory role in the story of how I base my belief that I'm officially married on my beliefs and memories concerning a ceremony, official signings of documents, etc.? There may be a large component of my belief system that is involved here, but it isn't very plausible to think that my belief that New Zealand is closer to Australia than Texas, plays a role in the psychological story of why I believe that I'm officially married.⁶

To respond to this objection, coherentists need to deny that the basing relation is a relation between a target belief and that which justifies that target belief. One way to do this is to distinguish between features of the belief system that are essential for that system to justify the belief in question and those features that are inessential. According to holistic coherentists, no part of a system of information is sufficient in itself to justify a belief, but some parts can be essential to the story of justification even if insufficient. By distinguishing parts of the system in this way, a coherentist can maintain that the basing relation is best understood as a relation between a target belief and items in the system of information that are essential to the justification of that belief, even though it takes the entire system to explain the nature of justification.⁷

B. The Isolation Objection

A second major objection to coherentism claims that coherentists cut off, or isolate, the story of justification from the story of the world, leaving the concept of justification too independent from reality itself. In some cases, this concern arises because coherentists typically insist that justification is solely a relation among beliefs. Conceived in this way, one can imagine cases where experience continues to vary across time, but a belief system (assumed for the sake of argument to be fully coherent) is frozen. Thus, Plantinga's example of a person initially in the middle of climbing a mountain, having a fully coherent set of beliefs while doing so, but having that set remain the same after descending the mountain and then experiencing a Mozart opera.⁸ If coherence is a relation solely on a system of beliefs, one can object that such a system might be fully coherent and yet be cut off from the world in virtue of eliminating the role of experience, understood in the obvious way to be the avenue by which we enter into contact with reality, from the story of justification.

As we have seen, however, there is no good reason for coherentists to rule out experience from the story of justification. At most, they have a reason to rule out uninterpreted experience, but when experience involves propositional content, i.e., interpretation, the relevant system of information with respect to which justificatory status is assessed can legitimately be included. When it is included, this version of the isolation objection is simply mistaken: justification is not cut off from the world in virtue of the discounting of experience from playing a role in the story of justification.

This input problem, however, threatens a large variety of coherentist theories. The usual coherentist response is to try to find some necessary effect by experience on the system of beliefs, perhaps in the form of spontaneous beliefs as found in BonJour.⁹ The problem with such attempts is that the connections posited seem at best necessary only in the causal or nomological sense. Perhaps it is nomologically necessary that we have some beliefs based in experience, and perhaps it is nomologically necessary that individuals who have experiences must have those experiences reflected in the system of beliefs at least in normal cases. Even if these claims are granted, however, they are not modally strong enough, since the account of justification in question is intended to be metaphysically necessary, covering all possible cases of the justification of belief, whether nomologically possible or not. So, as long as cases like Plantinga's are logically or metaphysically possible, the problem for coherentism remains even if experience and belief bear certain important nomological connections.

Even so, since coherentism should not be understood as a theory that requires that coherence be a relation only on a system of beliefs, the isolation objection does not threaten coherentism as such. Even so, it does threaten a major subtype, one adopted by most of the major coherentist approaches in the literature.

C. Problems Concerning the Truth Connection

Issues related to the truth connection arise initially when one notes that a good piece of fiction will display exceptional coherence in spite of having no connection at all to what is actually true. A traditional way of putting this concern is in the language of the alternative systems objection, according to which nearly any belief can be justified by simply embedding it in the right system of information.

In one way, this objection is simply misplaced. One of the platitudes about justification is that it is perspectival, that what a person is justified in believing is a function of their total perspective on the world. In light of this platitude it should not surprise us that many claims are justified for 21st-century westerners that were not justified for folk in the Middle Ages or even 21st-century non-westerners. Moreover, if we interpret the objection to mean that if a given individual were to have sufficient control over their system of beliefs, any belief could be defended by making enough adjustments elsewhere in the system, that point too is not obviously objectionable. It is a fact of epistemic life that recalcitrant experience does not carry with it any definitive pointer as to which elements in our belief system need to change in light of the experience. So the mere fact that adjustments to a system can be made in a variety of ways is part of the reality of epistemic life, rather than a basis for an objection to theories that accommodate such a point.

What would be a problem is if individuals with suitable control over their systems of belief were to be irrationally motivated to preserve certain cherished beliefs by making outlandish and extensive changes to a belief system when simpler changes could have been made. As should be obvious, however, there is nothing in coherentism as such that requires it to count such an individual as justified in making such changes. To hang onto cherished beliefs is to be motivated in certain non-epistemic ways to hold and retain a given belief, and even if such beliefs cohere with a system of information, they will have a hard time passing the tests for a properly based belief, since the non-epistemic motivation will interfere with basing the belief on the elements in the system of information that are essential for the justification of that belief.

Hence, when put in the language of the alternative systems objection, worries concerning the truth connection are not decisive against coherence theories. But forcing worries about the truth connection into the language of the alternative systems objection would be a mistake, for the heart of the concern is not about alternative systems at all, but about the fact that justification is supposed to be a guide to truth and the worry is that coherence itself is not an indicator of likelihood of truth.

One version of this concern traces to Lewis's endorsement of coherence as a guide to truth. Lewis thought that agreement among independent witnesses with some degree of initial credibility raises the likelihood of truth, and Lewis-inspired treatments of the notion of coherence in terms of probability appear to be promising approaches to addressing concerns about the truth connection.

Recent work, however, has questioned this idea, relying on formal impossibility results about the connection between coherence and likelihood of truth.¹⁰ The significance of the impossibility results depends, of course, on the plausibility of the regimenting of coherentism that is required to obtain the results. On this score, it is far from clear that such regimenting has revealed anything like the essence or core of coherentism. In fact, the notion that coherence needs to truth-conducive is only one way for coherentists to respond to the problem of the truth connection. Suppose, for example, we understand truth-conduciveness in this way: for any two sets of statements where the first is at least as coherent as the second, then the first is at least as probable as the second.¹¹ One might wonder why a coherentist should prefer this kind of formal regimentation to weaker ones which require only that changes in coherence are reflected in changes in justificatory status.¹² One way to formalize such a result would be to appeal to the Bayesian notion of confirmation-conducivity. If one set of statements is positively correlated and another is not, then the first is more coherent than the second (where a set of statements is positively correlated just when the probability of the conjunction of the members of the set is greater than the product of each of the individual members). More generally, we might want variations in coherence to track variations in degree of correlation, rather than probability itself. Careful coherentists will want this relationship to be *ceteris paribus* only, since connections between probability and coherence may only be one factor to consider in assessing the overall coherence of a set, and these two points in combination avoid the best-known impossibility results.¹³

There is much more work to be done on the formal side here, but whatever the results, their significance is muted to a significant degree. First, there is the general problem for Bayesians of not having an idea of how to model the inclusion of interpreted experiences. When it comes to beliefs and partial beliefs, one constructs a model in terms of the propositional content of such beliefs, but one can't adopt this approach for experience as well, since we sometimes rationally believe the opposite of what experience reveals (as when we know that conditions are deceptive). Second, it is not clear what amount of formalization can be defended regarding the topic in question. In order to prove impossibility results, the informal notion of coherence has to be formalized, and whatever precisifications are involved will be controversial.

Moreover, there is a general problem with the usual formulations of the truth condition that deserve attention. Casual formulations often maintain that justification must be understood in terms of the best guide to truth, but this formulation cannot be accepted as it stands. First, taken at face value, it threatens most foundationalist theories as well as coherentist theories, so it would not count as some special problem that coherentists face. Second, it has a special

difficulty in explaining why we are more justified in accepting a scientific theory than we are in accepting the conjunction of its empirical consequences. It is a theorem of the probability calculus that if A entails B, then the probability of A is less than or equal to the probability of B, so if justification is understood in terms of the best guide to truth, the safer bet would typically favor not accepting the theory itself but only the conjunction of its empirical consequences. Third, casual formulations ignore the new evil demon problem. It is intuitively obvious that denizens of the relevant demon worlds have roughly the same justified beliefs we do, since they could be us, and yet nearly all of their justified beliefs are false. So a casual statement of the truth connection ignores this important fact about justification.

There is, however, one special concern for coherentists regarding the connection to truth, arising out of the lottery and preface paradoxes. In each case, the central conclusion of concern to coherentists is that these paradoxes show that justified inconsistent beliefs are possible, and this point conflicts with the idea that whatever else coherence involves, a minimal condition for it is consistency. Here the problem is that coherentists have too strong a truth connection, restricting justification to systems of information that might have only true members, whereas the lessons of the paradoxes seems to be that such a requirement is too strong. In the lottery paradox, we imagine the lottery big enough that it is justified for you to think that your ticket will lose. The same grounds for this conclusion apply to the other tickets as well, so one can deduce that no ticket will win. Hence, the combination of justified beliefs is inconsistent. In the preface paradox, conscientious authors express humility in their prefaces, noting the obvious fact that in spite of our best efforts, we make mistakes. If we suppose that all the claims in the text are justified, this expression of humility in the preface generates an inconsistency.

This problem of justified inconsistent beliefs is a special problem for coherentists not faced by other theories of justification.¹⁴ One might try to avoid it by compartmentalizing beliefs systems and requiring consistency only within subcompartments,¹⁵ but such a suggestion faces the enormous difficulty of specifying how to carve out the subcompartments appropriately. There is, however, another approach that coherentists might take, appealing to the idea that the kind of justification they intend to characterize is one responsive to the platitude that justification is what knowledge looks like from the inside. On the basis of this platitude, coherentists may insist that the kind of justification relevant to an account of the nature of knowledge is one thing, and there may be other notions of justification, even alethic notions clarifiable in terms of the goal of getting to the truth and avoiding error, that are not. So they may claim that ordinary alethic justification is present in lottery and preface cases, but that ordinary alethic justification is not the epistemic justification they intend to characterize.

In order to defend this line, coherentists need to explain what distinguishes epistemic justification from ordinary alethic justification. There are two ways to do so, one stronger and one weaker. The stronger version maintains that when we are epistemically justified in believing a claim, we also have justification for concluding that we know the claim in question. One way to argue for this claim is to argue that each of the standard four conditions for knowledge are amenable to such a requirement. Thus, one might argue that if we justifiably believe p in the epistemic sense, then we have good enough evidence to conclude that we believe p and that p is true (the first would seem to be supportable by the first-person authority we each possess regarding the contents of our own minds, and the second would seem to be a consequence of the nature of truth). The key elements for this approach would then be to show that we have good enough evidence to conclude that we are justified in believing p when we are, and good enough

evidence to conclude that we are ungettiered.

To show the former, coherentists need an argument for the claim that if there is justification of the epistemic sort for p , then there is justification that there is justification for p (call this the “JJ principle”). This claim is one endorsed by access internalists, who maintain that justification is something that, when we have it, we can detect it by reflection alone. But the claim is independent of access internalism. In fact, one way to think about the relationship between access internalism and the JJ principle is that access internalism is what you get when you try to operationalize that principle, i.e., when you try to say what would happen in behavioral terms when attempting to detect whether you are justified in believing a given claim. So access internalism should be thought of as just a special case of a class of theories that endorses the JJ principle. Such a conclusion gives resources to coherentism for defending the JJ principle since many of the objections to this style of theory are really objections to access internalism rather than to the JJ principle itself.

A full defense of the JJ principle would be beyond the scope of this article, and it is here that the distinction between weaker and stronger ways of distinguishing ordinary alethic justification from epistemic justification is significant. Defending the JJ principle is only necessary for the stronger way of drawing the distinction, and the weaker way focuses on a feature of the kind of justification necessary for knowledge that is related to the Gettier problem. Knowing involves justification for closure to further inquiry, confirmed by the fact that it makes little sense to say, “I know that it is raining outside, but I think I should go check to make sure,” or to say, “I know today is Thursday, but further inquiry regarding today’s date is probably appropriate.” Just imagine a news conference in which the athletic director says, “We, together with the NCAA, have investigated all relevant charges of impropriety in the operation of our basketball program, and as a result, we now know that all of the charges are false and that our program is completely clean. But the investigation will continue...” Such a remark would be utterly bizarre, and the reason it would be bizarre is because knowledge involves a legitimate closure of investigation. In particular, it involves an inquiry that is of sufficient quality that it licenses the conclusion that any further learning could undermine one’s present opinion only by presenting one with misleading information.

This way of putting the point should not be confused with a stronger way of endorsing the present claim about the relationship between knowledge and further learning. One may be tempted here to say that inquiry adequate for knowledge licenses one to conclude that further learning could only confirm one’s present opinion, but that it is too strong. It is too strong because of the possibility of misleading pockets of information. A simple statistical case will suffice as an example. Suppose statistical knowledge is possible, and that one’s sample of tosses has given one knowledge that a given die is unlikely to come up a six when tossed. It is consistent with such knowledge that were one to investigate further, that any finite string of sixes could result on future tosses of the die. Were further inquiry to occur and such an improbable sequence of events happen, one’s opinion would have to change in order to be rational. What is licensed by one’s present body of evidence on the assumption that this body of evidence is strong enough to put one in a position to know, however, is not the conclusion that no further learning could rationally undermine present opinion, but rather the weaker claim that any further learning that would rationally undermine present opinion would involve misleading information.

This feature by itself is sufficient to distinguish ordinary alethic justification from epistemic justification. In the lottery case, closure of inquiry is not legitimated by the quality of

one's evidence for the claim that your ticket will lose. That is why it is fully appropriate to check the newspaper tomorrow to see what the winning ticket number is. Moreover, the expression of epistemic humility at the heart of the preface paradox involves a recognition of the difficulty and complexity of the subject matter of the book, sufficient to warrant further checking by anyone reading the book, including the author! So the closure of inquiry that is central to the kind of justification needed for knowledge is not present in these cases, and hence the kind of justification that is present is only ordinary alethic justification.

For those preferring the stronger response, the response that attempts to show that when epistemic justification is present, there is justification present for the claim that we know, the above point can be used in the following way. One would only need to tie the notion of misleading information to the defeasibility theory to complete the response. It is well-known that knowledge is different from undefeated justified true belief, since some defeaters are misleading. If one's total evidence confirms that further investigation could undermine present opinion only by uncovering misleading information, then one's total evidence confirms that further learning could at worst reveal only misleading defeaters. Hence, if one is epistemically justified in believing p , one is justified in concluding that one's justification is ungettiered, i.e., that one's justification is defeated at most by misleading defeaters. Adding the point to the defenses already pointed to regarding the other three conditions, the defender of the stronger approach would now have achieved the advertised result of arguing for the claim that ordinary alethic justification falls short of epistemic justification since only the latter involves justification for the claim that one knows.

Either the weaker or stronger routes gives a general coherentist strategy for dealing with the type of justified inconsistency involves in sets of statements of more than one member. There is, however, another type of justified inconsistent belief, one involving only a single proposition. For example, Frege was justified in accepting the comprehension axiom, in spite of the fact that it leads to Russell's paradox.

The best coherentist strategy for dealing with this problem involves noting that it is not so much an epistemological problem as it is a problem in the philosophy of mind. The problem is, at bottom, how to characterize the contents of representational states such as beliefs. The standard model here takes belief to be a relation between a person and a proposition having wide content (so that the sentences "I am tired" and "JK is tired" express the same proposition), but such an account is widely recognized to be unable to explain the cognitive significance of representational states. Addressing the problem of cognitive significance requires resources that go beyond this simple dyadic account of belief. One promising new approach is in terms of the variety of two-dimensional semantics that are being developed,¹⁶ but there is also the possibility of resurrecting a Fregean approach to propositional content in terms of the sense or meaning of the terms involved in the sentences in question, or adding a third relatum to the standard view, where the third element is designed to solve the problem of cognitive significance.¹⁷ In each case, the emendations to the standard view will involve representational and semantic features beyond those involved in wide-content propositions, and these additional features will provide coherentists with the resources needed to address the problem of justified inconsistent individual beliefs. For patterns of inference and explanation that count as fully justified will need to be explained on anyone's theory in terms that go beyond wide-content propositions, so the additional resources involved in solving the problem of cognitive significance can be used by coherentists to explain why, even though in terms of the wide-content proposition expressed, a belief is

inconsistent, a focus on the representational features involved in the patterns of inference and explanation that are justified for a person will involve a focus on representational features beyond wide-content propositions, and there is reason to expect these features to involve consistency even when the wide-content proposition is inconsistent.

Conclusion

Thus, in spite of the variety of problems and challenges facing coherentism, its current lack of popularity, in contrast to its heyday in the middle part of the twentieth century, should not be viewed as a sign that it is an all-but-dead theory. Coherentists have much yet to explain and explore, but the problems raised for the theory do not show that the path ought to be abandoned in favor of more promising approaches.

Endnotes

1. See, e.g., Donald Davidson, "A Coherence Theory of Truth and Knowledge," in *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*, Ernest LePore, ed., (New York: Oxford University Press, 1989); and Keith Lehrer, *Knowledge*, (Oxford: Clarendon Press, 1974) for such claims.
2. See, e.g., Wilfrid Sellars on the myth of the given and Keith Lehrer on the "prick of sense", especially p. 188.
3. For clarification and defense of such a view, see John Bengson, Enrico Grube, and Dan Korman, "A New Framework for Conceptualism," unpublished ms located at <https://netfiles.uiuc.edu/dzkorman/www/Conceptualism.pdf>
4. C.I. Lewis, *An Essay on Knowledge and Valuation*, (LaSalle, Ill: Open Court, 1946).
5. Laurence Bonjour, *The Structure of Empirical Knowledge*, (Cambridge, Mass.: Harvard University Press, 1985).
6. The source of this objection is John Pollock's *Contemporary Theories of Knowledge*, (Totowa, NJ: Rowman & Littlefield, 1985).
7. An account along these lines is developed in Jonathan L. Kvanvig, "Coherentists' Distractions," *Philosophical Topics* 23 (1995), pp. 257-75.
8. Alvin Plantinga, *Warrant: The Current Debate*, (Oxford: Oxford University Press, 1992), p. 82.
9. Laurence Bonjour, *The Structure of Empirical Knowledge*, (Cambridge, Mass.: Harvard University Press, 1985).
10. See, e.g., Erik Olsson, *Against Coherence: Truth, Probability, and Justification*, (Oxford: Oxford University Press, 2005); and Luc Bovens and Stephan Hartmann, *Bayesian Epistemology*, (Oxford: Oxford University Press, 2004).
11. This is the Lewis-inspired connection underlying the impossibility result proved in Chapter One of Bovens and Hartmann, *Bayesian Epistemology*, (Oxford: Oxford University Press, 2004).
12. Some formal results within this general approach that can help coherentists on this point can be found in Dietrich, F. and L. Moretti. "On coherent sets and the transmission of confirmation". *Philosophy of Science* 72(3), pp. 403-424; and in Moretti, L. "Ways in which coherence is confirmation conducive". *Synthese* 157(3), pp. 309-319.
13. The first idea can be found in Brandon Fitelson's review of Bovens and Hartmann, *Bayesian Epistemology*, *Mind* 114 (April 2005), pp. 394-99; and the second idea can be found in Michael Huemer's review of Olsson, *Against Coherence*, *Notre Dame Philosophical Reviews*, published May 10, 2006 at <http://ndpr.nd.edu/review.cfm?id=6583>.

14. The *locus classicus* of this objection is Richard Foley, *The Theory of Epistemic Rationality*, (Cambridge, Mass.: Harvard University Press, 1986), pp. 96-102.
15. A suggestion first made by Wayne Riggs and first published in William G. Lycan, *Judgment and Justification*, (Cambridge: Cambridge University Press, 1988).
16. For a nice overview of these approaches, see David Chalmers, "Two-Dimensional Semantics," forthcoming in LePore and Smith, eds., *Oxford Handbook for the Philosophy of Language*, a draft of which is available at <http://consc.net/papers/twodim.pdf>.
17. See, e.g., Nathan Salmon, *Frege's Puzzle*, 2nd edition, (Atascadero, Cal.: Ridgeview Publishing Co., 1986).