

# CONVERGENCE ARGUMENTS IN THE PHILOSOPHY OF SCIENCE

ROBERT HUDSON

*University of Saskatchewan, CANADA*

`robert.hudson@usask.ca`

<https://orcid.org/0000-0002-6713-1043>

---

**Abstract.** A convergence argument occurs when multiple people agree in their testimony about the occurrence of an empirical phenomenon, and then conclude that this phenomenon is real. It is often thought that convergence arguments provide special support for their conclusions. For example, in the philosophy of science, robustness arguments are convergence arguments where different experimental approaches display the same phenomenon, and in political science the Condorcet Jury Theorem is based on a convergence argument where majority decision making is claimed to point to true propositions. The plan of this paper is to assess the quality of convergence arguments. Two kinds of convergence arguments are identified: conditional independence convergence arguments and intersection of independent lies convergence arguments. Though both kinds of arguments are immensely popular amongst a variety of philosophers, my plan is to repudiate their epistemic significance, and instead argue that we should follow the strategy of coherence arguments, the format of which I discuss at the end of the paper.

**Keywords:** convergence arguments • conditional independence • intersection of independent lies • Wesley Salmon • coherence arguments

---

RECEIVED: 13/09/2024

REVISED: 22/09/2025

ACCEPTED: 27/02/2026

## 1. Introduction

Convergence arguments have a long history in philosophy and in human thinking generally. A typical convergence argument is one where, given the independent agreement of multiple witnesses about the occurrence of an event, the occurrence of this event is more likely, and so better confirmed, than it is relative to the individual testimonials themselves. For example, whereas an individual's report of a miracle does little to support its occurrence, a group consensus about a miracle provides comparatively strong support, assuming this consensus is independently arrived at by each individual consenter. In the philosophy of science, the term 'robust' is used to describe a convergence argument involving the results of multiple experimental strategies, where each of these strategies leads to the same experimental result, indicating that some physical phenomenon is real. Another example of a robustness argument



is the no miracles argument for scientific realism: the repeated, independent predictive successes of a theory demonstrate its truthfulness; such convergent successes are too unlikely, otherwise. In political science the Condorcet Jury Theorem involves a convergence argument where majority decision-making points to a true proposition, if individual voters act independently of one another (a phenomenon called ‘the wisdom of crowds’).

However, convergence arguments are not always viewed favourably. René Descartes is skeptical. He comments in *Discourse on the Method*:

although such be the ground of our opinions, . . . a plurality of suffrages is no guarantee of truth where it is at all of difficult discovery, as in such cases it is much more likely that it will be found by one than by many. (Descartes 1912 [1637], p.14)

With difficult topics, there are few experts — most people proclaim ignorance or consent unwittingly. So, if only, Descartes thinks, an expert could be found whose opinion he could trust. Finding none, he finds himself, as a methodological solipsist, “constrained, as it were, to use [his] own reason in the conduct of [his] life” (p.14). He only follows his individual judgment.

Bertrand Russell, too, expresses doubts about the value of convergence arguments. About a table,

to the eye it is oblong, brown and shiny, to the touch it is smooth and cool and hard; when I tap it, it gives out a wooden sound. Anyone else who sees and feels and hears the table will agree with this description, so that it might seem as if no difficulty would arise; but as soon as we try to be more precise our troubles begin. (1912, p.11)

The troubles Russell cites concern the variability of perceivers’ perspectives. Russell’s point is that, once we attend closely to our sense data — the ‘brown and shiny’, ‘the smooth and cool’, ‘the wooden sound’ — we realize that we see radically different tables, and our convergent agreement about the table’s properties is misleading.

Even more generally, in introductory critical thinking courses it is considered fallacious to blindly trust popular agreement about a topic – the so-called ‘bandwagon’ fallacy. Obviously, we cannot objectively adjudicate philosophical or scientific issues by means of democratic decision making, as though the production of knowledge could occur through a popular vote. It won’t even matter if we assume that voters stem from different, independent backgrounds. Their backgrounds can be as different as we like. If they lack knowledge, they lack knowledge, convergence aside.

There is, therefore, uncertainty about the value of convergence arguments. Contrary to the individualism of Descartes that prizes singular, expert approaches to gaining knowledge, and the private phenomenism of Russell which emphasizes the priority of personally held sense experiences, many current philosophers are socialist

in orientation, endorsing the value of multiple, diverse perspectives in the pursuit of knowledge. Helen Longino is representative of this emphasis on diversity. For her,

communities must themselves develop sites of critical discourse about their assumptions both through incorporating diverse perspectives within themselves as internal resources for criticism and by engaging with communities whose background is different. (1999, p.330)

Similarly, Elizabeth Anderson comments:

a scientific community composed of inquirers who share the same background assumptions is unlikely to be aware of the roles these assumptions play in licensing inferences from observations to hypothesis, and even less likely to examine these assumptions critically. Naturalized epistemology therefore recommends that the scientific community include a diversity of inquirers who accept different background assumptions. (1995 [2002], p.483)

Given Descartes' and Russell's demurrals about the value of divergent perspectives on issues of fact, the question raised by the viewpoints expressed by Longino and Anderson is, "What epistemic advantage is there in having multiple, diverse perspectives converge on a claim about the world?" Proponents of convergence arguments typically identify an epistemic advantage with multiple, diverse perspectives in terms of leading to the truth.

As such, the plan in this paper is to assess the quality of convergence arguments, of which two kinds are identified: first, conditional independence convergence arguments and secondly, intersection of independent lines convergence arguments. Though both kinds of arguments are immensely popular amongst a variety of philosophers, this paper repudiates their epistemic significance. Alternatively, it is suggested that one should advance coherence arguments, and I close the paper by describing how coherence arguments work and how they bypass the various problems cited with convergence arguments.

## **2. Conditional independence convergence arguments**

Let's start by clarifying the notion of a convergence argument. A convergence argument is an argument of the following form. Given procedure 1, claim A follows; given procedure 2, claim A follows; given procedure 3, claim A follows; and so on. Therefore, A is (likely) true. Here, what counts as a procedure is variable: it could be an inference, an experiment, the recording of a person's vote, the retrieval of a witness report, and other things. Whichever thing it is, it results in the production of A, a claim, that is asserted to be true. With a convergence argument, then, we find that a claim A is asserted to be true at the termination of multiple investigative procedures.

For a convergence argument to work, there are constraints on procedures that must be in place. First and foremost, the procedures must be different from one another; otherwise, it's just repetition, repeating the same process again and again with no informative value. Put another way, the procedures can't be exactly the same, that is, numerically identical. They can, nevertheless, be qualitatively the same, or the same kind of procedure. For example, two witness reports could serve in a convergence argument, or two instances of the same experimental procedure performed at different times. This is typically assumed to be acceptable so long as the kindred procedures are probabilistically independent of one another, in addition to being causally independent. To facilitate this requirement, a distinction is drawn between two kinds of probabilistic independence. Let  $A_i$  and  $A_j$  be two versions of a procedure that generates  $A$ .  $P(A_i)$  denotes the probability that procedure  $i$  generates claim  $A$ . We assume  $P(A/A_i) > P(A)$  — i.e., procedure  $i$  in generating  $A$  provides some reason to believe that  $A$  is true —, and  $P(A/A_j) > P(A)$  for the same reason. We further assume  $P(A_j/A_i) > P(A_i)$ : learning that procedure  $i$  generates claim  $A$  increases the chance that procedure  $j$  generates  $A$ , as well. After all, both procedures are designed to track  $A$ , so if one procedure generates  $A$  (i.e.,  $A_i$ ), one anticipates that the other procedure will also generate  $A$  (i.e.,  $A_j$ ). In this sort of case  $A_i$  and  $A_j$  are probabilistically dependent on one another unconditionally. This does not preclude the case where, even though  $A_i$  and  $A_j$  are probabilistically dependent on one another unconditionally,  $A_i$  and  $A_j$  are probabilistically independent of one another *conditionally*, that is, conditional on the truth of  $A$ : formally,  $P(A_j/A_i \text{ and } A) = P(A_j/A)$ , and vice versa,  $P(A_i/A_j \text{ and } A) = P(A_i/A)$ . In these cases, the truth of  $A$  screens off the probabilistic relevance obtaining between  $A_i$  to  $A_j$  (here, the term 'screen off' is used in the sense of Reichenbach 1956, p.189; see also Fitelson 2001, p.S128).

Accordingly, the strategy in probabilistic versions of convergence arguments based on the independence of the convergent procedures has been to assume the conditional, probabilistic independence of these procedures, not their unconditional independence. For example, Van Cleve (2011) in reviewing the history of convergence arguments identifies this conditional independence requirement as occurring in Keynes (2007 [1921]), Cohen (1976), Jeffery (1987), Huemer (1997), Blitstein (2001) and Olsson and Shogenji (2004). It is also fundamental to the arguments given in support of Condorcet's Jury Theorem, as seen in Romeijn and Atkinson (2011) and Dietrich and Spiekermann (2013a) and (2013b), and to the convergence arguments found in Salmon (1984) and Sober (2008).

In utilizing a conditional independence convergence argument, one needs to determine how likely is it that claim  $A$  will be produced with procedure  $j$ , given that it had already been produced with a different procedure  $i$ . If conditional independence holds, the likelihood of claim  $A$  given procedure  $j$ , given that  $A$  has already been produced by procedure  $i$ , is determined by whatever causal relation exists between

the situation described by  $A$  and the details of procedure  $j$ , a causal relation that dictates how reliable procedure  $j$  is in detecting  $A$  when  $A$  is the case. So, in effect, with conditional independence, we can just ignore the previous procedure  $i$  and the fact that it generated  $A$  and focus instead on procedure  $j$  and the probability with which it correctly detects the truth of  $A$ . Then, so long as our procedures are minimally reliable in that they generate claim  $A$ , when  $A$  is true, at a rate that is better than chance, it follows that when procedure  $j$  produces  $A$  we have further support for  $A$  than what is provided by procedure  $i$  alone, and that this support increases with every new procedure meeting the same conditions as procedure  $j$  (and  $i$ ).

As previously indicated, for this argument to work, the two procedures must be different from one another. For instance, if procedure  $j$  simply involves someone recalling that procedure  $i$  generates claim  $A$ , this does nothing to further support the truth of  $A$ , as essentially the same procedure is being used. Still, the two procedures needn't be too different from one another. For instance, let procedure  $j$  involve performing procedure  $i$  again at a later time. Doing this provides no certain guarantee that claim  $A$  will be generated again, not like when we made use of a numerically identical version of procedure  $i$ , since after an elapse of time the circumstances are slightly different, so there is a slight possibility that claim  $A$  won't be produced. So assume, once more, that conditional independence holds, and  $i$  generates  $A_i$ . It follows that the likelihood of claim  $A$  given a later, slightly different procedure  $j$ , given that  $A$  has already been produced by the earlier procedure  $i$ , is determined but whatever causal relation exists between the situation described by  $A$  and the details of the slightly different procedure  $j$ . This means we can just ignore the previous procedure  $i$  and the fact that it generated  $A$ , and focus instead on the slightly different procedure  $j$  and the probability with which it correctly detects the truth of  $A$  (which may be slightly different from the probability with which the earlier procedure  $i$  correctly detects the truth of  $A$ ). Moreover, we can perform this maneuver repeatedly, using yet another slightly different procedure, over and over again. In this way we continually, incrementally increase the support for claim  $A$ , once more under the assumption that the iterated procedures are minimally reliable. It follows that, in time, with such trivial modifications, support for  $A$  asymptotically and automatically reaches its highest limit, at some value below certainty.

This is a reduction to absurdity with conditional independence arguments. Given a procedure  $i$  generating  $A_i$ , one can arbitrarily invent a new, conditionally independent procedure  $j$  generating  $A_j$  that (along with  $A_i$ ) supports  $A$  to a greater degree than procedure  $i$  alone. That is because, to ensure conditional independence, all one needs do is assume  $A$ . If one assumes  $A$ , then any experimental procedure  $i$  generating  $A_i$  is conditionally independent of any other experimental procedure  $j$  generating  $A_j$ , no matter how trivial the difference between the two procedures. Suppose, for example, an experimental procedure  $i$  generates  $A_i$  in support of  $A$ , and one

is now looking for another experimental procedure  $j$  that generates  $A_j$  to provide enhanced support for  $A$ . Normally, finding this alternative experimental procedure involves detailed, scientific scrutiny — not just any experimental procedure can be used. However, given the conditional independence convergence argument, nothing special or elaborate terms of a detailed scientific scrutiny is required. This is because the original experiment and its trivial modification are conditionally independent of one another, given  $A$ . If  $A$  is the case, then the probability that a trivially modified experiment supports  $A$  is unaffected by whether the original experiment supports  $A$ , or not, or even on whether the original experiment ever took place. The trivially modified procedure generates the whole richness of support the original experiment had provided for  $A$ . By simply repeating a minimally reliable experimental procedure and trivially modifying it (for example, by only repeating it another time), we have a strategy for generating the highest probabilistic value  $A$  can have. If we know ahead of time that the evidence in favor of a hypothesis cannot go beyond a certain limit, then by repeating trivial modifications of an experiment supportive of this hypothesis one generates support for this hypothesis asymptotically approaching this limit.

The absurdity is that strong scientific support is not trivially generated by such a means. Scientists focus attentively and critically on the experimental procedures used in generating scientific support. There may be multiple minimally reliable experimental procedures available (including trivial ones). A scientist sometimes chooses experimental procedures for reasons other than (and additional to) minimal reliability. For example, an experimental procedure might be chosen because it deals with the particular subject matter of relevance to a theory. So what procedure is chosen is not a trivial matter. The conditional independence convergence argument fails to recognize this key, epistemic consideration.

Another way to see the given point is to look at it from a Bayesian perspective. Let the hypothesis under consideration be the truth of  $A$ . Evidence for the truth of  $A$  is  $A_i$ , the fact that procedure  $i$  generates  $A$ . So long as the prior probability of  $A_i$  is less than 1, it follows by Bayes theorem that  $A_i$  confirms the truth of  $A$  under appropriate conditions. That is, the evidence does not become ‘old’, which is the case where we use a numerically identical piece of evidence. This ‘does not become old’ condition is satisfied if we are using subsequent iterations of  $A_i$ , subsequent cases in which procedure  $i$  generates  $A$ , since these new iterations are, at least, slightly different from the original  $A_i$ , and the production of  $A$  is not thereby guaranteed.

Unfortunately, this result is problematic for conditional independence convergence arguments. For we can take a procedure and simply modify it slightly time and again, and still get the noted, valuable result of a conditional independence convergence argument, that is, the ability to incrementally increase the confirmatory support for a claim all the way to its prescribed probabilistic limit. This process may take a longer time than if the varying procedures were more different, for then the con-

firmatory results are less likely to be expected. Still, the underlying logic is the same, and so we have a conflict with the motivating idea behind convergence arguments — that, somehow, using a diverse set of procedures has a special value in establishing the truth of a claim. This does not seem to follow if minimally, hardly-at-all diverse procedures can have the same effect.

One diagnosis for why conditional independence convergence arguments lead to this problem is that they don't account for the presence of common causes unrelated to the truth of the claim under consideration that could produce convergent results. Again, take procedure *i* that generates claim *A*, and run it again at a later time so that it is slightly different from before. The later iteration of procedure *i* will not be much different; the two versions of procedure *i* will share a number of common elements, elements unrelated to its ability to reliably indicate the truth of *A*, and these common elements could be responsible for the production of claim *A*. Accordingly, to avoid this possibility, one should compare the results of procedures that are a lot different from one another, sharing few common elements.

Ideally, there is a way to represent this diversity requirement probabilistically. An attempt to do this is found in Dietrich and Spiekermann (2013a) and (2013b), who formulate a new kind of independence requirement conditionalizing over true states of affairs, as with conditional independence, and conditionalizing over common causes more generally in procedures that lead to convergent results. In this case a convergence argument can work, once again only if it is true that the procedures are accurate at a rate better than chance (that is, in the language of the Condorcet Jury Theorem, with which Dietrich and Spiekermann are working, the voters are 'competent'). Whether this is true is an open question if there are all sorts of common causes impacting procedures that are unrelated to the truth of *A*. Dietrich and Spiekermann acknowledge the problem here for their approach. They acknowledge that the relevant convergence argument will sometimes fail and occasionally lead to the "counterintuitive finding that large groups are worse [i.e., less reliable] than single individuals" (2013a, p.102).

One way to express the problem cited by Dietrich and Spiekermann is to note that two procedures may lead to the same result due to common factors unrelated to the theoretical significance of these results, and with respect to which they are trivial. For example, the results may both be written in blue ink, so have a common cause in this respect. Generalizing over such a fact, despite its indication of a common cause, has no significance for scientific inquiry as such inquiry focuses on the meaning of results, not on how these results are delivered, such as linguistically, by means of blue or some other colour ink. Convergence arguments, if they fail to distinguish between truth-related and non-truth-related common causes, will fail to capture the distinctive epistemic significance of diverse evidential sources. (Later on, I describe an approach by Wesley Salmon addressing this problem.)

Having contested the value of a conditional independence convergence arguments, let us examine a different kind of convergence argument, the intersection of independent lies convergence argument.

### 3. Intersection of independent lies convergence arguments

Suppose procedure 1 generates result  $R$  using modeling assumptions  $a$ ,  $b$ ,  $c$  and  $x$ ; procedure 2 generates result  $R$  using independent modeling assumptions  $d$ ,  $e$ ,  $f$ , and again  $x$ ; finally, procedure 3 generates result  $R$  using further different modelling assumptions  $g$ ,  $h$ ,  $i$ , and once more  $x$ . How is it that these procedures are able to generate the same result? These procedures have no modelling assumptions in common that would explain their convergent results, other than  $x$ . There is no requirement, in generating result  $R$ , that one maintain any of the independent assumptions, other than  $x$ . Thus, the non- $x$  assumptions could be falsehoods (“lies”) and still one can generate result  $R$ , so long as modeling assumption  $x$  is true.

This popular argument, involving modeling, is given succinct expression by Richard Levins. When using “the mathematical models of population biology”, he suggests “we attempt to treat the same problem with several alternative models [ $a$ ,  $b$ , ...] each with different simplifications but with a common biological assumption [ $x$ ]”. Now suppose “these models, despite their different assumptions, lead to similar results”. We have with  $x$  “a robust theorem ... relatively free of the details of the model”. In Levins’ catchy phrase, “our truth is the intersection of independent lies” (1966, p.423), where the independent lies are the idiosyncratic modeling assumptions contained in each of the procedures (for Levins, the different simplifications implicit in alternate models). These extra assumptions are false explanations for the production of  $R$ , since they are individually not needed to generate  $R$ . They are, as it were, washed out as explanations for the production of  $R$  since alternate procedures are able to generate  $R$  in circumstances where these assumptions fail to apply. So we conclude that  $x$  is the true modeling assumption that explains result  $R$ . This is the intersection of independent lies convergence argument, where the independent lies are the unneeded modelling assumptions,  $a$ ,  $b$ ,  $c$ ,  $d$ , and so on, leaving  $x$ , which is inferred to be true.

Is this an effective argument for the truth of a modelling assumption, such as  $x$ ? Here is an argument disputing its effectiveness. Suppose it is the 16<sup>th</sup> century, and we are interviewing people on the street about the reasons why they believe in geocentrism. In effect, we are cataloging procedures that generate the claim that geocentrism is true. Suppose the first person we interview says that their belief in geocentrism is caused by their faith in the views of the Catholic Church; a second person says that their belief in geocentrism is due to their disbelief in the arguments

provided by Copernicus for a Sun-centered system — they are in no way swayed by the Church; a third person has no interest in the Church, has never read Copernicus, yet has read Ptolemy and found his views convincing. Here are three procedures that independently generate the claim the geocentrism is true, where the three interviewees utilize different modeling assumptions — the model of Church doctrine, the model of Copernicus' blunder, and the model of a convincing Ptolemy –, each of them generating the result that geocentrism is endorsed. What explains this convergence? There is a modeling assumption each of the subjects shares that washes out their differences, that is the intersection of their independent lies. This modeling assumption is geocentrism, which we then infer to be true, by the intersection of independent lies convergence argument.

The inference is problematic. Of course, geocentrism is false, and the intersection of independent lies convergence argument readily leads, under imaginable circumstances, to a false claim. How can a proponent of this kind of convergence argument respond?

One might object that the cited procedures are not very reliable. Relying on the Catholic Church for instructions on scientific astronomy, not appreciating Copernicus' solid arguments for heliocentrism in *On the Revolutions of the Heavenly Spheres*, and studying Ptolemy's geocentrism, are not reliable passages to the truth; and, as with conditional independence convergence arguments, we should assume that the procedures at issue lead to true beliefs at a rate better than chance. Now this is definitely a point worth noting. Convergence arguments only work if we assume that the constituent procedures have at least minimal reliability. On this issue Dietrich and Spiekermann comment, "the Classical Competence assumption [in arguments for the Condorcet Jury Theorem] is not unreasonable; in fact it is plausibly true. It tells us that the voters are more likely to vote correctly than incorrectly" (2013a, p.98). This will not be the case where we are considering the epistemic reliability of voters on issues of any difficulty, as Descartes recognizes. Moreover, the reliability of procedures is perhaps something that should be evaluated after the empirical results have come in, not assumed ahead of time. The geocentrism case could then be an example of what happens if we aren't careful in this respect. Also, it won't do to conclude, for procedures converging in their results, that these procedures are reliable and the results likely to be true. Rather, this convergence should simply be used as one piece of evidence amongst a large collection of other pieces of evidence to prove a point, such as, in the geocentric case, the point that an Earth-centered universe best explains the apparent motions of the planets and the sun. In short, the fact that our three 16th century interviewees provide convergent support for geocentrism using independent procedures is a distinct justification of geocentrism. On the other hand, this convergence lacks the special significance proponents of convergence argumentation ascribe to it. There is no a priori, purely logical understanding of the value of a

convergence argument if, logically, it is possible for independently convergent beliefs to land on a falsehood.

Then again, are the three geocentrists all that unreliable? Taking a historically sensitive view of the matter, the first geocentrist is paying respect to the dominant intellectual tradition of the day, for which there is broad public and expert support; the second takes a rational, critical eye about a controversial work of astronomy; and the third studies a work of astronomy with a long history of intellectual acceptance. These procedures could, in general terms, be said to deliver truth at a rate better than chance. We are then returned to the teeth of our objection, that convergence arguments might well lead to falsehoods, and so are not reliable indicators of truth, even when the procedures themselves are reliable.

In defense of an intersection of independent lies convergence argument, one might object, alternatively, that the arguments of our geocentrists are not *causally* independent of one another, even if they are probabilistically independent. The interviewees live in a time and place dominated by the intellectual influence of the Catholic Church. Though the geocentrists, we can suppose, do not causally influence one another directly, there is a common cause to their arguments — the doctrinal geocentrism of the Church. In the way the case is set up, however, this is not a factor. The latter two geocentrists, for their reasonings to be compelling, need not be swayed by the Church; one can critique Copernicus' reasonings and endorse Ptolemy's arguments, independently of the Church's influence. So there is no reason to suspect a Church-induced common cause.

Let us suppose, nonetheless, that there is a Church-induced common cause underlying their arguments. We infer that the three procedures lack independence and so form no counterexample to either conditional or intersection of independent lies convergence arguments. Now, let us change the case to one where we are today interviewing people on the street about the reasons why they believe in heliocentrism. Similar to the above, we catalog procedures that generate the claim that heliocentrism is true. Moreover, we assume a case similar to the Church-dominated 16<sup>th</sup> century, this time assuming, instead, the current intellectual hegemony of modern science with its heliocentric worldview. In our interviews, we learn that one person supports their belief in heliocentrism based upon their appreciation of Copernicus' treatise. Another is compelled by Galileo's telescopic observations, and a third is convinced by the expert advice of their university's renowned astronomy professor. If each of the interviewees is unaware of the others' considered support for heliocentrism, their views are to that extent causally independent of each other. Of course we live today in an intellectual climate where heliocentrism is taken as undoubtedly true, just as people in the 16<sup>th</sup> century lived in an intellectual climate where geocentrism was taken as undoubtedly true. Also, we can suppose this intellectual climate to be a common cause of the arguments provided by the three interviewees on behalf

of heliocentrism, just as we assume to be true in the geocentric case. The problem is that, in the heliocentrism case, the relevant convergence argument should work. It is compelling for a belief in heliocentrism that the interviewees cite the authority of three distinct, reliable individuals acting independently of one another – Copernicus, Galileo and the renowned university professor. Thus, the relevant convergence argument gets the right answer, albeit for the wrong reason. The agreement compels a belief in heliocentrism not because of the bare agreement of reliable authorities. The agreement compels because of the status of the three authorities as highly reliable reporters on astronomy. The agreement is a successful replication of expert affirmations of heliocentrism.

In both the geocentrism and heliocentrism cases, the common cause cited, either the Church or hegemonic science as a societal force, needs to be shown to exist and not just assumed to exist. People in expressing their considered, rationally justified opinions are autonomous enough to override the persuasive pressure of a widely-held belief, whether about God, heliocentrism or whatever. They are able to decide for themselves whether the evidence for heliocentrism is compelling, its status as a widely-held belief notwithstanding. And they were the same way in the 16<sup>th</sup> century as regards their considered, rationally justified opinions about geocentrism. From their perspective, one could formulate good arguments for geocentrism, adjusting for the state of the available evidence. People were also manifestly aware of the intellectual pressure imposed by the Church (as Copernicus was), allowing them to discount that pressure in their thinking. So, to write off an intersection of independent lies convergence argument for the truth of geocentrism because of the presence a Church-imposed common cause, risks writing off similar convergence arguments that we should be comfortable with (if we believe in convergence arguments to begin with), such as the convergence argument cited in support of heliocentrism.

To summarize, it has been argued that an intersection of independent lies convergence argument can feasibly be used to support a geocentric viewpoint, or indeed any false claim. Investigative procedures are fallible enough to produce false claims from a variety of angles. Defenders of such an argument, I have suggested, might respond by attributing a convergence on false claims as due to the presence of a non-truth-related, common cause, such as the Church's 16<sup>th</sup> century intellectual dominance in the geocentrism case. The problem is that such non-truth-related common causes could be present with a legitimately compelling convergence on a true claim, such as with modern heliocentrism.

#### 4. Wesley Salmon's truth-related common causes

Nevertheless, there are some compelling convergence arguments involving truth-related common causes for widely-held claims, and where only the truth of the claim, as generated by idiosyncratic reasoning procedures, explains the convergence of beliefs. This sort of case is discussed by Wesley Salmon in *Scientific Explanation and the Causal Structure of the World*. Salmon describes a situation where several independent witnesses at a crime scene concur that John Doe was nearby when the crime was committed, and concur further on what John Doe was wearing, what he was carrying, that he had arrived in a cab displaying a particular license number, and so on (see Salmon 1984, p.220). Such concurrence, Salmon believes, provides strong evidence that their reports, as convergent, are correct. In fact, he continues, "even if none of the witnesses could be considered particularly reliable, the agreement among their reports would greatly enhance our confidence in the veracity of their account" (p.220).

To clarify Salmon's argument, let us consider a case where ten witnesses arrive at a consensus about the license plate number of the cab containing John Doe, a plate that contains a random mix of six alphanumeric characters. That an individual witness correctly identifies, purely by chance, and in the right order, all six characters of the license plate, is enormously small. That all ten witnesses converge in their views, and correctly identify, purely by chance, and in the right order, all six characters of the license plate, is enormously smaller. It is so small that, in applying a convergence argument, one does not need to assume that a procedure gets the right answer at a rate better than chance. Even if a witness is hardly reliable, such as in guessing license plate numbers, their possession of a true belief needs an explanation. That a set of unreliable witnesses comes to the same agreement needs even further explanation. With a convergence argument, this explanation is the *truth* of the belief, either upon which we conditionalize or as the intersection of a set of independent beliefs.

As per the latter, let us assume the absence of any obvious common cause for their convergent claims, such as someone informing them in advance of the license plate number. Assume further that their visual systems are working normally and that the viewing conditions for the license plate number are optimal. The witnesses are not subject to psychological, hallucinatory interventions, perpetrated by manipulative scientists. Each of them, moreover, has their own unique private causes or procedures for their claims, such as seeing the license plate at a certain angle, with a certain pattern of colour. These private causes will differ for each witness, and so cannot be used to explain their agreement. So surely, one might argue, the agreement of the witnesses about the license plate number can only be explained in one way, by the *truth* of their beliefs, the fact that each of them saw the same number. This truth is the intersection of their independent lies (i.e., the private causes of their beliefs,

generated by idiosyncratic procedures). Despite their varying misapprehensions, they agree on the license plate number, and the explanation for this is that the number is correct, a feature of the real world.

As per the former, involving conditional independence, we assume the unconditional dependence of witness reports  $A_j$  and  $A_i$ :  $P(A_j/A_i) > P(A_j)$ . Witness  $j$ 's report ( $A_j$ ) is confirmed by witness  $i$ 's report ( $A_i$ ) since witnesses are chosen for their ability to testify about the license plate number, and if witness  $i$  reports  $A_i$ , this increases, if only slightly, the probability that witness  $j$  reports  $A_j$ :  $P(A_j/A_i) > P(A_j)$  holds, unconditionally. One can reconstruct this evidential relation by assuming it contains two parts, both of which are reasonable to assume, and from which the unconditional dependence of  $A_j$  on  $A_i$  follows by means of the transitivity of probabilistic support, given a screen-off condition. Where  $A$  is the true state of the license plate,

- 1)  $P(A/A_i) > P(A)$
- 2)  $P(A_j/A) > P(A_j)$ .

In other words, testimony  $A_i$  lends support for  $A$  (as the true state of world), and where  $A$  is the true state of the world, this lends support for  $A_j$ .  $A_i$  makes  $A$  more likely to be true, and  $A$ 's truth makes it more likely to expect  $A_j$ . The unconditional dependence of  $A_j$  on  $A_i$  then follows transitively from 1) and 2) if we assume that  $A_i$  and  $A_j$  are screened off from one another by  $A$  in the following sense:

$$P(A_j/A_i \text{ and } A) = P(A_j/A) \text{ and } P(A_j/A_i \text{ and not-}A) = P(A_j/\text{not-}A).$$

This feature of probabilistic evidence, demonstrated by Tomoji Shogenji (2003, p.614) and accordingly labelled 'Shogenji's theorem', is a feature of convergence arguments involving a series of evidential judgements, as above: an evidential chain from  $A_i$  to  $A$  followed by a chain from  $A$  to  $A_j$ . These chains lead to the unconditional dependence of  $A_j$  on  $A_i$  if we assume the truth of  $A$ , where  $A$  screens off  $A_i$  and  $A_j$  from each other in Shogenji's sense. Accordingly,  $A_i$ 's evidential support for  $A_j$  is explained by the truth of  $A$ , where  $A$  is a common cause for both  $A_i$  to  $A_j$  and where  $A$  explains the (unconditional) evidential relevance of  $A_i$  to  $A_j$ .

The argument by Salmon for the claim that the truth of  $A$  best explains the evidential relevance of  $A_i$  to  $A_j$  (and vice versa) depends on what other common causes need to be in place for  $A_j$  to converge with  $A_i$ . For example, the witnesses' visual systems are assumed to work in similar ways. To see the same numbers on the license plate, their visual physiology and related cognitive processing work in a similar manner. The numbers appear to them in similar ways, and they attach similar concepts to numeric and alphabetic symbols. When witnesses arrive at the same view about the license plate number, such physiological and conceptual factors are common causes of their arriving at the same view, causes additional to the truth of their reports (expressed conditionally as  $A$ ).

For example, the relevant symbols mean the same things for the witnesses. The symbol ‘6’ means the number 6, the symbol ‘5’ means 5, and so on. If, for example, the symbol ‘6’ means for some witness the number 5 (the witness utters “6” upon seeing 5) and for another witness ‘5’ means the number 5 (the witness utters “5” upon seeing 5), then agreement upon seeing 5 (i.e., both uttering “5” upon seeing 5) is impossible for them. Similarly, for such agreement, they need to have similar ideas about the nature of the setting. If one witness thinks they are in a deceptive context and another witness thinks their situation is perfectly normal, then once more their agreement is hard to achieve. So a number of common causes, additional to the objective nature of the license plate, must be in place for there to be agreement amongst the witnesses about the license plate number.

Now imagine a situation which is deceptive. Suppose, for example, the license plate being witnessed is a trick license plate, one that exhibits different alphanumeric characters at the behest of the driver. The license plate has a set of real characters, and the driver has the ability to modify how these characters appear, perhaps as a way to evade capture by the police. So suppose the driver initiates the trick license plate, hiding its true value. Again, all the witnesses agree about the alphanumeric characters on the license plate, and do so because of the common causes we have cited — except for the fact that they are mistaken about what the real license plate characters are, objectively speaking. So we can explain the convergence of their views about the license plate without making any claims about the true state of the license plate, thereby rebutting Salmon’s argument. The truth of *A* is not assumed in explaining the convergence of their reports, not even conditionally. Instead, we explain this convergence by means of the other common causes cited above, i.e., having similar visual systems and similar ideas about the nature of the visual setting (either normal or deceptive). These other conditions are comparatively difficult to dispense with when explaining the convergence of their reports.

To save Salmon’s argument, one can hedge on the truth that underlies conditionally the convergence of the witnesses’ opinions (or is the intersection of their opinions). Originally, the convergence of their views on the license plate is conditional on the objective truth that the license plate number contains a particular configuration of alphanumeric characters. We now see that this claim is difficult to establish, given trick license plates and other deceptive factors. So, a weaker claim is established: the witnesses converge on what they subjectively think the license plate number to be, not what it actually is. Given their convergent reports, we conclude that they all have the same perceptual belief about the plate’s particular configuration of alphanumeric characters, leaving aside whether this particular perceived configuration is accurate, or not. They have this same belief, despite their beliefs having been generated by different procedures, different private causes. One witness sees the characters on the license plate with a certain shading of color, whereas a different (agreeing) witness

sees a different shading. One witness, because they're a little farther away, sees the alphanumeric characters as quite small, whereas a different witness closer to the license plate sees the characters as proportionately bigger. Despite these differences, the witnesses nevertheless arrive at the same belief about the license plate number (again, accurate or not). Of course, arguing in this manner is far too weak to give a convergence argument much value. It is similar to the geocentrism case. People in the 16<sup>th</sup> century arrived at convergent beliefs about the truth of geocentrism. Some were convinced by the teachings of the Church, others by their dissatisfaction with heliocentric arguments, still others by their reading of Ptolemy. Overall, their convergent views are a product of their common awareness of a universe perceptually and cognitively configured in a certain way. Without this shared subjective perspective there would be no convergence of astronomical views; with this shared subjective perspective, their convergence of views is achieved, the falsity of their views notwithstanding.

Our assessment of conditional independence and intersection of independent lies convergence arguments can be summarized in the following way. With the former, since conditional independence is achieved using new procedures that are only slightly different from an original procedure, it follows that we have little assurance regarding the accuracy of the result. Alternate procedures only slightly different from the original procedure can be harnessed in support of a conclusion. To be clear, this is not to say that there is no benefit whatsoever in generating results with procedures that are only slightly different. It is the point of experimental replication to ensure that a result comes about through the use of a procedure that is only slightly different from how it was used initially; one attempts to repeat the original procedure in as identical a fashion as possible. Replication in itself, though, does not guarantee the truth of the result. What assures us of the truth of the result is the quality of the experimental procedure, that is, how reliable it is at discerning the phenomenon under consideration. A successful replication simply indicates that the result is, indeed, what the original application of the procedure had previously shown.

Similarly, with intersection of independent lies convergence arguments, there is no guarantee that a claim is nearer to the truth if it is arrived at by multiple procedures than by a singular procedure. A straightforward way to make this point is to consider a case where procedures generate claims informed by misleading concepts. For example, consider a set of procedures used to identify instances of a fictional kind of animal, a kind composed of horses and dogs. A procedure generates an affirmative claim that an object is a member of this kind if the object meets the criterion of being either a dog or a horse. Here, for some object, we might find that each member of this set of horse-dog procedures generates an affirmative claim when presented with this object, despite the presence in individual procedures of idiosyncratic, private causes. Some procedures identify members of the kind by reference to genetic factors (ei-

ther horse or dog genotypes); others by stereotypic, observable features of dogs and horses; others according to the counsel of scientific experts. In explaining this surprising convergence of opinions about an object one needs, as suggested above for the case of human, visual witnesses of a license plate number, the presence of common causes corresponding to shared visual and cognitive abilities, such as similar conceptualizations of what constitutes a dog or a horse, and a shared commitment to either the normalcy or deviance of the perceptual situation. These common causes are necessary if we are to explain a convergence of results. What is not necessary, *contra* Salmon, is for these procedures to generate true claims. Obviously, in the horse-dog case, the categorization leads to a convergence of claims that are all false, as by hypothesis dogs and horses do not form a single kind. Such falsehoods pose no obstacle to procedures generating the same result, with each of them doing so in diverse ways, given the satisfaction of certain common causes.

How can a defender of the epistemic value of convergence arguments (such as Salmon) respond to the horse-dog case? Recall the requirement in conditional independence convergence arguments that reporting procedures have a better than chance likelihood of generating true results. Clearly, in the dog-horse case, this requirement is violated: none of the procedures generate true claims as they all refer to a fictional kind. And, generally speaking, our argument against conditional independence convergence arguments has been that they do not guarantee the truth of a converged-upon claim since the common causes that are necessary and sufficient to ensure convergence apply independently of the truth of the claim. We now rebut this general argument on the grounds that the procedures generating a claim must be at least minimally reliable. Doing this, one might argue, restores the connection of the noted convergence to the truth of the claim being converged upon.

For this minimal reliability requirement to apply, the non-truth-related common causes of the procedures that lead to convergent claims must, in fact, be truth-related. Again, these common causes apply in each instance where a procedure generates a claim; other causes are idiosyncratic and private to the procedures. So it must turn out that the sorts of common causes identified above for human perceivers, from shared visual abilities to similar conceptual frameworks, lead to procedures that are better than chance at producing correct results. Of course, there's room to dispute such a claim of minimal 'competence' (again, Dietrich and Spiekermann's terminology). Leaving that point aside, a minimal reliability requirement leads to the result that the emphasis on diversity underlying convergence arguments is dispensable. So long as the relevant non-truth-related common causes are present, procedures will lead to convergent results and the results will be accurate at the rate better than chance, according to minimal reliability, whether or not diverse private causes are present, whether the procedures are distinctive in any way. There's no need to diversify procedures at all: one can use the same procedure again and again for the same

effect since the procedure works by means of a process that is minimally reliable. We cited above a problem for conditional independence convergence arguments stemming from the fact that such arguments work even if we use new procedures that are only slightly different from previous procedures. It was for this reason that we turned to the intersection of independent lies convergent argument as a way to ground the need for diverse procedures. We are now led back to where we started: we once again have the result that repeating the same procedure again and again — with only the slight difference that this procedure occurs at a different time — boosts the likelihood that a generated claim is true, approaching its probabilistic limit, whatever this might be, as with standard, conditional independence convergence arguments.

Thus, when it comes to the value of explaining the value of diverse evidential sources, it does not appear that convergence arguments have the answer. Such arguments work with absolutely the bare minimum of diversity, simply using the same procedure at a different time, which is hardly better than no diversity at all. We need a better explanation for the epistemic value of diversifying evidential sources.

## 5. Coherence arguments

The question we posed at the beginning of the paper was, “What epistemic advantage is there in having multiple, diverse perspectives converge regarding a claim about the world?” The answer suggested by convergence arguments is that diversity is a means to generate claims that have a greater likelihood of truth. If it’s the case, as I argued above, that neither conditional independence convergence arguments nor intersection of independent lies convergence arguments are successful in this respect, then perhaps there’s a better way to answer this question. The approach to be defended in the final part of this paper explores this advantage in terms of a coherence argument. The idea of considering coherence instead of independent convergence is suggested at the end of Van Cleve (2011, p.372-373), and motivates the following approach. A similar approach is also defended in Hudson (2020, p.45-46) with application to a historical case study in science.

The structure of a coherence argument is fairly simple. Suppose we have three properties, 1, 2, and 3, consistent with each other, that are each indicative of an A-kind of object. For example, a sweet taste, a whitish interior and a crunchy texture, are all indicative of an apple-kind of object. Then, suppose we have three minimally reliable procedures that generate the following results: for some object *X*, procedure 1 generates a result that *X* has property 1; procedure 2, the result that *X* has property 2; and procedure 3, that *X* has property 3. We conclude that *X* is an A-kind object, with coherent 1, 2 and 3 properties.

Of course, even though we only suggest here three properties, any larger number

of properties approaching infinity strengthens the argument. The underlying idea is that any real object in the world has an infinity of properties that are consistent with one other. The more properties we confirm an object as having, the closer we are to understanding the object as it exists in the world. The benefit of having diverse perspectives about an object thus rests on the ability to attach to an object diverse properties, not on the ability to attach to an object the same property in a variety of ways. This is not to suggest that the diverse properties are independent of one another. As Van Cleve suggests, “here independence is not wanted; indeed, independence of content is the antithesis of coherence” (2011, p.373). For example, if it is noted that an object has a sweet taste, a whiteish interior and a crunchy texture, then one (fallibly) concludes that the object is an apple. In no way are these properties independent of each other. If we know that the object has the first two properties it becomes more likely that it has the third. These properties all occur together in apples, and so are dependent on one another.

A key part of this approach is that different procedures generate different claims, with each claim attaching to an object a different property. This contrasts with a convergence argument in which the same claim is derived in multiple ways. It is a weakness of a convergence argument that some procedures are more reliable than others. Intuitively, where one has a reliable procedure for proving a claim, it doesn’t make much sense to use a less reliable procedure for proving the same claim. This sort of weakness doesn’t apply to a coherence argument. Procedure 1 may be more reliable in reporting that an object has property 1 than procedure 2 is in reporting that this object has property 2. Even then, there’s an advantage in using procedure 2 since we’re learning something new about the object not revealed by procedure 1, rather than confirming what procedure 1 already told us.

As we saw above, one of the obstacles to an intersection of independent lies convergence argument is that multiple procedures convergently generate the same claim only if one assumes the presence of a set of common causes underwriting this convergence. The problem with this set of common causes is that it is designed to ensure the convergence of results, and can succeed at this independently of the accuracy of these results and the reliability of the procedures generating them. The convergent results of multiple procedures is achievable, without these results being true. This sort of consequence does not arise with a coherence argument since the different procedures used to attach different properties to an object need not have anything in common. The procedures are chosen solely for their reliability; it is no detriment if the procedures involve different assumptions or mechanisms, so long as this does not undermine the claim that the multiple properties thus attaching to an object by means of these procedures cohere, that is, are consistent with one another.

Recall the concern with conditional independence convergence arguments that one can achieve diversity in one’s procedures simply by performing minimal revisions

to these procedures. By comparison, with a coherence argument, it is likely that the relevant procedures will be much different from one another since they're being used to identify different properties, such as properties involving different sensory modalities or referring to different physical phenomena. The sense of diversity at play in a coherence argument is not necessarily at odds with the conception of diversity propounded by Helen Longino. Having diverse perspectives as regards what procedures one should use in defending a claim is not necessarily a benefit if one of these procedures is guaranteed to be more reliable than the others. On the other hand, this diversity is a valuable tool if it plays a role in identifying the different properties an object possesses. For example, it may be that only someone from a distinct social class recognizes an object as having a certain property. Individuals from other social classes are blind to these properties. Again, there may be properties whose identification is gender-related. Diversity in this respect is much like interdisciplinarity where it takes a different kind of discipline to see an object in a certain light. In this sense, where subgroups have special abilities in identifying certain properties of objects, it doesn't make much sense to look at different procedures, adopted by different social groups, as converging on the same claim. It is better to just let these special subgroups do the job they're best at, identifying properties of a certain kind, with the overall goal of filling out the details of what an object is like. And what an object is like is something with an enormous diversity of properties. For this reason, we need a plethora of diverse procedures, each capable of identifying one or more of distinct properties of an object. In this way Longino's 'perspectival realism' (on which, see Julian Go 2016, p.162) grounds the value of diversifying our reporting procedures, inclusive of subaltern ones.

The key problem for the intersection of independent lies convergence argument is not only that truth is not guaranteed. Manifestly false claims, too, can be diversely justified, a point illustrated by the convergent support enjoyed by geocentrism in our 16<sup>th</sup> century 'interviewee' case. A coherence argument does not lead to this result. We suppose there to be multiple, diverse procedures ascribing to the universe multiple, confirmed properties. Some of these procedures generate the claim that the universe is geocentric, as revealed in our interviews. As we suggested, these procedures could be minimally reliable, such as the minimally reliable procedure of trusting the astronomical experts of the day (16th century Italy). The question then is not, "What other procedures generate the same results?" The question is, "What other properties of the universe do other minimally reliable procedures inform us about?" This sets the stage for a negative assessment of geocentrism since some of these other procedures involve Galileo's telescopic observations of Venus' phases (which are difficult to account for with geocentric astronomy) and of Jupiter's moons (which show that, contrary to geocentrism, not every celestial body goes around the Earth). In the end, geocentrism along with other confirmed claims about the universe do not lead to a

very coherent picture, a fact we learn by using a diverse set of procedures. Such a scientific approach is more fruitful than simply taking one claim and repeatedly showing that alternate procedures generate the same claim, as with convergence arguments. Again, the underlying idea to such arguments is that, with increasing convergence, the probability of a converged-upon result increases up to its highest probabilistic value, at which point it becomes compelling. Unfortunately, for intersection of independent lies convergence arguments, this increase is not guaranteed since, as shown above, convergent views can easily lead to falsities. Alternatively, with conditional independence convergence arguments, the increase in probabilistic value is too easy since such arguments work with hardly any diversity. With coherence arguments, on the other hand, this probabilistic increase is moot. We are not striving to become increasingly sure, up to the highest probability possible, about a handful of claims about the world. Rather, our goal is to construct richer, more detailed accounts of the world. We do this because the world is, in fact, extremely rich and detailed. For this reason our richer, more detailed accounts are closer to the truth.

## 6. Conclusion

At the start of the paper, we posed the question, “What epistemic advantage is there in having multiple, diverse perspectives converge on a claim about the world?” We have answered this question negatively: there is no epistemic advantage in using multiple, diverse procedures to generate the *same* result. As Descartes notes, some of these procedures are difficult to perform, so there may be little convergent testimony about them. Also, some of them involve patently false beliefs and are convergent anyways, as suggested by Russell’s account of the table. On the other hand, we have identified an epistemic advantage in using multiple, diverse procedures to generate *different* results. To this extent, Longino and Anderson have the right view. The value of diversity rests not in continually increasing the probability of an individual claim, as with a convergence argument. The value of diversity rests in giving a voice to multiple procedures generating varied results, each with a distinctive view of the world. Many of these results, to be appreciated, need a specialized awareness, one that is cultural, and possibly scientific. In each case, the awareness is irrevocably fallible, and the task of perpetually increasing their probability by means of convergence arguments is not the goal of inquiry. Rather, the goal is to formulate a picture of the world that is rich and detailed, coherently containing a vast number of confirmed beliefs about the properties of objects, under the assumption that the world itself is highly rich in the properties it coherently possesses.

## References

- Anderson, E. 1995 [2002]. Feminist Epistemology: An Interpretation and a Defense. In: Y. Balashov; A. Rosenberg (ed.), *Philosophy of Science: Contemporary Readings*, p.459–488. New York: Routledge.
- Blitstein, M. 2001. Do We Need Self-Justified Knowledge to Know at All? Unpublished manuscript, originally prepared as a term paper for an undergraduate course at Brown University in the spring of 2001, cited by van Cleve (2011, p.379).
- Cohen, L. J. 1976. How Can One Testimony Corroborate Another? In: R. S. Cohen et al. (ed.), *Essays in Memory of Imre Lakatos*, p.65–78. Dordrecht, Holland: D. Reidel.
- Descartes, R. 1912 [1637]. *Discourse on the Method of Rightly Conducting the Reason, and seeking Truth in the Sciences*. Trans. J. Veitch. London: J. M. Dent and Sons.
- Dietrich, F.; Spiekerman, K. 2013a. Epistemic Democracy with Defensible Premises. *Economics and Philosophy* 29(1): 87–120.
- Dietrich, F.; Spiekerman, K. 2013b. Independent Opinions? On the Causal Foundations of Belief Formation and Jury Theorems. *Mind* 122(487): 655–685.
- Fitelson, B. 2001. A Bayesian Account of Independent Evidence With Applications. *Philosophy of Science* 68(3): S123–S140.
- Go, J. 2016. *Postcolonial Thought and Social Theory*. New York: Oxford University Press.
- Hudson, R. 2020. The Reality of Jean Perrin's Atoms and Molecules. *British Journal for the Philosophy of Science* 71(1): 33–58.
- Huemer, M. 1997. Probability and Coherence Justification. *The Southern Journal of Philosophy* 35(4): 463–472.
- Jeffery, R. 1987. Alias Smith and Jones: The Testimony of the Senses. *Erkenntnis* 26(3): 391–99.
- Keynes, J. M. 2007 [1921]. *A Treatise on Probability*. New York: Cosimo Classics.
- Levins, R. 1966. The Strategy of Model Building in Population Biology. *American Scientist* 54(4): 421–431.
- Longino H. 1999. Feminist Epistemology. In: J. Greco; E. Sosa (ed.), *The Blackwell Guide to Epistemology*, p.327–352. Oxford: Blackwell Publishers.
- Olsson, E. J.; Shogenji, T. 2004. Can We Trust Our Memories? C.I. Lewis's Coherence Argument. *Synthese* 142(1): 21–41.
- Reichenbach, H. 1956. *The Direction of Time*. Berkeley: University of California Press.
- Romeijn, J.; Atkinson, D. 2011. Learning Juror Competence: a Generalized Condorcet Jury Theorem. *Politics, Philosophy & Economics* 10(3): 237–262.
- Russell, B. 1912. *The Problems of Philosophy*. New York: Henry Holt and Company.
- Salmon, W. 1984. *Scientific Explanation and the Causal Structure of the World*. Princeton: Princeton University Press.
- Shogenji, T. 2003. A Condition for Transitivity in Probabilistic Support. *The British Journal for the Philosophy of Science* 54(4): 613–616.
- Sober, E. 2008. *Evidence and Evolution*. Cambridge: Cambridge University Press.
- Van Cleve, J. 2011. Can Coherence Generate Warrant Ex Nihilo? Probability and the Logic of Concurring Witnesses. *Philosophy and Phenomenological Research* 82(2): 337–380.

**Acknowledgments**

The author is grateful for the helpful comments of anonymous reviewers from Principia. The author's research is supported in part by funding from the Social Sciences and Humanities Research Council of Canada.