

# Method Coherence and Epistemic Circularity

This is a pre-print of an article published in *Erkenntnis*. The final authenticated version is available online at:  
<https://doi.org/10.1007/s10670-017-9967-2>

## Abstract

Reliabilism is an intuitive and attractive view about epistemic justification. However, it has many well-known problems. I offer a novel condition on reliabilist theories of justification. This *method coherence* condition requires that a method be appropriately tested by appeal to a subject's other belief-forming methods. Adding this condition to reliabilism provides a solution to epistemic circularity worries, including the bootstrapping problem.

Reliabilism is motivated by a compelling thought: we care about the reliability of the sources of our beliefs. After all, epistemic norms are plausibly truth-oriented in some way or another, and so how frequently our belief-forming methods deliver the truth is of paramount importance. This intuitive thought has yielded a variety of reliabilist theories which purport to solve many epistemological problems. However, reliabilism faces a number of significant difficulties, as has been well-recognized for decades. In order to solve many of these difficulties, I will propose a new reliabilist theory of justification.

My new theory adds a novel condition for justification. It is inspired by the same intuitive notion just mentioned: that we care about how we form beliefs. It is a kind of coherence condition, but the kind of coherence I am after is distinct from previous proposals in that it focuses on the coherence of belief sources or methods, rather than on sets of beliefs. As I will argue, this new *method coherence* condition solves a number of extant problems for reliabilism. Most importantly, it provides a solution to epistemic circularity and bootstrapping objections. Although I have a favored version of this new theory of justification, what I will provide is really a schema for creating new reliabilist theories incorporating this novel condition. .

# 1 Reliabilism and Epistemic Circularity

## 1.1 Reliabilism and Reliabilisms

Reliabilism seeks to explain the justification of a belief in terms of that belief's history. To paraphrase Alvin Goldman's classic "process reliabilist" account, a belief is justified iff it is caused by a reliable belief-forming process (1979). This is an externalist theory, requiring only that the subject actually use a reliable method. It does not require that the subject knows, or can offer an argument to justify, that the method is reliable. Nor does it require the subject to have evidence for what she believes (though many reliable methods will utilize such evidence).<sup>1</sup>

There have been a variety of different reliabilist theories, besides Goldman's influential version.<sup>2</sup> These views differ regarding how they understand belief-formation, and subsequently in what the appropriate target is for evaluation of reliability. Thus, what distinguishes such views is their characterization of the way beliefs are formed. For instance, Goldman focuses on cognitive processes, while Ernest Sosa's virtue reliabilist view focuses on virtues or competences, understood as dispositions of the subject (2007, 2009).<sup>3</sup>

Reliabilism provides an attractive account of justification. It is intuitively well-motivated, parsimonious, and allows for solutions to numerous problems in epistemology. It solves the epistemic regress problem (i.e., the Agrippan trilemma), as reliabilist justification for a belief need make no appeal to other beliefs.<sup>4</sup> Moreover, it gives us a principled way of articulating a Moorean response to the skeptic.<sup>5</sup> Reliabilism is also consonant with both scientific and everyday practice. People are concerned with reliability of witness testimony, and in general with the reliability of their sources of belief. Scientists are concerned with the reliability of tests and experimental methods.

Despite these advantages, reliabilism is faced with a slew of difficult problems and objections. The new coherence condition I propose here will solve some of these problems. This condition is applicable to any view of reliabilism, and so I will largely attempt to remain neutral on the appropriate characterization of belief formation, and will speak neutrally in terms of belief-forming *methods* or *sources*.

The notion of "justification" at issue here is knowledge-level justification, sometimes called "warrant." It is the kind of justification relevant to knowledge, though for Gettier reasons it is not all that distinguishes true belief from knowledge.<sup>6</sup>

---

<sup>1</sup>For an overview of the internalist/externalist debate in epistemology, see Pappas 2014.

<sup>2</sup>Perhaps the first clear statement of reliabilism is in Ramsey's *Truth and Probability* (1926). For an overview, see Goldman and Beddor 2015.

<sup>3</sup>Armstrong (1973) and Dretske (1981) focus on reliable indicators, and Plantinga's proper-functional account focuses on cognitive mechanisms functioning as designed (1993).

<sup>4</sup>Cf. Goldman 1979, Goldman and Beddor 2015 and Sosa 2007, 2009.

<sup>5</sup>cf. Sosa 2009.

<sup>6</sup>I also intend to remain neutral on the way in which this kind of epistemic norm is related to things like "rationality" or "reasons." Some epistemologists also distinguish a notion of "justifica-

## 1.2 Two Kinds of Epistemic Circularity

The first problems for reliabilism to concern us involve epistemic circularity. There are two kinds of epistemic circularity objections to reliabilism. These two types are not always distinguished, and it is true they are related. However, many solutions offered to one do not serve as solutions to the other. My method coherence condition is motivated in large part as an attempt to solve both of these circularity worries.

The first epistemic circularity problem was pressed forcefully by Richard Fumerton (1995).<sup>7</sup> I will call this *direct circularity*, as it involves direct appeal to a method in justifying itself. It arises from the question of how we can justify a higher-order belief that a particular method is reliable. Reliabilism does not require that a subject, in order for her belief to be justified, know whether a method is reliable. Nor does she need to provide a reason for treating it as reliable. Thus, the theory seems to allow for a facile, circular justification of a method via application to itself. Compatible with reliabilist justification, the subject can form a belief that her method is reliable by using that very method. For instance, a subject could form the belief that a particular visual method is reliable by the use of this very visual method. On a reliabilist account, this higher-order belief would be justified as long as this visual doxastic method is actually reliable. Yet it seems viciously circular to suggest that we can justify the reliability of a particular method by appealing to that same method. “You cannot use perception to justify the reliability of perception!” (Fumerton 1995, 177). The same point applies to induction. What reason do we have to think that past experience is a guide for future experience? We have only the reliability of inductive inference, but this reliability is exactly what is at issue. Since this procedure seems much like begging the question, it seems we are engaging in a fallacious procedure when we use a particular method to justify its own reliability.

Furthermore, Fumerton points out, reliabilism not only allows such a procedure, it seems that the reliabilist *must* permit it use when inquiring into the reliability of a method. “There is no conceptual basis for the reliabilist to get cold feet when epistemological questions are raised to the next level up. Either reliability alone is sufficient or it is not” (1995, 178). If induction really is reliable, then according to reliabilism there is no reason why it should not be utilized to form a justified belief about the reliability of induction. This, of course, is meant to sound absurd. The suggestion is we are really no closer to understanding justification after the reliabilist account. Fumerton claims

---

tion” that is distinct from warrant or knowledge-level justification, which has to do with the actual reasons subjects give for their beliefs (e.g., Sosa 2009 makes this distinction). I am hopeful that at the end of the day we won’t need this distinction, but for my purposes here I just want to focus on knowledge-level justification. I use this terminology in order to be clear that I am concerned with reliabilism about justification (Goldman 1979), rather than reliabilism directly about knowledge (Armstrong 1973, Dretske 1981). For an overview of that distinction, see Goldman and Beddor (2015).

<sup>7</sup>Barry Stroud offers a related argument against externalism in general, including reliabilist theories (Stroud 1989, 15).

“there is no philosophically interesting concept of justification or knowledge that would allow us to use a kind of reasoning to justify the legitimacy of using that reasoning” (Fumerton 1995, 178). Epistemic circularity is vicious, and this casts doubt on the ability of reliabilism to explain justification.

The purported viciousness of direct circularity is well illustrated by the example of the person who gazes into a crystal ball to predict the future (Sosa 2007). Let us agree that in our world, at least, crystal gazing is not a reliable method for forming beliefs. So, of course, reliabilism does not suffer from the difficulty of countenancing justified crystal gazing. However, the crystal gazer can appeal to reasoning that is identical to the kind of circular reasoning that Fumerton has argued is allowed by reliabilism. The gazer can form the belief that crystal gazing is reliable by gazing into his crystal. To put some flesh on the case:

**Crystal Ball Gazer Case** Al is a fortune teller, and gazes into his crystal ball to form beliefs about the future. One day, Larry challenges Al on his justification for this. Al peers into his crystal ball, and the ball tells him that “gazing into a crystal ball is a reliable way to form beliefs.” And on this basis he forms the belief that crystal ball gazing is reliable. Al tells Larry “As long as my crystal ball is reliable, then my belief in its reliability is justified. And my crystal ball is reliable, it just told me so!”

Al uses the same kind of circular reasoning that reliabilism must allow. Per reliabilism, the reliable perceiver can, merely by appeal to perception, justify her belief in the reliability of perception. Similarly, Al the gazer can, merely by (direct) appeal to crystal ball gazing, “justify” his belief that gazing is reliable. And this similarity of reasoning, (especially the similarity from the inside) looks bad for reliabilism.

A second kind of circularity worry, known as the “bootstrapping” or “easy knowledge” problem, was first leveled by Jonathan Vogel (2000) and Stewart Cohen (2002). Vogel articulates the bootstrapping problem with what he calls the “Gas Gauge” example (Vogel 2000, 612).<sup>8</sup> In its mature form (modified from Vogel 2008, 518):

**Gas Gauge Case** Roxanne’s car has a reliable gas gauge. She has never investigated the gauge or others like it; she has no information whatsoever on the subject. Rather, Roxanne automatically forms beliefs about the level of gas in the car’s tank simply by consulting the gauge. For example, if the gauge reads “F” she immediately and directly forms the belief that the car’s tank is full. Given that the gauge is reliable, it seems clear that Roxanne’s belief that the car’s tank is full is formed by a reliable process. Now, Roxanne can also observe what the state of the gauge itself is, if she chooses to. Roxanne notes that the needle reads “F” at the time when she

---

<sup>8</sup>Per Vogel, this is taken and modified from Michael Williams.

believes, by reading the gauge, that the tank is full. Roxanne conjoins her belief that the gauge reads “F” with her belief that the tank is full, and deduces that the gauge reads accurately on this occasion. We can suppose that Roxanne repeats this strange procedure a good number of times, accumulating beliefs that the gauge reads accurately at various times. Roxanne then concludes by induction that the gauge is accurate in general, that is, that the gauge is reliable.

In the example, Roxanne forms beliefs by reliable methods (the gauge, deduction and induction). The gauge is reliable *ex hypothesi*, so her beliefs in the tank level are justified, per reliabilism. Moreover, given the background assumptions of the story, it is a logical consequence of the statements that “the gauge reads ‘F’” and “the tank is full,” that “the gauge is reading accurately.” Thus, Roxanne’s belief at each time in the accuracy of the gauge is gained by competent deduction. Furthermore, if Roxanne repeats this procedure with the same results, she can make an inductive inference that on many occasions her fuel gauge has been accurate, and therefore this gauge is a reliable one. Thus, all Roxanne’s methods in this procedure are reliable: the gas gauge is reliable by stipulation, and deduction and induction are commonly accepted as reliable. According to the reliabilist, Vogel thinks, Roxanne must now have a justified belief that the gauge is reliable, *without having done anything to check whether it is reliable beyond having blind faith*. But this is clearly problematic; this “bootstrapping” procedure is reliable but seems illegitimate. The only way to deny Roxanne justification is to deny that she knows that the tank is full in the first place. The reliabilist cannot do this, however, without giving up the central tenet of reliabilism: that a belief is justified iff it is formed by a reliable method.

Direct circularity and bootstrapping are not often distinguished.<sup>9</sup> They are clearly related: each involves the use of a method to determine its own reliability. However, the two problems are distinct. Direct circularity involves undue trust of a method, based entirely and directly on that method itself. Bootstrapping, on the other hand, requires more than just the method in question in order for the problem to arise: Roxanne appeals to deduction, induction, and memory, in addition to her gas gauge, in order to bootstrap to the reliability of the gauge. The bootstrapping problem is also easily generalized to any epistemological theory which admits of basic beliefs (Cohen 2002). It is less clear how many theories besides reliabilism are subject to the direct circularity worry.

There is a significant literature regarding epistemic circularity, primarily focused on finding a solution to the bootstrapping problem. Many of the proposed solutions will not work as a response to the direct circularity problem.<sup>10</sup>

---

<sup>9</sup>For an exception to this, see Barnett (2014), who distinguishes between bootstrapping and direct circularity, which he calls “self-verifying” circularity. Cohen (2002) also distinguishes two ways in which circularity problems arise, one via closure from basic beliefs, and the other via bootstrapping, though this latter distinction is different than the one I draw here.

<sup>10</sup>For example, Titelbaum’s (2010) solution involves a ban on “no-lose investigations”: ways

Moreover, each proposed solution has been met with objections and difficulties even in attempting to solve the bootstrapping problem.<sup>11</sup> In light of these difficulties, some have even argued that the reliabilist should accept circularity as unproblematic, on pain of skepticism (e.g., Van Cleve (2003) and Bergmann (2004)). I find this an unwelcome result.

In what follows, I argue that reliabilist views incorporating a method coherence condition can deal with both epistemic circularity objections. Method coherence blocks bootstrapping of all kinds, and provides the reliabilist with a well-motivated response to direct circularity.

## 2 Method Coherence

One attractive avenue for a rescue of reliabilism from circularity and insufficiency objections is by appeal to *coherence*. After all, coherentism about justification is an attempt to redeem a certain kind of circularity in response to the epistemic regress problem (BonJour 1985). Coherence is generally understood as a property of sets of beliefs, a property a set has when it is self-supporting in some sense. It has seemed plausible to many that certain sets of beliefs can be justified simply by appeal to the support the members offer each other, even if this is in some sense circular.

However, coherence is a disputed concept. Some theorists treat it as a kind of robust connection between consistent sets of beliefs (cf. BonJour 1985, Thagard 2000), while others have understood it as a relation of explanation (cf. Elgin 2005, Poston 2014), while still others understand it in terms of mutual probabilistic support (cf. Lewis 1946, Olsson 2005, Fitelson 2003). It is difficult to define, and when definitions are given, it becomes unclear what the relationship is between coherence and other epistemic values (Olsson 2014). In particular, whether coherence is truth-conducive has proved controversial (Klein and Warfield 1994, Olsson 2014, Olsson and Shubert 2007). Nonetheless, I think that appeal to coherence can save reliabilism from many of the difficulties raised above. However, the kind of coherence I will appeal to is distinct from that of traditional coherentist views: it is a coherence of belief-forming methods, rather than of beliefs.<sup>12</sup>

Why are we so certain that crystal gazing is not a reliable belief forming method? Why is it effective as an example of an absurd method? I would

---

of forming beliefs which cannot result in disbelief or disconfirmation. This solution rules out bootstrapping (and perhaps some other things we want to allow in), but does not rule out direct circularity. Al's crystal ball *might* tell him it isn't reliable. Similarly, Weisberg's "no-feedback" rule (2010), to be discussed below, also fails to rule out direct circularity. Of course, neither of these solutions was aimed at solving direct circularity, so this is less a criticism than a clarification.

<sup>11</sup>For an overview of the various proposed solutions, and their individual difficulties, see Weisberg (2012). Additional problems for extant solutions are raised in Barnett (2014), Douven and Kelp (2013), and Van Cleve (2011).

<sup>12</sup>I am not the first to suggest linking reliabilism with coherence. Ernest Sosa appeals to coherence in his virtue reliabilist account of knowledge and justification (2007). However, Sosa appeals to a traditional kind of coherence.

like to suggest that we could come to a more satisfying reliabilist theory of justification by attending to the “common sense” answers to these questions.

Why are we certain that crystal ball gazing is unreliable? Because gazing does not work. The crystal ball does not give us true predictions about the future in a reliable manner. How do we know it does not work? When we attempt to test the predictions offered us by the gazer, we find that they do not correspond to what we experience. Since we form beliefs about our experiences using our other methods, what this means is that the predictions offered by the crystal gazer do not *cohere* with beliefs we form using other methods like vision or hearing. Thus, the primary reason we are certain that crystal gazing is not a reliable belief-forming method is that *its predictions are incoherent with the deliverances of our other methods*.

When I am told by the crystal gazer that I will meet a beautiful stranger this week, and no such stranger appears, then the crystal gazer is “proven wrong.” To put it in the terms of the current discussion: the belief formed using crystal gazing, “I will meet a beautiful stranger this week,” is not coherent with the beliefs I form (after the allotted time has passed) with my other methods, such as hearing, vision, and memory, that “I did not meet a beautiful stranger this week.” In this sense, the crystal gazing is incoherent with my other methods. This specific sense of coherence is the one relevant to justification and knowledge. Utilizing this common sense intuition about how we actually identify reliable methods, I would like to suggest an additional condition necessary for a reliabilist theory of justification.

I will argue that adding a requirement of “method coherence” to a reliabilist theory of justification provides a better response to the circularity objections. Below, I will give a more rigorous explanation of the notion of method coherence. To see how it fits into a theory of justification, however, I am first going to provide the schema I promised above, an ecumenical version of a method coherence reliabilism:

**MCR:** A belief held by a subject is justified iff it meets the following two conditions:

1. **Reliability Condition:** It is formed by a reliable belief-forming method.
2. **Method Coherence Condition (MCC):** This method is sufficiently coherent with the other methods used by that subject.

Thus, to obtain a new theory of justification, simply plug your favored version of reliabilism into the reliability condition, and your favored characterization of belief-forming methods into both conditions. Any reliabilist can appeal to this strategy to build a new theory of justification which solves the difficulties raised above.<sup>13</sup>

---

<sup>13</sup>Note that some reliabilist theories are intended to be more generally applicable than others. That is, some reliabilist theories are only meant to apply for certain kinds of belief-formation. For example, Lyons provides a reliabilist account that only applies to certain kinds of empirical, basic beliefs (2009). I intend to remain neutral on just how general the a reliabilist theory is meant to

Before elaborating on this account, it is worth noting that one way that MCR is schematic is regarding considerations of defeat. Most reliabilists think that reliabilism needs some additional condition about how to handle defeaters in order to deal with troublesome counter-examples, like BonJour's clairvoyance cases (BonJour 1980, Goldman and Beddor 2015). Goldman initially appealed to a requirement that the subject be sensitive to other belief-forming processes available to her. The MCR schema is compatible with a variety of ways to exclude cases of defeated justification. For our purposes, we can treat the simple version of MCR as giving an account of *prima facie* justification. We can then leave open what other conditions need to obtain for *ultima facie* justification.<sup>14</sup> In what follows, I will largely set considerations of defeat aside.

There are several concepts here that must be explained. The most important are the very idea of methods being coherent, and what is required in order to be "sufficiently coherent."

One thought behind the MCC is that testing methods against each other is our only real indicator of reliability: how else are we to tell whether some method is reliably giving us true beliefs than by checking at least some of those beliefs against those formed by our other methods? We have no other option in testing our methods than by appeal to other methods. Luckily, at least in the actual world, this is a good way to tell whether a method is reliable.<sup>15</sup>

A first stab at a definition of Method Coherence is as follows:

**Method Coherence (MC):** A belief forming method *P* is coherent with another method *Q* insofar as *P* and *Q* tend to produce identical or sufficiently mutually supportive beliefs concerning the same times.

The motivation here is to account for the claim argued for above: that the reason crystal gazing is not considered a reliable method is that it does not generate beliefs that cohere with those generated by our other belief-forming methods. Underlying this definition is the notion that we can test the coherence of two methods based on the beliefs they generate regarding some particular time. Here there are three important concepts that need to be explained: "identical", "sufficiently mutually supportive" and "concerning the same times."

---

be, and the MCR schema is compatible with a variety ways of interpreting the theory. I am simply advocating that, if you are a reliabilist about justification (of some kind of knowledge), then you should also adopt the MCC (for that kind of knowledge).

<sup>14</sup>Even without adding an additional defeat condition, MCR might already provide a way out of some of the counter-examples. For instance, it's not clear that Norman the clairvoyant (BonJour 1980) meets the method coherence condition. *Ex hypothesi*, though his clairvoyance is reliable, he has never tested its reliability. However, there is not space here to consider how method coherence can help with this and similar counter-examples, and ultimately I suspect that we will still need some form of defeat-exclusion condition.

<sup>15</sup>Note that I am motivating the principle by appeal to instances when we do, consciously, judge a method reliable by testing it against other methods. However, MCC itself does not require awareness or accessibility. While we often are aware of our testing, I do not think this is a requirement of justification.



## 2.1 Mutual Support

The definition of method coherence utilizes the idea that beliefs can be “identical” and “sufficiently mutually supportive” because appeal to the “coherence” of the beliefs, in the traditional sense (as typified by BonJour), will make the MC condition too strong. Traditional coherence theories apply to a set of beliefs of significant size in order to establish coherence among them. This is because mere consistency of beliefs is not enough; for this kind of broad coherence, there must be a sufficient number of appropriate relationships between beliefs. The kind of test I am proposing involves a set of beliefs too small to establish that kind of coherence.

I take it that it is possible for us to form type-identical beliefs using two different belief-forming methods. The relevant type-identity is identity of content. For example, there are at least two ways I can discover that there is an end table in front of me: I can see the end table, or I can stub my toe on it in the dark. I can form the belief that there is a dog nearby by seeing it or hearing it bark. I can form the belief that there is something rotten in my apartment by smelling it, or by remembering that I forgot to bring the pears with me when I went for a week’s vacation.

Even when the content is not identical, we can sensibly talk about beliefs being mutually supportive. For example, I can form a belief with vision that there is a brick wall in front of me. However, if it is dark, I may only be able to form the belief with my haptic (touch) sense that there is something hard and rough in front of me. These two beliefs are mutually supporting: each is more probable in light of the other. This mutual support is relevant to judging how well the methods that formed them are functioning together.

This understanding of mutual support (and similarity) seems reasonably straightforward and unproblematic to me. However, some readers might still be concerned that this notion of support is too imprecise. In order to assuage such doubts, I will provide one way of making it more precise. One obvious way to elucidate the relevant notion of similarity is by appeal to probabilistic notions of dependence, evidential support, and coherence. This provides a more precise, formalized definition of method coherence, for those skeptical of the more intuitive notion.

The idea of mutual support can be cashed out probabilistically, in terms of evidential support. Two beliefs *A* and *B* are similar enough just when they each support each the other to a high enough degree.<sup>16</sup> The two example beliefs are mutually supporting in this sense: the probability of a brick wall in front of me is higher, given that there is something hard and rough in front of me. And that “there is something hard and rough” is entailed by “there is a brick wall.” Each belief supports the other, and it is plausible in this case that they do so to a high enough degree. Thus the two beliefs are mutually supportive

---

<sup>16</sup>I am remaining neutral about precisely which probability function we appeal to. It seems plausible that we could appeal to the subject’s own credence function, or the probability function of some (perhaps ideal) evaluator. Fans of evidential probability could even appeal to that sort of probability function.

enough, in this sense. This can be generalized: for instance, if we have a set of three beliefs,  $\{A, B, C\}$ , sufficient mutual support requires that each one is supported enough by each of the others individually, and by the conjunction of the other two. So,  $A$  would need to be supported by  $B$ ,  $C$ , and  $B \wedge C$ . And *mutatis mutandis* for sets with four members, and so on and so forth.

Helpfully, this way of precisifying mutual support turns out to be a simplified version of a probabilistic measure of coherence developed by Branden Fitelson (2003). Fitelson uses a probabilistic definition of evidential support (borrowed from Kemeny and Oppenheim (1952)) to define a measure of coherence over a set of propositions  $E$ ,  $C(E)$ . It works by looking at each proposition  $e$  which is a member of  $E$  and determining how much  $e$  is supported by every subset of  $E$ . Basically, it looks to see how well supported  $e$  is by every combination of the other propositions in  $E$ . If each member  $e$  is supported by every such combination (subset),  $E$  gets a high degree of coherence. Thus, the measure is a helpful way of determining how mutually supporting a set of propositions is.<sup>17</sup>

Fitelson's measure has a number of nice features. For instance, it turns out that sets of logically equivalent statements get a maximal measure of coherence,  $C(E) = 1$ , no matter how small the set (so my independent appeal to identical beliefs above is unnecessary on this precisification). Moreover, whenever a set is composed of sentences all of which are positively probabilistically dependent on each other (and all the conjunctions of the others), it receives a positive value of coherence.

So we can just appeal to Fitelson's measure in giving a precisified version of our notion of mutual support: a set of beliefs  $E$  is sufficiently mutually supportive just when  $C(E) > t$ , where  $t$  is a contextually determined threshold.<sup>18</sup>

Thus, a more precise account of mutual support is provided by appeal to evidential support, and a probabilistically defined notion of "coherence" for sets of beliefs. However, we need not be committed to this particular precisification. I offer it simply as an illustration of how we might make the notion precise. There are other measures of coherence that we might appeal to, or

---

<sup>17</sup>The details of Fitelson's measure are somewhat "intricate," as he puts it, though the basic idea is not. Fitelson defines a measure of evidential support (called "factual support" by Kemeny and Oppenheim (1952)),  $F$ .  $F$  is a two place function, where  $F(X, Y)$  gives the degree to which  $Y$  supports  $X$ .  $F(X, Y) = 1$  if both propositions are necessary truths, 0 if  $X$  is necessary and  $Y$  is contingent, and  $-1$  if  $Y$  is a necessary falsehood. In all other cases  $F(X, Y) = \frac{Pr(Y|X) - Pr(Y|\neg X)}{Pr(Y|X) + Pr(Y|\neg X)}$ . Once the  $F$  measure of support is defined, this can be used to determine the degree of coherence of any set of evidence  $E$ . Fitelson does this by building a set  $P_i$  which, for each  $E_i \in E$ , is the powerset of the complement  $E \setminus \{E_i\}$  (which just isolates  $E_i$  and the other members of the set  $E$ ). Then, he uses these elements to consider how much each  $E_i$  is supported by the other members of  $E$  (and all their possible conjunctions), given by  $S = \bigcup \{ \{F(E_i, X) | x \in P_i\} | E_i \in E \}$ . Finally, the degree of coherence is obtained by simply averaging the support that each member receives from the others, that is the mean of  $S$  for  $E$ :  $C(E) =_{df} \text{mean}(S)$ . For more details, see Fitelson (2003).

<sup>18</sup>This appeal to contextual sensitivity here is similar to the way reliabilist views generally appeal to contextual features to determine the relevant threshold of reliability. Since so many epistemological theories are committed to such features playing a role in setting various thresholds, I don't think this is a worrisome contextual sensitivity, or at least something particularly worrisome for my view.

other formal tools we could employ. Moreover, I think the intuitive, informal characterization is probably enough to support the view. But I am attracted to the precisified view presented.

Method coherence is a matter of degree. This explains my earlier appeal in MCR to methods being “sufficiently coherent.” We considered one way in which method coherence could be degreed: regarding how mutually supportive the produced beliefs are to one another. But we might also ask how many of our methods are producing mutually supportive beliefs. How often does this method produce beliefs that cohere well with our other methods? Although it may be difficult to specify a procedure for determining the exact degree of coherence, the fact that method coherence is graded fits well with the idea that methods will have varying degrees of reliability. Induction has a great deal less conditional reliability than deduction, in the sense of percentage of true beliefs generated from true premises.

The method coherence condition captures our intuitive sense that testing our methods against each other is our only access to their reliability. Since this kind of coherence is proposed (in part) as an indicator of reliability, it only makes sense that the coherence would admit of degree just as reliability does. It may be difficult to nail down marginal cases (Vegas odds in sports predictions, divining rods, etc. . . ), but it is clear that we are often quite easily able to distinguish between methods with a low degree of coherence (crystal gazing), and those with a high degree of coherence with our other methods (vision).

## 2.2 Method Coherence as a Test

The idea that the beliefs formed should “concern the same time” is also crucial to this conception of method coherence and to the MC condition. We want to compare beliefs about how things stand in the world at a given time, where that time might be the current time, or some future time. This allows for determining how well methods cohere both by making predictions and by more simultaneous belief formation.<sup>19</sup>

The crystal gazing case provides an example of prediction. In that case, the belief formed is a prediction about the future, i.e., a belief concerning a future time. In this case, the beliefs formed will fail to be identical, or sufficiently mutually supportive of, the beliefs formed by vision, hearing, introspection, or reasoning. Other methods that we use, however, form predictive beliefs that are sufficiently mutually supportive to the beliefs that we form when the time comes. For instance, we make predictions using deductive and inductive inferences. Using induction, I predict that the sun will rise tomorrow morning. If I see the sunlight tomorrow, and feel its warmth, then the beliefs formed by

---

<sup>19</sup>On some views of mental content, the proposition a belief is held toward will already include a specific time (e.g., if a proposition is a set of possible worlds). On views like this, this added stipulation that the beliefs concern the same time will be redundant, and could be dispensed with. However, I want to remain neutral on questions regarding the nature of propositions and of mental content.

these methods will be identical with, or will mutually support, the prediction. Then, my simple enumerative induction method is, in this instance, coherent with my other methods.

Whether methods are coherent depends only on beliefs which are formed concerning the same time. This requirement is helpful in two ways. First, it allows us to treat method coherence as a kind of test, and second, it makes the requirement of method coherence less stringent and more realistically attainable.

Method coherence is not a global coherence, as Bonjour would require, or the kind of broad coherence Sosa suggests for reflective knowledge. It is a narrow kind of coherence. Instead of testing the coherence of a belief against an enormous web of beliefs, a subject is instead testing the coherence of her methods by comparing the beliefs delivered by them in specific situations. This corresponds to our common sense procedure of testing a method, as in the crystal gazing case. We are certain that crystal gazing is not a reliable method, because when we test the beliefs it generates by comparing them to beliefs generated by our other methods, it fails the test. The beliefs crystal gazing generates are not coherent with those generated by our other methods. Induction, on the other hand, coheres well with our other methods, because it produces predictions that, when tested in this way, are often enough identical or sufficiently mutually supportive to our other methods' deliverances.

Method coherence, then, requires being able to pass a certain kind of test. A method is coherent when it can pass such a test, and more coherent the more it tends to pass. Passing the test requires that the method delivers a belief that is identical to or sufficiently mutually supportive with the beliefs delivered at the same time by the subject's other methods.

Of course, this kind of test is precisely what we use whenever we consider the reliability of some new instrument, exam, or procedure. Anytime we have some new way of collecting data or evidence, we test it by appeal to our existing belief-forming methods. When the first radar speed gun was developed, it was tested for reliability using existing methods, probably using a car on a road with a measured length and a stop watch, as well as using a car's internal speedometer. Only after a method has shown itself capable of repeatedly passing such tests do we come to trust it. What my account does is extend this notion so that it applies not just to new methods, but to all of our methods. After all, *pace* Descartes, we do not have some perfectly reliable method we can use to calibrate each of our new methods against as we begin using them. Instead, we "calibrate" our methods by appeal to our other methods.

Method coherence is thus similar to a kind of calibration of our methods, but without any external standard to calibrate with. It is similar to the way a group might calibrate their watches before engaging in some highly time-sensitive activity. Each must be appropriately calibrated with the others, but not necessarily with some standard timepiece (such as the U.S. Atomic Clock). It also echoes the notion of "robustness" appealed to by William Wimsatt (2007), in which we are more confident in a theory when there are multiple kinds of tests confirming it. Moreover, the motivations for method coherence

are similar to our motivation for thinking that independent sources of evidence are superior, a fact which Bayesian confirmation theory helps explain (Fitelson 2001, Earman 1992).

Method coherence requires passing of a test, where passing means producing sufficiently mutually supportive beliefs. One might wonder, however, which beliefs need to be mutually supportive in order to pass the test. Presumably, during any particular episode of belief formation, an agent may form an enormous number of beliefs. Which beliefs count toward determining method coherence? The notion of a method coherence test, coupled with our idea of mutual support, provides an answer. Specifically, any belief formed during a relevant testing episode which concerns the relevant time will appropriately count toward determining coherence. Whether or not the set of beliefs thus formed do count as coherent then just depends on whether they are adequately mutually supportive. If the relevant beliefs formed for the test ( $E$ ) are adequately mutually supportive ( $C(E) > t$ ), then the test is successful and the methods being tested are adequately coherent. This set of beliefs might be quite large, but this poses no obstacle to evaluating them for mutual supportiveness.<sup>20</sup>

I have been clarifying the central notions employed in the definition of method coherence. We have established that MC requires that a method pass a test, whereby it produces mutually supportive beliefs with other methods, and that it tends to do so sufficiently often. The question now is about how to understand “tends to.” I think the best understanding of “tends” is the following: the method in question tends to pass method coherence testing when it has actually done so frequently enough in the subject’s history. MC thus requires that the subject has actually had her methods pass such tests in the past, and done so frequently enough.<sup>21</sup> These past tests need not be anything the subject is aware of, nor do they need to be something the subject remembers having done. It is simply required that the subject has done so, and the method has passed frequently enough.

One might worry that this understanding of “tends” makes method coherence too strong a requirement. One might even worry that it seems quite internalist, contra the spirit of reliabilism. However, it is not nearly so strong a requirement as that of accessibilism or mentalism, the two standard forms of internalism (cf. Conee and Feldman 2001). A method coherence test does not have to be done deliberately or consciously: the subject need not have access to the fact that the test has been passed. So accessibilism is not required. Nor does the subject need to remember, or have any current evidence (that supervenes on her current mental state) that their methods have passed coherence tests. Thus mentalism is not required either. Instead, MC only requires that the agent have actually tested her methods.<sup>22</sup>

---

<sup>20</sup>Thanks to an anonymous reviewer for raising this issue.

<sup>21</sup>Whatever contextual, social, evolutionary or other features that determine the threshold of reliability for a method can also determine the frequency required for passing method coherence tests. As above, cf. Sosa 2007, 2011 for more on how thresholds of reliability are set.

<sup>22</sup>It does, however, require more than that the subject simply notice if a method is incoherent.

MC is also not as strong a requirement as traditional coherentism (BonJour 1985), as I pointed out above (§2.1). It does not require that all of a subject's beliefs are consistent, interconnected by inference and explanation, are mutually probabilifying, and avoid anomalies. Rather, it merely requires a small subset of beliefs concerning a particular time to be the same or to support one another. As I will argue below, the strength of the requirement is necessary for solving both kinds of epistemic circularity worries.<sup>23</sup>

So although method coherence is clearly a stronger requirement than mere reliability, it is not so strong as internalism or coherentism. It thus avoids raising the bar implausibly high for what kinds of creatures are able to have justified beliefs. Although it is unlikely that most children will be able to articulate anything as sophisticated as “method coherence testing,” children are still capable of passing method coherence tests.

Method coherence is thus an externalist requirement. Satisfying the MCC merely requires that the subject have actually successfully tested her method for coherence in the past. It need not be the case that she knows this, or has any kind of access to or awareness of this fact. As I will argue below, adding this requirement provides an answer for the reliabilist to both kinds of circularity worries, without the move to a genuinely internalist theory that would present a more stringent requirement.<sup>24</sup>

### 2.3 Method Coherence: Precise Version

The preceding discussion provides for a more precise version of method coherence:

**Method Coherence (Precise):** A belief forming method *P* is coherent with another method *Q* insofar as *P* has (frequently enough) passed method coherence testing with *Q*, where passing such a test means the two methods produced beliefs concerning the same time which have identical content or have content which is adequately mutually supporting (i.e., “Fitelson-coherent”).

---

Cf. Kent Bach's “take-for-granted principle” which essentially requires that the subject would become aware if something in a situation was “out of the ordinary” (1985, 261), or “No Defeater” conditions (Beddor 2015).

<sup>23</sup>One might be tempted by a counterfactual understanding of “tends,” as a different way of avoiding the problems I have just mentioned. However, such a counterfactual version of the MC requirement fails to address the bootstrapping problem, as it would allow that Roxanne is justified, because *if* she checked the gas gauge, it would be coherent (since, by hypothesis, it is reliable). The problem is that she does no such testing.

<sup>24</sup>It is worth noting, however, that there is a nearby version of the view that *does* involve having a limited accessibility requirement. This minimally internalist version of the view would be more stringent than MCR, but still be significantly less stringent than a view which involves full accessibilism. This version of MCR would simply involve the further requirement that the subject have access to the fact that her methods are coherent. That is, a subject must be able to become aware of the fact that she has tested her methods in the past. This version might be amenable to those who seek a more robust conception of justification, while maintaining many of the benefits of reliabilism, and avoiding some difficulties for full-blown accessibilism. Thanks to an anonymous referee for discussion on this point.

This more precise version is optional, given the way the terms in the initial version have been subsequently defined and explained.

### 3 Method Coherence: Solutions for Reliabilism

Before moving on to detail how MCR provides solutions for reliabilism, it is worth noting the tight connection between method coherence and reliability. Given an overlap in the content of beliefs produced, the more reliable a set of methods is, the more method coherent it will be. This is most clear at the extreme end: if a set of methods is perfectly reliable, and overlaps completely in the relevant content area, then the method is guaranteed to be perfectly method coherent. This illustrates the connection between the notions. Moreover, I suspect that there will be interesting connections between the three factors (reliability, content, and method coherence). As you vary the degree of reliability of the methods in question, for instance, I expect you will also see a variation in coherence. Exploring these connections in greater detail is an avenue for future research. What is important at this juncture is that, far from being a radical departure from reliabilism, MCR is a natural extension of it. Method coherence is connected with reliability. Moreover, the view continues the emphasis that reliabilism has regarding the importance of how we form beliefs.

The addition of method coherence to reliabilist views provides solutions to both kinds of epistemic circularity. I will address each in turn.

#### 3.1 Direct Epistemic Circularity

A reliabilism incorporating the method coherence condition does not allow a method to be justified *merely* by appeal to itself, and therefore avoids vicious direct circularity. The MCC requires that a belief about the reliability of a method be formed by a method which has been tested against other methods. So, even if one were to engage in the type of reasoning Fumerton is concerned with, it would not be merely the reliability of the method in question doing the justificatory work. Consider the toy case where a subject is forming a belief about the reliability of visual perception by using the method of visual perception. The visual doxastic method, if it is to satisfy the requirements of MCR, must have previously been tested against other doxastic methods. But then its power to confer justification on beliefs is partially due to its method coherence.

On this account, then, a belief is justified in virtue of the fact that the method by which it was formed has been method coherence tested by the subject using their other methods. So the belief is justified partially in virtue of these other methods, rather than merely by appeal to the method in question. Thus, it is not true that a subject may justify her belief in the reliability of a method *merely* by appeal to that very method.

Relying on method coherence means that there is still some sense of circularity present. After all, there is no assumption of a Cartesian-style perfectly reliable method by which we can calibrate our other methods. So the methods are merely coherent with one another.

Fumerton objects to the use of a method in justifying itself. Perhaps, one might think, the use of one method to justify another, which would justify another, etc..., would be just as objectionable. Method Coherence, however, involves holistic, non-linear justification: the coherence of the set of methods justifies the use of a particular method. Instead of allowing the direct, linear, and two-step circularity that Fumerton finds objectionable, the MC condition instead widens our scope: the subject is justified in the use of some method by appeal to its coherence with other methods. In other words, beliefs are justified partially in virtue of their method of formation being justified (i.e., successfully tested) by other methods. This seems to be a significant improvement over a standard reliabilist account, which Fumerton rightly points out, must resort to linear epistemic circularity.

Still, Fumerton (or someone with similar commitments) may have a lingering worry about direct circularity. Method coherent reliabilism states that a belief is justified just when it is formed by a method that is both reliable and method coherent. On this account, if some method *M* meets both conditions, then a belief formed using *M* about the reliability of *M* will count as justified. That is, there still *appears* to be some direct circularity: it seems that a belief in the reliability of a method is justified in virtue of that very same method.

I think this lingering appearance is mistaken. MCR does not allow a belief in the reliability of a method *M* to be justified *merely* in virtue of being formed by *M*. While it is true that the belief will receive justification through being formed by *M*, the fact that *M* is capable of conferring justification is due to its coherence with other methods. That is, the belief in method *M*'s reliability is justified partially in virtue of the subject's other methods. In order for *M* to get in the game of conferring justification, it must first pass the MC test. In some sense, then, *M* itself must first be justified, must first receive support from other methods. *M*'s ability to justify thus depends on the subject's other methods. Any justification conferred by *M* must then depend, at least in part, on those other methods.

Still, a subject employing a method need not be aware of its reliability or of its method coherence. According to MCR, if I appeal to a method of visual perception in judging the reliability of that very method, I need not have any evidence, memory, or awareness of the fact that this method meets the MC condition. I need not remember having tested the method for coherence. From the inside (that is, from the subject's first-person perspective), it will seem just like the original case of direct circularity.

At this point, however, it seems that the objector is simply demanding some form of accessibilist internalism. That is, the objector wants the subject to always be able to tell that there is no direct circularity. But MCR is an explicitly externalist view. I do not expect that the arguments I have provided here will convince someone to give up internalism. I take it that there are strong reasons



to give up internalism, though this is not the place to rehearse them. Instead, I am providing a solution for a purported problem for reliabilism: that the reliabilist allows direct circularity. MCR forbids vicious direct circularity. This means the reliabilist is not forced to give up reliabilism in the face of the objection (though obviously the solution involves a modification of the view).<sup>25</sup>

To be clear, I suspect that this line of reasoning will not persuade the committed internalist, such as Fumerton, to give up their theory in favor of method-coherent reliabilism. Instead, I offer this as a solution for the reliabilist who still feels the pull of Fumerton's objection. Once we adopt MCR, the objection is defused, for the reliabilist. After we adopt MCR, what remains of the objection just is the internalists' commitment to accessibilism. It is not my project here to dissuade them from that commitment.<sup>26</sup>

The method coherence condition also has the advantage of denying justification to subjects like the crystal gazer, who are obviously unjustified. As Sosa (2009) points out, the crystal gazer can have an extremely coherent (in the traditional sense) belief system that nevertheless seems wholly divorced from reality. This results from the ability of the gazer to consider only certain inputs, or to entertain only beliefs that fit with his chosen world-view, and thus avoid any possibility of contradiction. In the case of the gazer, this is cashed out as speaking in riddles, or making unspecific predictions, in order to maintain the illusion that the method is reliable and coherent. A view incorporating the MCC as part of the definition of justification will avoid this problem because it requires the testing of predictions in a very specific, concrete way. It requires testing of particular beliefs concerning particular times, by appeal to other belief-forming methods used by the subject. So, it precludes the gazer's method of avoiding incoherence: when the gazer makes an unspecific prediction, this will be tested by appeal to actual beliefs formed using perception, induction, and other methods. In order for the predictions to be coherent with later beliefs, they would need to be adequately mutually supporting. And the less specific a prediction, the less support it will gain or receive from these other, specific beliefs.<sup>27</sup>

---

<sup>25</sup> If one is still concerned about the use of a method to justify itself, this discussion suggests another solution that makes use of method coherence, although it departs from my schema for justification. One could argue that any time a subject forms beliefs about the reliability of a method, they must do so by the method coherence testing criterion. That is, second order beliefs about whether a method is reliable can only be justified by the process that ensures they are method coherent. This has the benefit of the intuitive appeal of the method coherence procedure, and of making it harder to form justified beliefs about one's own reliability. However, it has the drawback of having to postulate a special kind of justification for a certain set of beliefs. The justification schema I offer has the benefit of making justification a unified notion.

<sup>26</sup> Some reliabilists are not perturbed by the direct circularity objection in the first place. Some are convinced by Van Cleve (2003) that circularity is both unavoidable and acceptable, as discussed below in §4. The solution presented here to the direct circularity objection may seem less pressing to the reliabilists who are so convinced. However, MCR does also provide a solution to the bootstrapping problem. And it seems nice to have more than one available response to the direct circularity problem. So I think that method coherence should still be attractive to these reliabilists, as well.

<sup>27</sup> The intuitive version of the view gets us this result: being less specific means a proposition

This accounts for what it is about the gazer's reasoning that is illegitimate in a way that appeals to broad coherence cannot. Thus, we now have a way of faulting the direct circularity illustrated in the crystal gazer case.

### 3.2 Bootstrapping and Easy Knowledge

A reliabilism that includes the MCC rules out the bootstrapping procedure in the gas gauge case. Roxanne will not be able to "bootstrap" to knowledge about her gauge's reliability. She will not have knowledge (or be justified in believing) that her tank is full in the first place because her belief fails the MC condition. By stipulations necessary for the case, she has not (consciously or otherwise) tested the gas gauge against her other methods to establish its coherence. While the gas gauge is reliable, and would thus satisfy the reliabilist condition of MCR, Roxanne's use of it does not satisfy the MC condition.

MCR blocks Roxanne's illicit bootstrapping procedure from the first step: she does not know that the tank is full, when her gauge reads "full," because this method has not been tested for coherence with her other methods. Therefore, Roxanne would not have knowledge that "the tank is full" to use as a basis by which to deduce the fact that the gauge reads accurately at any particular time, let alone enough times to underwrite an inductive inference that the gauge is reliable.

Notice that the MC condition does not deny the average motorist knowledge about whether her gauge is reliable. Roxanne has been described as a bad epistemic agent: she makes no attempt to check the gauge's reliability. I think it is reasonable to suspect that most motorists will check their gauges for reliability, simply through the common everyday activities of driving and filling the car with gas. If the gauge is not reliable, this will become quickly apparent to the driver (e.g., when the car runs out of gas). An attentive driver also notices the number of gallons that she puts into the gas tank when the gauge reads "half," "quarter," or "empty," and checks the owner's manual of the car for the tank size. These procedures allow a driver to test the coherence of the gas-gauge-based doxastic method with her other relevant methods, and thus establish the necessary level of coherence for the use of the gauge to meet the MCC.

Thus, a reliabilist theory patterned on MCR, incorporating both a reliability requirement and a method coherence requirement, accounts for why boot-

---

rules out fewer possibilities. It will thereby receive and provide less support to other propositions, and will thus be inadequately mutually supportive of the deliverances of other methods. However, MCC will also rule out methods which produce only unspecific predictions on the precisification I give above in terms of Fitelson's measure of coherence. (At least, this is true as long as the belief-forming methods standardly used by subjects do not also result in beliefs with *only* unspecific content that is logically equivalent to the predictions offered by the crystal gazer. And I think this is very plausible to assume about ordinary subject's perceptual belief-forming methods.) This is because Fitelson's coherence measure appeals to a probabilistic support relation. An unspecific theory, which is compatible with a great deal of evidence, is supported less by evidence *E* than a more specific theory which predicts *E*, due to the law of likelihoods (for details on this feature of factual support, see Fitelson (2003) and Kemeny and Oppenheim (1952)).

strapping is an illegitimate procedure, without denying knowledge to those subjects who intuitively have it. MCR explains the illegitimacy of the first stage of the bootstrapping procedure. However, as we will see in a moment, the method coherence will also block the inductive stage of the bootstrapping procedure, since Roxanne also seems to be using an illicit inductive procedure.

### 3.3 Slight Bootstrapping

Jonathan Weisberg (2010) raises an additional problem for solutions to the bootstrapping problem: these solutions fail to block cases of “slight bootstrapping.” This problem illustrates something seemingly illegitimate about subjects who use an inductive procedure like Roxanne’s. These are cases where, for instance, the subject moves from knowledge about reliability to “super-reliability” in a seemingly suspect, bootstrapping way. “Super-reliability” denotes extremely high reliability.

**Super-Reliable Gas Gauge:** Charlie knows that the gauge in his car is reliable, and it is in fact super-reliable. On one occasion the gauge reads *F*, leading him to believe that the tank is full, which it is. He notes that on this occasion the tank reads *F* and is full. He then repeats this procedure many times on other occasions, coming to believe that the gauge is not only reliable, but super-reliable (Weisberg 2010, 532).

Weisberg’s Super-Reliable case shows that other solutions, including a retreat to internalism, fail to block these cases of bootstrapping. Requiring antecedent knowledge of reliability is not enough to avoid bootstrapping problems.

Weisberg offers a solution to bootstrapping which blocks bootstrapping at the inductive step, including in such slight bootstrapping cases. His solution involves a new condition on inductive methods, which he calls “No Feedback” (533).

In a nutshell, Weisberg’s No Feedback condition outlaws inductive methods with illicit lemmas, that is, lemmas which provide greater support to the conclusion than the base evidence already does.<sup>28</sup> So any conclusions reached by inductive inference cannot be simply added to the inductive base in order to make a further inference, because this seems to be illicitly adding evidence to one’s inductive base. As a classic example, suppose I have come across 100 ravens and each one has been black. I now have a strong inductive argument for the conclusion that the next raven I see will be black. Suppose I form a belief that the 101st raven is black, on that basis. If I then go on to add this

---

<sup>28</sup>The precise version is:

**No Feedback** : If (i)  $L_1 - L_n$  are inferred from  $P_1 - P_m$ , and (ii)  $C$  is inferred from  $L_1 - L_n$  (and possibly some of  $P_1 - P_m$ ) by an argument whose justificatory power depends on making  $C$  at least  $x$  probable, and (iii)  $P_1 - P_m$  do not make  $C$  at least  $x$  probable without the help of  $L_1 - L_n$ , then the argument for  $C$  is defeated.

belief to my evidence (my inductive base), I have an even stronger argument that the 102nd raven I encounter will be black, and so come to believe this, and so I will come to believe this and add this premise to my base as well. I can repeat this procedure as many times as I want. This seems clearly illicit; only the ravens I have encountered should count as part of my evidence.

This kind of feedback procedure can lead to absurd results, and this is precisely what we see in the slight bootstrapping case. What *No Feedback* rules out is taking the belief “The tank is full and the gauge is reliable” as a lemma, and adding it to the inductive base for the final inference to “the gauge is super-reliable.”

One limitation of No Feedback, from the perspective of the reliabilist looking to employ it, is that it does nothing to block Roxanne from having knowledge of the first step of the bootstrapping. That is, Roxanne counts as justified in believing that her tank is full when it reads full. But she is a bad epistemic agent, *ex hypothesi*. Method coherence, on the other hand, blocks Roxanne from being justified in her belief about the fullness of the tank.<sup>29</sup>

I think that Weisberg is right that No Feedback is a rule that good inductive methods respect. However, MCR will ensure that one is only justified in using methods which respect No Feedback. This is because adding lemmas to one’s evidence before making a further induction will lead to *weak*, and thereby unreliable, inductive methods. When such feedback-allowing inductive methods are given a method coherence test they will (at least usually) fail, because the methods being used to test them will be reliable. A reliable method and an unreliable one will tend to fail to produce sufficiently mutually supportive beliefs when applied in the same situation regarding the same topic. And so a feedback-allowing method will be method incoherent with reliable methods, at least in the actual world. So, any inductive method that fails to respect No Feedback will fail to be method coherent, because it will allow one to produce beliefs which tend to fail method coherence testing.

Thus, feedback-inclusive inductive methods will qualify for neither the reliability condition, nor the method coherence condition in my justification schema. While I think that Weisberg’s No Feedback principle is correct, it does not provide a full solution to bootstrapping. But in the actual world (and worlds like it), it will be captured by a method coherent reliabilist view.

In virtue of requiring No Feedback, MCR theories will block slight bootstrapping, while also ruling out problematic inductive methods. This, without having to postulate an additional requirement on a theory of justification: methods that are method coherent and reliable will respect No Feedback. MCR will thus deny justification to bad epistemic agents like Roxanne and Charlie.

---

<sup>29</sup>Note that this is a drawback for the reliabilist hoping to employ the No Feedback solution. Weisberg is defending evidentialism, rather than reliabilism. Evidentialism does not grant Roxanne justification at the first step, so the objection does not apply to an evidentialist account of the Roxanne case. However, there are other cases of basic knowledge that many evidentialists do grant, which allow for similar kinds of bootstrapping or “easy knowledge” (see Cohen 2002). No Feedback might suffer from this drawback in some of these cases: it fails to block justification that is too easy at the first step.

## 4 Objections

There are several possible objections to MCR, and to the solutions it provides for the circularity problems, that are worth addressing immediately.

First, the characterization of belief-forming methods here is quite vague, and this may affect whether the type of testing procedure necessary for the MCC is plausible. However, I have allowed the description of methods in this account to be non-specific because I think the MC condition is widely applicable to reliabilist theories regardless of their generality or specificity in characterizing methods. Whether a method is characterized broadly (e.g. “vision” and “induction”) or very narrowly (e.g. “the faculty of recognizing blueness in objects under twilight conditions”), the MCC could be used effectively to block the epistemic circularity objections and garner the other advantages I have suggested.

This worry might be strengthened by pointing to the fact that often beliefs are formed by appeal to multiple perceptual modalities, and in fact people are more reliable and successful at tasks when utilizing multi-sensory integration (Nagel 2014). Thus, it may be difficult to pull apart independent belief-forming methods in order to engage in the relevant testing. However, I think it is plausible that a subject will often be able to individuate between different methods, where we have cases of multiple methods operating together. Moreover, in many of these multi-modal cases, it is very plausible that there is really only one *belief-forming* method, even when multiple *perceptual* modalities are involved. So there are some difficult technical issues that arise if we attempt to individuate doxastic methods by appeal to psychology. Ultimately, my view requires a solution to the generality problem, and it is far from being alone in this requirement. But there is no special reason from multimodal perceptual processing to think that we can’t distinguish different belief-forming methods that can be tested against one another.<sup>30</sup>

These considerations lead to a second potential objection: that my reliance on using different methods to test each other requires that there is a solution to the generality problem, since we must be able to distinguish these methods from one another. This seems correct. However, reliabilists are already on the hook for providing a solution to this problem, and I do not have space here to solve this thorny problem. Perhaps it is worth pointing out, however, that recently it has been argued that the generality problem applies to any view of doxastic justification (Comesaña 2005, Bishop 2010). Thus, I do not think the objection is particularly worrisome for my view, compared to other reliabilist views, and perhaps compared to any current theory of doxastic justification. Everyone needs a solution to the problem. I am optimistic that a solution will eventually be found, but I suspect that this solution will require significant em-

---

<sup>30</sup>Moreover, there is some evidence from the metacognition literature that we receive feedback on how well different methods perform, including how well they perform at the same time (Proust 2013). And this kind of inward tab-keeping on the reliability of methods may very well provide support for method coherence. Though see Nagel (2014) for an opposing view of this research. There is much left to be said on this topic, but I will set it aside for future work.

pirical research into the actual psychological and neurological processes that underly belief-forming methods. How to type human belief-forming methods seems like a question that depends on the way the human psychology works.<sup>31</sup>

Another objection involves the possibility of untestable methods: beliefs formed by one method that are not amenable to testing by other methods. For instance, we could imagine that there is a particular belief about a smell or about a color that is not testable by comparison to our other beliefs. These would be beliefs such as “I smell pie” or “this shade of blue looks like the ocean,” or perhaps most worrisome of all “this is blue.” These kinds of beliefs may seem to be in principle untestable, and thus to be something that could not count as justified under MCR. There are several possible responses to this objection.

First, I suspect there are few belief-forming methods that would qualify as untestable. When confronted with a belief about a particular shade of a color, for example, generally I would compare my perception of the color with the reports of other people’s perceptions to determine if I was correct in comparing it with other instances of color. This would qualify as testing the belief, since testimony<sup>32</sup> is generally a reliable belief-forming method. Episodic memory and induction are both methods as well, so in fact there would be several methods that could be used to test this color perception, even though beliefs about color are characteristically formed by vision. Note that this is similar to how some people might actually come to find out that they are color-blind. One might consistently judge that two lights are the same color, until they are told otherwise. This would then show that the belief-forming method in question is incoherent and thus suspect.

Moreover, if we have a belief formed by one method that cannot be validated by other methods, we often take this as evidence that there is something wrong with that method, whether that is a temporary unsuitability to the situation, or something more drastic. Just as in the case of the carefully vague beliefs formed by crystal gazing, the fact that the beliefs are untestable is a reason to think that something may be wrong with the method. This seems akin to problems we have with scientific theories which are untestable, or unfalsifiable.

If one is not yet convinced by these considerations, I wonder how problematic the cases left over will really be for the view. The most worrisome kind of case will involve judgments about phenomenal character or our mental life, such as “this appears blue,” or “blue is like *this*,” or something like “I think, I exist.” But such judgments are those of which we can (perhaps) be most certain. However, reliabilism, and MCR, are theories about how we can come to know things, even though we are fallible. It doesn’t seem like much of a concession to admit that MCR doesn’t apply to those (very few) cases in which we are most certain. The project of trying to understand all of knowledge and

---

<sup>31</sup>For more on the generality problem, its application to non-reliabilist theories, and on potential solutions, see Goldman and Beddor 2015.

<sup>32</sup>This could also be made more specific, as in “testimony from John.”

justification in terms of the way we come to know this small set of things is impossible. So we need some other account of our less certain knowledge and justification, and that is what MCR provides.<sup>33</sup>

Another worry one might have about the MCC is that it is too stringent for very simple creatures to satisfy it. If a creature had only one belief-forming method, for instance, then it would be unable to achieve method coherence. This does not seem like a problem to me, however. I doubt that creatures with only a single method for forming “beliefs” would actually be cognitively sophisticated enough to count as believers at all. This is because having beliefs seems to require having mental states with a sophisticated functional profile. Beliefs are mental states of creatures capable of cognition. It is unclear that a system without significant interaction between states would count as cognitive. If a creature only has one kind of “belief”-forming method, it doesn’t seem like it will have enough variety in its belief-like states in order to count as being cognitive, and thereby even count as having genuine beliefs.

For instance, consider a magnetotactic bacterium, which orients itself based on earth’s magnetic field. It seems incapable of having beliefs, and of having epistemic justification. This is despite the fact that it has a reliable way of determining the direction of oxygen, and acts so as to satisfy its survival requirements on the basis of this determination. It is not sophisticated enough to count as a believer, because it does not seem that any cognition is occurring. Rather, it just reacts to external stimuli. Moreover, even if such a creature did count as having beliefs, it seems intuitive to me that they would not count as being epistemically justified in their beliefs. Epistemic justification is another feature that seems restricted to relatively sophisticated creatures. After all, even committed reliabilists recognize the epistemic importance of a subject having sensitivity to defeat (e.g., Goldman 1979). Simple creatures relying on a single method will not have such sensitivity.

Conversely, one might worry that the method coherence condition is too easy to achieve. Specifically, if it were possible for a method to count as “sufficiently coherent” when it has only been tested by a single other method, then the MCC would be worryingly permissive. However, I intend sufficient coherence to be more stringent than this, as my appeal to the notion of “holistic” coherence above is meant to indicate. A subject’s methods must be tested against a number of others in order to achieve this holistic coherence. Precisely how many will depend on a number of factors: the kind of reliability requirement one prefers (i.e., process, or competence), the answer to the generality problem, and perhaps the situation and context of the subject. Since I want to remain ecumenical about these topics, I will not endorse a particular answer here. Nonetheless, I do not think this objection poses a serious problem for the view.

---

<sup>33</sup>I talk of more and less certainty here to avoid saying “infallible.” If one is not convinced by anti-luminosity arguments that we are fallible about even our own mental states (Williamson 2000), then this response should look even better. If we have luminous access to our own mental states, we should not expect that our fallible external-world knowledge will be gained in the same way that we infallibly come to know things about our own mental states.

Another objection concerns the Roxanne case. Above, I claim that it is a benefit of MCR that it avoids granting Roxanne justification even at the first step of her bootstrapping process: it denies that she is justified in believing that the tank is full. One might worry that this goes against the intuitions and commitments of reliabilists. Isn't this just the kind of basic belief reliabilists want to justify?

I want to grant that there are many cases which are similar to Roxanne's where the subject is justified (as I point out at the end of §3.2). However, Roxanne has been described as a particularly bad epistemic agent. She does not have the sensitivity to her other available methods which we expect an agent to have. That there is an expectation of heeding one's available reliable methods has been recognized by reliabilists as a way of responding to certain purported counter-examples (Goldman 1979, Goldman and Beddor 2015). MCR provides an account of one aspect of that expectation: Roxanne is using a method which she has not successfully tested for coherence with her other methods. Thus, I don't think reliabilists should be happy to treat Roxanne as justified in this case. The MCR appropriately allows us to avoid this result.

Another worry one might have is related to the last objection. This objection is that we do not really need a solution to the direct circularity problem, and that MC is thus less necessary than I suggest. Some reliabilists, including Van Cleve (2003) and Sosa (2009), have argued that the similarity the circular reasoning between the perceiver and the crystal gazer should not perturb the reliabilist. After all, the difference between the perceiver and the gazer is that perception really is reliable, whereas crystal gazing is not. Moreover, in light of the bootstrapping/easy knowledge problem, we might be stuck with skepticism if we don't allow this kind of circular reasoning (Van Cleve 2003). We should simply be happy to point to the fact that Al in the Crystal Ball Gazer case (from pg. 4) employs false premises in his reasoning: crystal gazing is not in fact reliable. This distinguishes Al and the perceiver adequately so that reliabilists should not be perturbed by the similar form of the arguments.

Nonetheless, I remain perturbed. I'm perturbed because, generally, we are able to distinguish good reasoning from bad, quite apart from the actual truth value of our premises. This is true even in cases where the quality of the reasoning does not depend merely on formal structure; i.e., we are generally good at distinguishing strong inductive inferences from weak ones. Al does not just have false premises, he seems to engage in bad reasoning. If reliabilism really does permit the same kind of reasoning for Al and for the perceiver, then it seems to permit bad reasoning. A solution to epistemic circularity that simply appeals to the truth of the premises fails account for the fact that Al's reasoning is bad, not just his beliefs.

The method coherence condition provides an explanation of how Al's belief-formation went wrong, without appeal simply to the falsity of premises in his reasoning. His reasoning is unacceptable, over and above the falsity of his premises, because the method he is using to evaluate itself is not method coherent. Suppose Al has a twin in another possible world, Al\*. Al\* lives in a world where crystal gazing is reliable. If Al\* forms his beliefs just like Al does



(and fails to adequately method coherence test the crystal ball by, e.g., only making vague predictions), he will still be doing something wrong.

Recall that the kind of reasoning in question is an inference we can express thusly: method  $x$  says that method  $x$  is reliable, therefore method  $x$  is reliable. What is wrong with this kind of inference? One clear problem is that it is prone to false positives. It can be employed with respect to an unreliable method, and unreliable methods (like AI's crystal gazing) can provide the result that they are reliable, a false conclusion. The MCC provides a principled way of limiting which methods should be trusted, and so helps avoid these false positives.

Importantly, the MCC limits such inferences in a way which does not require us to presuppose the truth or falsity of the conclusion of the reasoning in question. A restriction of the inference that instead appealed to the reliability of the methods would require us to presuppose the truth or falsity of the conclusion in question. Furthermore, the method coherence restriction seems more easily action-guiding for the subject. Method coherence does not require the subject have access, but it does seem more directly accessible than reliability.

Thus, I take it that we do need a solution for the direct circularity problem, and method coherence provides one that is acceptably non-circular and action-guiding.

Finally, one might worry about that a reliabilist who adds the MCC to her account of justification may no longer be able to provide a Moorean response to the skeptic (such as the one in Sosa 2009). The Moorean argument requires that the account allows for a subject to be justified in believing (e.g.,) "here is a hand" without presupposing the conclusion "I am not a hand-less brain in a vat." Since the MCC requires more than mere reliability, one might worry that we cannot obtain this result. However, MCC can still deliver this result. Satisfying the method coherence condition does not require presupposing that one is not brain in a vat (in fact, a brain in a vat can satisfy the MCC). All that is required is that one successfully test one's relevant perceptual belief-forming method before deploying it to form the belief "here is a hand." This requires no presupposition that one is not a brain in a vat.

## 5 Conclusion

A reliabilism incorporating a method coherence condition will be better able to meet the direct circularity objections raised by Fumerton. Method coherence also accounts for the illegitimacy of bootstrapping. Furthermore, method coherence respects our intuitions about the importance of the way we form beliefs, and it fits with our everyday and scientific practices of justification. The addition of the method coherence condition offers us a way to capture the key advantages of reliabilism, while avoiding objections, in a manner that coheres well with our intuitive notion of appropriate epistemic behavior. Arguing that this solution to epistemic circularity problems is superior to all others on offer must await future work.

**Acknowledgments** For a great deal of helpful feedback, I would like to thank Walter Ott, Bill FitzPatrick, Steve Daskal, Ted Parent, Ernest Sosa, Branden Fitelson, Susanna Schellenberg, Andy Egan, Georgi Gardiner, David Black, Lisa Miracchi, Bob Beddor, and Megan Feeney.

## References

- Armstrong, David M. (1973). *Belief, Truth and Knowledge*, Cambridge: Cambridge University Press.
- Bach, Kent. (1985). "A Rationale for Reliabilism." *The Monist*. 68, pp. 246-263.
- Barnett, David J. (2014). "What's the Matter with Epistemic Circularity." *Philosophical Studies*, 171 (2):177-205.
- Beddor, Bob (2015). "Process Reliabilism's Troubles with Defeat," *The Philosophical Quarterly*, 65(259): 145–159.
- Bergmann, Michael (2004). "Epistemic circularity: Malignant and Benign." *Philosophy and Phenomenological Research*, 69 (3): 709–727.
- Bishop, Michael A. (2010). "Why the Generality Problem is Everybody's Problem." *Philosophical Studies*, 151:285–298
- BonJour, Lawrence (1980). "Externalist Theories of Empirical Knowledge." *Midwest Studies in Philosophy 5: Studies in Epistemology*. Ed. By P.A. French, et al. Minneapolis: University of Minnesota Press, Pg 53–73.
- BonJour, Laurence (1985). *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.
- Cohen, Stewart (1984). "Justification and Truth," *Philosophical Studies*, 46: 279–296.
- Cohen, Stewart (2002). "Basic Knowledge and the Problem of Easy Knowledge," *Philosophy and Phenomenological Research*, 65(2): 309–329.
- Comesaña (2005) "A Well-Founded Solution to the Generality Problem," *Philosophical Studies*, 129(1): 27–47.
- Conee, Earl and Richard Feldman, (1998) "The Generality Problem for Reliabilism," *Philosophical Studies*, 89(1): 1–29.
- Conee, Earl and Feldman, Richard (2001). "Internalism Defended," *American Philosophical Quarterly*, 38(1): 1?18.
- Douven, Igor, and Kelp, Christoph. (2013). "Proper Bootstrapping," *Synthese* 190: 171–185.

- Dretske, Fred (1981). *Knowledge and the Flow of Information*, Cambridge, MA: MIT Press.
- Earman, John (1992). *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. Cambridge, MA: MIT Press.
- Elgin, Catherine Z. (2005). "Non-Foundationalist Epistemology: Holism, Coherence, and Tenability." *Contemporary Debates in Epistemology* ed. Matthias Steup and Ernest Sosa. Boston: Blackwell, pp. 156–167.
- Fitelson, Branden (2001). *Studies in Bayesian Confirmation Theory*. Ph.D. Thesis, University of Wisconsin-Madison.
- Fitelson, Branden (2003) "A Probabilistic Measure of Coherence," *Analysis*, 63: 194–199.
- Fumerton, Richard (1995). *Metaepistemology and Skepticism*. Lanham, MD: Rowman and Littlefield,
- Goldman, Alvin (1979). "What is Justified Belief?" G.S. Pappas (ed.), *Justification and Knowledge*. Dordrecht: D. Reidel.
- Goldman, Alvin and Beddor, Bob, (2015). "Reliabilist Epistemology," *The Stanford Encyclopedia of Philosophy* (Winter 2015 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/win2015/entries/reliabilism/>>.
- Greco, John (1993). "Virtues and Vices of Virtue Epistemology." *Canadian Journal of Philosophy*. Vol 23, 3.
- Hájek, Alan (2007). "The Reference Class Problem is Your Problem Too." *Synthese* 156 (3):563–585.
- Kemeny, John G. and Oppenheim, Paul (1952). "Degrees of Factual Support." *Philosophy of Science*; 19: 307–324.
- Klein, Peter., and Warfield, T. A., (1994) "What Price Coherence?" *Analysis*, 54: 129–132.
- Lehrer, Keith (1990). *Theory of Knowledge*, Boulder, CO: Westview Press.
- Lewis, C. I. (1946) *An Analysis of Knowledge and Valuation*, LaSalle: Open Court.
- Lyons, Jack (2009) *Perception and Basic Beliefs*. Oxford: Oxford University Press.
- Nagel, Jennifer (2014). "The Meanings of Metacognition." *Philosophy and Phenomenological Research* 89 (3): 710-718.
- Nozick, Robert (1981). *Philosophical Explanations*. Cambridge, MA: Harvard University Press.

- Olsson, Erik J. (2005) *Against Coherence: Truth, Probability, and Justification*, Oxford: Clarendon Press.
- Olsson, Erik (2014) "Coherentist Theories of Epistemic Justification," *The Stanford Encyclopedia of Philosophy* (Spring 2014 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/spr2014/entries/justep-coherence/>>.
- Olsson, E. J., and Schubert, S., (2007). "Reliability Conducive Measures of Coherence," *Synthese*, 157 (3): 297–308.
- Pappas, George (2014). "Internalist vs. Externalist Conceptions of Epistemic Justification," *The Stanford Encyclopedia of Philosophy* (Fall 2014 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/fall2014/entries/justep-intext/>>.
- Plantinga, Alvin (1993), *Warrant: The Current Debate*, New York: Oxford University Press.
- Pollock, John (1984). "Reliability and Justified Belief," *Canadian Journal of Philosophy*, 14: 103–114.
- Poston, Ted (2014). *Reason and Explanation*. Palgrave Macmillan UK.
- Proust, Joelle (2013). *The Philosophy of Metacognition: Mental Agency and Self-Awareness*. Oxford: Oxford University Press.
- Pryor, James (2000). "The Skeptic and the Dogmatist," *Noûs* 34:4, 517–549.
- Ramsey, Frank P. (1926) "Truth and Probability," in Ramsey, 1931, *The Foundations of Mathematics and other Logical Essays*, Ch. VII, p.156-198, edited by R.B. Braithwaite, London: Kegan, Paul, Trench, Trubner & Co., New York: Harcourt, Brace and Company.
- Sosa, Ernest (2007). *A Virtue Epistemology*, Vol 1 . Oxford: Clarendon Press,
- Sosa, Ernest (2009). *A Virtue Epistemology*, Vol 2 . Oxford: Clarendon Press.
- Ernest Sosa (2011). *Knowing Full Well*, Princeton: Princeton University Press.
- Stroud, Barry 1989. "Understanding Human Knowledge in General." *Knowledge and Skepticism*. Ed. by M. Clay and K. Lehrer. Boulder: Westview Press
- Thagard, P. (2000). *Coherence in Thought and Action*, Cambridge, Mass.: The MIT Press.
- Titelbaum, Michael. (2010). "Tell Me You Love Me: Bootstrapping, Externalism, and No-Lose Epistemology." *Philosophical Studies*, 149, (1): 119–134.
- Van Cleve, James. (2003) "Is Knowledge Easy—Or Impossible? Externalism as the Only Alternative to Skepticism." *The Skeptics: Contemporary Essays*, Ashgate, 45–60.

Van Cleve, James. (2011) "Sosa on Epistemic Circularity." *Philosophical Studies*, (2011) 153: 19–28.

Vogel, Jonathan (2000). "Reliabilism Leveled." *The Journal of Philosophy*, Vol. 97, No. 11 pp 602–623.

Weisberg, Jonathan (2010). "Bootstrapping in General." *Philosophy and Phenomenological Research* 81 (3): 525–548.

Weisberg, Jonathan (2012). "The Bootstrapping Problem." *Philosophy Compass* 7(9): 597–610.

Williamson, Timothy (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.

Wimsatt, William (2007). *Re-Engineering Philosophy for Limited Beings*. Cambridge, MA: Harvard University Press.