Patrick, K. (2018). Unity As An Epistemic Virtue. *Erkenntnis*, 83(5), 983-1002. <https://doi.org/10.1007/s10670-017-9923-1>

Publisher's PDF, also known as Version of record

License (if available):
CC BY

Link to published version (if available):
[10.1007/s10670-017-9923-1](https://doi.org/10.1007/s10670-017-9923-1)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the final published version of the article (version of record). It first appeared online via Springer at <https://link.springer.com/article/10.1007%2Fs10670-017-9923-1> . Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Unity As An Epistemic Virtue

Kit Patrick¹

Received: 10 May 2016 / Accepted: 20 June 2017 / Published online: 7 July 2017
© The Author(s) 2017. This article is an open access publication

Abstract It's widely supposed that unification is an epistemic virtue: the degree to which a theory is unified contributes to its overall confirmation. However, this supposition has consequences which haven't been noted, and which undermine the leading accounts of unification. For, given Hempel's equivalence condition, any epistemic virtue must be such that logically equivalent theories must equally well unify any body of evidence, and logically equivalent bodies of evidence must be equally well unified by any theory. Yet the leading accounts of unification in Bayesian terms, or those in terms of argument patterns, cannot satisfy these constraints conditions. The reason for this runs deep: these accounts of unification make unity depend on factors that vary between equivalent theories: the probabilistic relations of their components, or their relations to argument patterns. The solution is to abandon such accounts and instead adopt an account of unity based on worldly relations such as causation, rather than inferential relations. Such an account effortlessly satisfies the equivalence conditions, and so may describe the epistemic virtue of unity.

Suppose that unification is an epistemic virtue: the degree to which a theory unifies the evidence contributes to its overall confirmation. This has consequences for what account of unity we can accept. For unity can be an epistemic virtue only if it satisfies the 'equivalence conditions', i.e. only if logically equivalent theories equally well unify any given body of evidence, and only if logically equivalent bodies of evidence are equally well unified by any theory. Accounts of unity have been developed by Friedman (1974), Kitcher (1981, 1989), Schurz (1999), McGrew (2003) and Myrvold (2003). All these accounts of unity have something in common:

✉ Kit Patrick
Kit.patrick@bristol.ac.uk

¹ Department of Philosophy, University of Bristol, Cotham House, Bristol BS6 6JL, UK

they determine the degree to which a theory unifies the evidence by the properties of inferences licensed by the theory to the evidence. These accounts are often taken to describe the epistemic virtue of unity. However, none of these accounts can succeed in describing that epistemic virtue, because none satisfy the equivalence conditions. This is because they make how well a theory unifies the evidence depend on properties of those inferences that vary between equivalent bodies of theory or equivalent bodies of evidence.¹ The solution is to abandon the attempt to determine unity by the properties of inferences altogether. An account of unity based on Salmon's (1984, p. 276) hint that unity is determined by causal mechanisms, which has been partially sketched by Skipper (1999) and Skipper and Millstein (2005), would satisfy the equivalence conditions.

The paper proceeds as follows: Sect. 1 clarifies and motivates the claim that unification is an epistemic virtue. Section 2 argues for the equivalence conditions, and Sects. 3–5 show that our best accounts of unification cannot satisfy these conditions. Section 6 shows that the remedy lies in abandoning the inferential approach to unity and instead making unity depend on causation or some similar worldly relation. I argue that any such account will satisfy the equivalence conditions, though the full development of such an account is left to a later paper.

1 The Epistemic Virtue of Unity

Assume, for the sake of argument, that unification is an epistemic virtue. That is, unification contributes to a theory's overall confirmation.² If we make the very plausible assumption that:

Strong Pareto: for any two theories T and T' , if T has more of some epistemic virtue or virtues and less of no epistemic virtues than T' , then T is better confirmed than T' .

Then it follows that:

More unification means more confirmation: for any theories T and T' and sets of evidence claims E and E' , if T unifies E better than T' unifies E' , and they have the same amount of all other virtues, then T is better confirmed by E than T' is confirmed by E' .

At least for scientific realists who already accept the existence of epistemic virtues, there's good reason to include unity amongst them. For scientists often appeal to

¹ All references in this paper to equivalence are to logical equivalence. For economy of expression, I drop 'logically'.

² There are two types of concepts of confirmation: formal and informal. On formal concepts, 'confirmation' refers to a relation in a formal system between sentences or propositions defined in a formal language. On informal concepts, 'confirmation' refers to a certain relation between scientific evidence and scientific theories. For example, Carl Hempel gave the formal notion of confirmation as the entailment relation. But he clearly had an intuitive informal notion of confirmation that he wanted this formal notion to model; that's why he looked to scientific practice to shape his formal notion. We want a (formal) model of the intuitive informal confirmation relation on which unification is an epistemic virtue.

unity to resolve disputes between theories. Newtonian mechanics unified sub-lunar and super-lunar evidence. Plate tectonics unified evidence of fossils, volcanic activity, and the shape of continents. Both theories triumphed partly as a result of how well they unified the evidence. For example, geologist Andreason asserts:

The ultimate goal of science is unification. . . . The general theory of plate tectonics unifies plate tectonics and so-called midplate phenomena, and explains the diversity of magmas and other phenomena labelled as anomalies in the standard model. (2003, p. 36)

Other geologists, such as Foulger (2010, p. 287) and Camp and Ross (2004), argue for their own theories on the basis that they are more unifying than the alternatives. And philosophers often include such unity amongst the epistemic virtues (see e.g. Thagard 1978; Lipton 2004; Myrvold 2003 and many more). Nonetheless, this paper doesn't attempt to show that unity is an epistemic virtue. Instead, the argument of this paper is that if unity is an epistemic virtue then it must satisfy the equivalence conditions, and that the leading accounts of unity fail to do this.

2 The Equivalence Conditions

If unification is an epistemic virtue, then it must be invariant over equivalent theories and bodies of evidence. More precisely, let a theory be a sentence, and a body of theory be a set of such sentences. Let an evidence claim be a sentence describing our evidence, and a body of evidence be a set of such sentences. It will not matter for our purposes whether evidence must be observable or not.³ This section shows that logically equivalent bodies of theory must equally unify any given body of evidence, and logically equivalent bodies of evidence must equally unify any given theory. This follows more or less directly from extensions of Hempel's (1945a, b) famous equivalence condition. The section proceeds as follows: Sect. 2.1 recalls Hempel's persuasive arguments for his original equivalence condition. These same arguments justify extensions of his principle to the degree of confirmation of a theory (Sect. 2.2) and to the degree of confirmation by bodies of evidence (Sect. 2.3). Then we derive the resulting constraints on an account of unity as an epistemic virtue (Sect. 2.4).

2.1 Hempel's Equivalence Condition

Hempel (1945a, b) introduced the equivalence condition, which we may state thus:

Equivalence condition: for any two logically equivalent bodies of theory, any body of evidence that confirms one also confirms the other.

³ These assumptions are made to keep the discussion in line with the already existing literature. Although theories and evidence claims are widely treated as sentences, this approach has some well-known problems. See e.g. French (2013) for a discussion of the problems of this approach and its rivals.

The arguments in favor of the equivalence condition are so persuasive that Forster is exaggerating only slightly when he writes that “[n]obody seriously entertains denying the Equivalence Condition” (1995, p. 412). Fitelson calls the arguments for it “compelling” (2008, p. 2), and Swinburne writes that the equivalence condition is “certainly taken for granted in scientific practice” (1971, p. 321). Morgenbesser (1962) and Hanen (1967) object to the equivalence condition in an attempt to rescue Nelson Goodman’s account of confirmation. However, their arguments have been widely rejected. See Swinburne (1971) for a summary.

Hempel has two arguments for his equivalence condition. First, logically equivalent theories are different formulations of the same idea. They “have the same content” and “are different formulations of the same hypothesis” (1945a, b, p. 263). Since the content of a theory is the real object of confirmation, and since equivalent theories in this sense of ‘content’ have the same content, we must treat equivalent theories as interchangeable for the purposes of confirmation.

Second, if the equivalence condition were false then we could increase or decrease how well a theory was confirmed by replacing it with a logically equivalent one. But scientists clearly don’t concern themselves over which logically equivalent theory to endorse, nor which logically equivalent theory to use for prediction or explanation. On the contrary, they don’t distinguish at all between equivalent theories, i.e. different formulations of the same content.

We can add a third argument to Hempel’s original two arguments: if almost any sort of Bayesianism is correct, then the equivalence condition must be true.⁴ For Bayesians make confirmation depend entirely on probabilistic properties. And, by the axioms of probability, logically equivalent theories always have the same probabilistic properties.

2.2 Extending the Equivalence Condition to Degree of Confirmation

The arguments for Hempel’s equivalence condition that compel us to accept a constraint on whether evidence confirms also compel us to accept a constraint on the degree to which evidence confirms. Namely:

Equivalence condition’: for any two logically equivalent bodies of theory, any body of evidence that confirms one, also confirms the other equally well.

Each of the compelling arguments for the original equivalence condition can also be used to argue for equivalence condition’. First, given that equivalent theories are just different formulations of the same content, they must be equally well confirmed by that evidence. For the content of a theory is the real object of confirmation. So we must treat equivalent theories as interchangeable for the purposes of confirmation. Second, if equivalence condition’ were false then we could increase or decrease how well a theory was confirmed by replacing it with a logically equivalent one. Yet, as argued above, scientific practice and common sense show that this is absurd.

⁴ Not quite any sort of Bayesianism: Bayesians such as Garber (e.g. 1983) reject the claim that our probabilities ought to obey the axioms of probability. See Glymour (1980) for problems with this view. The appeal to Bayesianism is consistent with an account of confirmation that uses epistemic virtues. See e.g. Okasha (2000) and Lipton (2004).

Third, almost any form of Bayesianism implies equivalence condition' because Bayesianism makes the degree of confirmation depend on probabilistic properties and equivalent theories have the same probabilistic properties.

2.3 Extending the Equivalence Condition to Confirmation by Evidence

The arguments that compel us to accept constraints on whether and how much a theory is confirmed by a body of evidence also compel us to accept parallel constraints on whether and how much a body of evidence confirms a theory, namely:

Equivalence condition'': for any two logically equivalent bodies of evidence, any theory that is confirmed by one is also equally well confirmed by the other.

Again, the compelling arguments for the original equivalence condition can also be used to argue for evidence condition''. First, given that equivalent sets of sentences are just different formulations of the same content, they must provide equal amounts of confirmation to any theory. Second, if equivalence condition'' were false then we could increase or decrease how much confirmation our theories received by replacing our body of evidence with a logically equivalent one. Again, scientific practice and common sense show that this is absurd. Third, Bayesians make the degree to which evidence confirms a theory depend entirely on probabilistic properties, and equivalent bodies of evidence have the same probabilistic properties.

In summary, both equivalence condition' and equivalence condition'' are supported by the same arguments that made the original equivalence condition so compelling. They are amongst the most secure principles in epistemology.

2.4 The Equivalence Conditions Applied to Unity As An Epistemic Virtue

Given our assumption that more unification means more confirmation (see Sect. 1), the extended equivalence conditions place constraints on any plausible account of unity. In particular:

The theory equivalence condition for unity (TEU): equivalent theories must equally unify any given body of evidence.

The evidence equivalence condition for unity (EEU): equivalent bodies of evidence must be equally unified by any given theory.

Given equivalence condition', the stipulation that unity is an epistemic virtue makes TEU nearly irresistible. For, given that equivalent bodies of theory must be equally well confirmed by any evidence, the fact that unity contributes to confirmation shows that equivalent bodies of theory must equally unify any body of evidence (i.e. TEU). We could avoid TEU only if on every occasion where a theory has less unity than another, equivalent theory, this is off-set by that first theory having a greater amount of some other epistemic virtue or virtues. Indeed, all such theories will have to have exactly the right amount more of the other virtue(s) required to make the theories equally well confirmed, and thereby satisfy equivalence condition'.

Without any plausible candidate for such an offsetting epistemic virtue, this is merely special pleading. Henceforth, we will assume that there is no such off-setting epistemic virtue, and thus any account that allows equivalent theories to differ in unity will also allow equivalent theories to differ in degree of confirmation. Given that we already accept equivalence condition', the claim that unity is an epistemic virtue, together with the assumption that there isn't always an off-setting epistemic virtue commits us to TEU.

For parallel reasons, accepting equivalence condition", together with the assumption that unity is an epistemic virtue, make EEU nearly irresistible. Again, we can only resist EEU if there is an off-setting virtue of the sort we have already ruled out. It's worth noting that similar constraints apply not just to unity, but to any epistemic virtue realists adopt, such as simplicity, novel prediction, and non-ad-hocness. Further study is required to establish whether our explications of these epistemic virtues can satisfy the corresponding equivalence constraints. The following sections show that our accounts of unity, at least, cannot satisfy the constraints.

3 Inferential Accounts of Unity and the Equivalence Conditions

Accounts of unity have been developed by Friedman (1974), Kitcher (1981, 1989), Schurz (1999), McGrew (2003) and Myrvold (2003). Each of these accounts has been said to describe an epistemic virtue. And each make the degree of unity depend on intrinsic properties of the inferences licensed by theory to our evidence. Call any account that satisfies this description an 'inferential' account. The following sections show that these accounts don't satisfy TEU and EEU.

The problem isn't that all inferential accounts are logically incompatible with TEU and EEU. For example, a theory that determined degree of unity only by whether there were any valid deductive arguments from the unifying theory to the body of evidence or vice-versa, so that a theory with any such inferential relation to the evidence is maximally unifying and a theory without has no unity at all, would be an inferential account.⁵ Furthermore, this account would satisfy TEU and EEU since the deductive relations between theory and evidence are invariant over logically equivalent bodies of theory and bodies of evidence. Yet this account of unity is fatally undermined by, amongst other things, its implausible notion of degrees of unity, with some theories having maximal unity and all others being not unifying at all. To find a plausible account of unity, we must appeal to other properties of the inferences to determine the degree of unity. The difficulty is that these other properties may vary between equivalent bodies of theory, or between equivalent bodies of evidence.

As a simple example of an account of unity that is based on properties of inferences that vary between equivalent theory and evidence, take the naïve account suggested and rejected by Friedman (1974). On this account, the degree of unity is a function of the number of sentences in the body of theory used as premises in our

⁵ The resulting account of confirmation was suggested and rejected by Hempel (1945a, b)

inferences, and the number of sentences in the body of evidence that feature as conclusions of such inferences. But obviously, these quantities vary between equivalent bodies of theory and evidence. For instance, a body of theory may consist of many different sentences only some of which are used to derive the evidence; an equivalent body of theory consists in a single long conjunctive sentence which is used to derive the evidence. Friedman's suggested account would make the latter of the equivalent theories more unified, and thereby would violate TEU.

The following sections demonstrate that the leading accounts of unity also fail to satisfy the equivalence conditions, because they also determine degree of unity by properties of inference that vary between equivalent bodies of theory or bodies of evidence. Our strategy will be to examine the developed accounts of unity that come closest to satisfying both TEU and EEU, and showing why their attempts to do so must fail. First, Sect. 4 examines the Bayesian accounts that make unity depend on inductive (probabilistic) relations between theory and evidence. The section shows that whilst their use of Bayesian mechanics allows them to satisfy TEU, they cannot satisfy EEU without abandoning the central idea that makes them plausible as accounts of unification. Then, Sect. 5 examines Kitcher's argument pattern account that makes unity depend on features of deductive arguments licensed by the theory. Again, this account cannot satisfy TEU or EEU whilst remaining plausible. The failure of other accounts of unity can readily be demonstrated by parallel arguments.

4 Bayesian Accounts of Unification

McGrew (2003) and Myrvold (2003) describe their accounts of unity using Bayesian epistemology. They explicitly advocate these accounts of unification as describing an epistemic virtue. Lange (2004) has objected that Myrvold's account has counterintuitive consequences. Their views have been defended and shown to be closely connected by Schubach (2005).

McGrew (2003, p. 562) describes consilience in Bayesian terms, stating that one theory is more consilient than another (in our terms, more unifying) when it makes our evidence claims more positively relevant to each other. That is, for any theories T and T' , T unifies any body of evidence $\{E_1, E_2, \dots, E_n\}$ better than T' just in case:

$$\frac{P(E_1 \& E_2 \& \dots E_n | T)}{P(E_1 | T) \times P(E_2 | T) \times \dots P(E_n | T)} > \frac{P(E_1 \& E_2 \& \dots E_n | T')}{P(E_1 | T') \times P(E_2 | T') \times \dots P(E_n | T')}$$

In the same spirit, Myrvold (2003) measures unifying power by the amount of change in positive relevance a theory provides:

$$\log_2 \frac{P(E_1 \& E_2 \& \dots E_n | T)}{P(E_1 | T) \times P(E_2 | T) \times \dots P(E_n | T)} - \log_2 \frac{P(E_1 \& E_2 \& \dots E_n | T')}{P(E_1 | T') \times P(E_2 | T') \times \dots P(E_n | T')}$$

As Schubach observes, these are similar measures. For when we are comparing bodies of evidence that are equally positively relevant, then Myrvold's measure implies that the unification of a theory is proportional to a factor, and McGrew's

measure implies that one theory is more unifying than another just in case it has a greater amount of that same factor. This factor is:

$$BAYESUNITY : \frac{P(E1 \& E2 \& \dots \& En | T)}{P(E1 | T) \times P(E2 | T) \times \dots \times P(En | T)}$$

The differences between the Bayesian accounts of unity won't matter for our discussion; their relation to BAYESUNITY is sufficient to ensure that EEU is violated.

The Bayesian accounts are attractive. Their central idea is that highly unifying theories are those which show that the evidence claims are highly interrelated and dependent upon one another. This is very intuitive. The Bayesian accounts cash out this intuition using probability: a more highly unifying theory is one which makes evidence claims more highly probabilistically dependent upon one another. Furthermore, the Bayesian accounts come with arguments that show that theories which better unify a body of evidence are better confirmed by that body of evidence, *ceteris paribus*. This seems to show that unification as described by the Bayesian accounts is an epistemic virtue. Nevertheless, the Bayesian accounts cannot describe an epistemic virtue of unity, because although they satisfy TEU (Sect. 4.1), they violate EEU (Sect. 4.2). The *ceteris paribus* arguments don't show that unification is an epistemic virtue (Sect. 4.3).

4.1 The Bayesian Accounts Satisfy TEU

The Bayesian accounts satisfy TEU because they ensure that all equivalent theories equally unify any given evidence. For each of the terms of BAYESUNITY will be the same if we swapped T for any logically equivalent theory. Indeed, this holds true no matter what we make conditional on T.

To prove for any sentences T, T' and E, if T and T' are logically equivalent then $P(E|T) = P(E|T')$.

Proof Since T and T' are logically equivalent

$$P(T \& \neg T) = P(\neg T \& T') = 0 \quad (1)$$

From (1)

$$P(T \& \neg T' \& E) = P(\neg T \& T' \& E) = 0 \quad (2)$$

By definition

$$P(E|T) = \frac{P(T \& E)}{P(T)} = \frac{P(T \& T' \& E) + P(T \& \neg T' \& E)}{P(T)} \quad (3)$$

From (2) and (3):

$$P(E|T) = \frac{P(T \& T' \& E)}{P(T)} \quad (4)$$

By an argument parallel to 3–4:

$$P(E|T') = \frac{P(T \& T' \& E)}{P(T')} \quad (5)$$

Since T and T' are logically equivalent:

$$P(T) = P(T') \quad (6)$$

From (4), (5) and (6):

$$P(E|T) = P(E|T') \quad (7)$$

□

Since (7) is true for any sentence E, all the terms in BAYESUNITY will be unchanged if we exchange T for any logical equivalent theory. Therefore, BAYESUNITY will be the same no matter which amongst equivalent theories we use, and this is sufficient to ensure that the Bayesian accounts satisfy TEU. More informally, they satisfy TEU because the only property of theories they use to determine unity is how likely the theory makes other sentences. And such properties remain the same for all logically equivalent theories.

4.2 The Bayesian Accounts Violate EEU

The Bayesian accounts violate EEU: they allow equivalent bodies of evidence to be differently unified by the same theory. Any logically equivalent bodies of evidence must make the numerator of BAYESUNITY the same, for the simple reason that the conjunction of any equivalent bodies of evidence will always have the same probabilistic properties, as can be proved with an argument parallel to that in Sect. 4.1. However, the denominator of BAYESUNITY may vary between logically equivalent bodies of evidence.

To prove $\{E_1, E_2, \dots E_n\}$ being logically equivalent with $\{E_a, E_b, \dots E_x\}$ is consistent with $P(E_1|T) \times P(E_2|T) \times \dots P(E_n|T) \neq P(E_a|T) \times P(E_b|T) \times \dots P(E_x|T)$.

Proof by example For bodies of evidence be $\{E_1, E_2\}$ and $\{E_a, E_b\}$, where $E_a = (E_1 \& E_2)$, and E_b is trivially entailed by T and $(E_1 \& E_2)$. So defined, the two bodies of evidence are logically equivalent.

Suppose also that T unifies E1 and E2 because:

$$P(E_1|T) \times P(E_2|T) < P(E_1 \& E_2|T) \quad (8)$$

Since, by stipulation, $E_a = (E_1 \& E_2)$:

$$P(E_a|T) = P(E_1 \& E_2|T) \quad (9)$$

Because E_b is entailed by T:

$$P(E_b|T) = 1 \quad (10)$$

By substitution of (b) and (c) into (a):

$$P(E1|T) \times P(E2|T) < P(Ea|T) \times P(Eb|T) \quad (11)$$

□

Therefore, the Bayesian accounts suffer violate EEU: they allow that two equivalent bodies of evidence may be differently unified by the same theory.⁶ More informally, they violate EEU because they make the degree of unity depend on the way we divide the body of evidence up into evidence claims with different probabilities, and this varies between equivalent bodies of evidence.

The problem is deep; I can see no way to avoid it whilst retaining the central idea that unifying theories probabilistically unite evidence claims in the body of evidence. There's no way to avoid the possibility that equivalent bodies of evidence whose members have different probabilities are probabilistically interrelated by theories in different ways. In short, the violation of EEU arises from the Bayesian accounts' central idea that unification is a matter of probabilistically relating the elements of our body of evidence.

4.3 The Ceteris Paribus Arguments

As noted above, the Bayesian accounts of unification were initially attractive because there are arguments which show that more unified theories are better confirming, all else being equal. It seems to follow that unity as described by the Bayesian accounts must be an epistemic virtue. Both McGrew (2003) and Schupbach (2005) provide such a ceteris paribus argument. I give a simplified form of the ceteris paribus argument, adapted from McGrew (2003, p. 562):

To prove theories and evidence with a greater BAYESUNITY will be better confirmed, ceteris paribus. More precisely, for any theories T , T' and any bodies of evidence $\{E1, E2...En\}$, $\{D1, D2...Dn\}$, the required ceteris paribus conditions are:

- i. $P(T) = P(T')$
- ii. $P(E1 \& E2 \& \dots \& En) = P(D1 \& D2 \& \dots \& Dn)$
- iii. $P(E1|T) \& P(E2|T) \& \dots \& P(En|T) = P(D1|T') \& P(D2|T') \& \dots \& P(Dn|T')$

Under condition iii the denominator of BAYESUNITY is equal between T , $E1, E2...En$ and $T', D1, D2...Dn$. So $T, E1, E2...En$ will have greater BAYESUNITY just in case:

- iv. $P(E1\&E2\&...En|T) > P(D1\&D2\&...Dn|T')$

⁶ This argument complements Novack's (2007) argument that the Bayesian accounts of variety make confirmation language dependent. The argument presented in this paper relies on the independently plausible evidence equivalence condition, rather than on an assumption about language dependence that is denied by some followers of Carnap working in formal epistemology. The equivalence conditions discussed here are also more general than the language dependence constraints on Novack's argument, showing that nothing that can vary between logically equivalent bodies of evidence can influence their confirmation.

But since conditions i. and ii. fix the priors of the theory and evidence as the same in both cases, by Bayes' theorem, it follows from iv. that: $P(T|E1 \& E2 \& \dots E_n) > P(T'|D1 \& D2 \& \dots D_n)$

And this implies that T is better confirmed than T'.⁷

These arguments can seem puzzling. For they seem to commit the Bayesian to the following apparently inconsistent claims:

1. **The Bayesian accounts violate EEU:** equivalent bodies of evidence are differently unified by the same theory, on the Bayesian accounts of unity.
2. **The equivalence condition":** for any two logically equivalent bodies of evidence, any theory that is confirmed by one is equally well confirmed by the other.
3. **The ceteris paribus arguments:** for any theories T1, T2 and bodies of evidence {E1, E2...En}, {D1, D2...Dn}, if T1 unifies {E1, E2...En} more than T2 unifies {D1, D2...Dn} then T1 is better confirmed by {E1, E2...En} than T2 is by {D1, D2...Dn}, all else being equal.

These seem inconsistent, but there is no fault with the ceteris paribus arguments, so it seems that the Bayesian must reject either 1 or 2. Alas, they cannot reject either: they must accept 1 given the proof in Sect. 4.2, and they must accept 2 on the grounds discussed in Sect. 2.3 that it follows from orthodox Bayesianism. Something is amiss. But it's neither equivalence condition" nor the claim that the Bayesian accounts violate EEU that's at fault.

The puzzle is solved once we realise that 1–3 are consistent with one another. For the ceteris paribus conditions mentioned in 3 exclude precisely those cases which lead to the violation of EEU mentioned by 1, i.e. they exclude precisely those cases where the bodies of evidence have members which are probabilistically interrelated by the theory in different ways. So the ceteris paribus proofs merely show that the theories are equally well confirmed by bodies of evidence in unproblematic cases. They do not show that unity in the Bayesian sense is an epistemic virtue, or the Bayesian accounts can satisfy EEU, or that the equivalence condition" is suspect.

5 The Argument Pattern Account of Unity

Kitcher (1981, 1989) sketches the 'argument pattern' account of unity, which develops and corrects Friedman's (1974) account. Kitcher's central insight is that unification is a matter of how we derive our knowledge; we unify better if we derive more of our knowledge in more similar ways. The details required for the argument below are as follows: we have a set of sentences we accept. Using these we can construct accepted arguments, whose premises and conclusions are accepted sentences. The similarity of sets of accepted arguments is assessed using argument patterns: deductive arguments with some of the terms replaced by dummy letters,

⁷ This holds for all the major Bayesian measures of degree of confirmation, e.g. all those discussed by Tentori et al. (2007). Any measure which don't make T better confirmed than T' are thereby highly suspect. I know of none that have been proposed.

together with substitution instructions telling us how we can fill in those dummy letters, and a classification telling us which sentences in the argument pattern will be premises and which conclusions. Kitcher provides the following illustration of an argument pattern:

- (1) G and G* descend from a common ancestral species S.
 - (2) Almost all organisms in S had property P.
 - (3) P was stable in the lineage leading from S to G . . .
 - (4) P was stable in the lineage leading from S to G*.
 - (5) P is heritable: that is, almost all offspring of parents both of whom have P will have P.
 - (6) Almost all members of G have P and almost all members of G* have P.
- Filling Instructions: P is to be replaced by the name of a trait, G and G* by the names of groups of organisms (populations, species, genera, higher taxa), S by the name of a species.
- Classification: (1)–(5) are premises; (6) is derived from (1)–(5) using mathematical induction on the lineages. (Kitcher 1989, p. 443)

The unity of a set of argument patterns is a function of three factors.

Argument pattern unity: the degree to which a set of argument patterns unifies our body of evidence is a function of:

- a) the number of conclusions of accepted arguments that are instances of argument patterns in the set.
- b) the number and similarity of argument patterns in the set.
- c) the stringency of the argument patterns in the set.

More stringent argument patterns are those with more similar possible instances. An instance of an argument pattern is simply an argument which matches its non-logical vocabulary, classification, and obeys its filling instructions. Kitcher observes that stringency is a function of the constraints placed on what instances an argument pattern may have by the filling instructions, non-logical vocabulary, and classification.

The charm of the argument pattern account is that it links unity with the details of scientific practice. Working scientists often use something like argument patterns. And Kitcher's account implies that unification is a matter of using those argument patterns in systematic and wide-ranging ways. He gives several detailed case studies to this effect (e.g. 1989, pp. 438–447).

Thus far, Kitcher has sketched an account of the unity of sets of argument patterns, rather than the unity of theories. But Kitcher intends the account to extend to the unity of theories. As he summarizes his key idea: “a *theory* unifies our beliefs when it provides one or more (generally a few) pattern(s) of arguments which can be used in the derivation of a large number of sentences which we accept.” (1981, p. 514, emphasis added). And he recognizes that extending the idea of unificatory power to theories is necessary to make sense of scientific practice (1981, p. 512, 1989, p. 438). Philosophers who have discussed Kitcher's work commonly assume that his account extends to an account of the unity of theories and evidence, and

often assume that it may be extended to give an account of the epistemic virtue of unity (e.g. Day and Kincaid 1994, p. 278, Lipton 2004, p. 139). It is this extension of the argument pattern account to give an account of the epistemic virtue of unity of theories by evidence that this section assesses.

Kitcher doesn't explicitly give details about how to extend his account to the unity of theories; he doesn't say precisely how a theory "provides" argument patterns.⁸ Nevertheless, there is an obvious connection between the unity of argument patterns and theories and bodies of evidence (in our sense of sets of sentences). Our accepted theories and bodies of evidence provide the premises and conclusions of the accepted arguments respectively. So which theories and evidence we accept will determine which arguments are accepted. And unity in turn is determined by the number of conclusions of those accepted arguments (a), and their similarity as measured by argument patterns (b and c). In short, a theory helps us unify by providing the premises of arguments that systematize our knowledge well.

So a theory's unity should be determined by how much it helps us by providing systematizing arguments. In Kitcher's terms, let a theory's unity be determined by the unity of the most unified set of argument patterns whose instances it makes acceptable; call this the argument pattern set associated with the theory. The more unified the associated set, the more unified the theory. In the context of epistemic virtues and theory choice, we will be interested in theories we don't accept yet. Kitcher makes plain that his account should be extended to talk about the unity of theories which we don't accept (1981, p. 519). So for theories T1 and T2 none of whose members we accept, T1 unifies our evidence better than T2 if the pattern set associated with T1 is more unified than the pattern set associated with T2.

Likewise, we could extend Kitcher's account to give a comparison of how well a theory unifies different bodies of evidence. For any two bodies of evidence E1 and E2 none of whose members we accept, E1 is more unified by our theories than E2 if the most unified argument pattern when we accept E1 is more unified than the most unified argument pattern when we accept E2.

It's a strength of our argument that we don't need to settle how to combine factors a–c to determine overall unity. Kitcher foregoes specifying how these factors trade-off against one another, i.e. he doesn't specify which function of these factors gives the unity of a set of argument patterns or a theory (1981, p. 552, 1989, p. 478). Many of the objections to the argument pattern account claim that it fails to match our intuitions about which theories unify. For example, Woodward (2005, pp. 360–370) suspects that the argument pattern account makes common mathematical structures or Linnaean classifications highly unifying, yet Woodward holds that they are neither explanatory nor highly unifying in the relevant sense. Yet

⁸ Especially given his broader project and more recent work, Kitcher may be best taken as analysing theories not as truth-apt claims but rather as non-truth-apt parts of practice. If so, then he does not take the argument pattern account of unity to give an account of the epistemic virtue of theories in our sense, for theories are not truth-apt. Further work would be required to show that a reduction of theories to practice would avoid problems parallel to those raised in this paper. Here we focus on whether argument patterns and Kitcher's core thought can give us an adequate account of unity as an epistemic virtue of (truth-apt) theories, as it has been widely supposed to do. Thanks to an anonymous reviewer for these points.

Woodward's suspicion cannot be checked because, without fleshing out the argument pattern account, we cannot determine whether mathematical structures and Linnaean classifications provide the most unifying argument patterns. Likewise, Steel (1998) and Mäki and Marchionni (2009) argue that Kitcher's account makes the wrong judgments about what theories are unified in the history of science. Yet, as Jones (1995) observes, these objections are indecisive until Kitcher's account is fleshed out to allow us to tell more precisely which function of factors a–c is unity. Similar problems plague Roland's (2008) and Humphrey's (1993) objections to Kitcher's account. There is always the possibility that Kitcher's account can be fleshed out to avoid these objections. However, the arguments in this paper are different. No matter what how a–c trade off against one another, the resulting account of unity cannot plausibly satisfy TEU and EEU.

5.1 The Argument Pattern Account of Unity and TEU

Kitcher's (1981) argument pattern account of unity violates TEU, because it makes unity depend on the non-logical vocabulary used in the theory. Logically equivalent theories may have different, even distinct, non-logical vocabulary. Such logically equivalent theories will make different arguments acceptable. And, since argument patterns must share some non-logical vocabulary with their instances, such logically equivalent theories will have different associated pattern sets. The different sets can be differently unifying, implying that the two equivalent theories are differently unifying, in violation of TEU.

To illustrate this, we can construct cases of equivalent theories that would be differently unified if the argument pattern account were correct. This is simple when one of the equivalent theories uses projectable predicates and the other does not. For, to avoid mischief caused by argument patterns with gerrymandered predicates, Kitcher requires that the non-logical vocabulary of argument patterns must be in terms of projectable predicates (see Kitcher 1989, pp. 481–482 for details). If we accept a theory with projectable predicates, this would allow us to accept arguments using those projectable predicates, and the resulting set of argument patterns may lend the theory some unity. However, an equivalent theory with no projectable predicates cannot be associated with any argument patterns, and so cannot be equally unified. And there are other ways of cooking up equivalent theories that would be differently unified on the argument pattern account.

The problem is deep: the argument pattern account violates TEU because it makes unity dependent on the non-logical vocabulary used to express a theory, which varies between equivalent theories, rather than making unity depend on the features which do not vary; i.e. what Hempel calls "content" (see Sect. 2.1). The problem arises, not from any technical details, but rather from Kitcher's central thought that highly unifying theories are those that systematize by showing how to derive much of our evidence in a similar way, i.e. using similar argument patterns.

We might hope to modify the argument pattern account to avoid violating TEU by dropping the natural assumption that theories are associated with a set of

argument patterns only if they provide the relevant non-logical vocabulary.⁹ However, this modified argument pattern account loses the attractiveness of the original account, because it loses the close connection between the argument patterns that scientists use and the argument patterns that are used to assess their theory's unity. On the original account, it's easy to see why a unified theory is desirable: because it provides the language and argument patterns that scientists use to derive a large amount of the evidence. On the modified account, the argument patterns that scientists actually use might be entirely different from the argument patterns that are used to assess their theory's unity. For example, although evolutionary biologists use the argument patterns such as the one quoted above, there may be an entirely different set of argument patterns that should be used to judge the unity of the theory of evolution. Furthermore, the modified account leaves the role of theory in unification mysterious. On the original argument pattern view, our theory gives us the premises and non-logical vocabulary required to unify. But on the amended view, this non-logical vocabulary isn't provided by the theory. This makes it difficult to see in what sense it's the theory that unifies the evidence. And this shows that the argument pattern account has little hope of satisfying TEU.

In summary, without the constraints on argument patterns provided by the non-logical vocabulary of the theory, the argument pattern account cannot give a plausible account of unity. But with those constraints, the argument pattern account violates TEU.

5.2 The Argument Pattern Account and EEU

On Kitcher's (1981) account, the argument pattern account of unity violates EEU, for the same reason that it violates TEU: it allows the non-logical vocabulary of the evidence to influence unity. Logically equivalent bodies of evidence that use different languages will make different arguments accepted, and thereby will change how unified theories are.

As in Sect. 5.1, we can avoid violating the equivalence condition by taking argument pattern unification to be determined by the deductive closure of our accepted sentences. Trivially, logically equivalent bodies of evidence have the same deductive closures, and so accepting one or the other of two equivalent bodies of evidence cannot affect unity.

Although Kitcher (1989, p. 431) treats using the deductive closure of our body of evidence (and our accepted sentences in general) to calculate unity as a convenience, this move avoids violating EEU only at the cost of facing an unpalatable dilemma. For any argument pattern AP there is a 'logically expansive' alternative argument pattern AP* which consists of all the same elements but adds

⁹ Careful attention shows that Kitcher's later work (1989, p. 431) has this consequence, although perhaps by accident. Kitcher, apparently as a matter of convenience, stipulates that our set of accepted sentences is deductively closed. This idealization has profound consequences. For logically equivalent theories entail one another. On the assumption that theories in our sense are part of the accepted sentences it follows that, given Kitcher's idealization, for any two logically equivalent theories if we accept one then we accept both. Therefore, accepting either will lead to the same set of premises for accepted arguments with the same vocabulary, and therefore the same associated set.

an additional conclusion ‘X’ with the filling instruction ‘X is a deductive consequence of the conclusion of AP.’ The argument pattern account now faces a dilemma: on the one hand it can make a theory’s unity determined by sets of argument patterns whose members are logically expansive; on the other hand it can stipulate that a theory’s unity isn’t determined by such sets. Because of the appeal to deductive closures in calculating unity, neither option is plausible. We examine each horn of the dilemma in turn.

The first horn of the dilemma is that a theory’s unity will be determined by sets of argument patterns whose members are logically expansive. This implies that if some sentence is a conclusion of an instance of one of those argument patterns, then any deductive consequence of that sentence is also a conclusion of an instance of one of those argument patterns. It follows that factor a) is almost always incomparable between theories. The argument for this, like Tichý (1974) and Miller’s (1972) argument concerning verisimilitude, exploits the fact that any claim in a deductively closed set of sentences has infinitely many deductive consequences in that closure. Take any two logically expansive argument patterns AP1 and AP2. For any sentence S1 that is a conclusion of an instance of AP1 but not of AP2, there are countably infinitely many more such sentences that are instances of AP1 but not of AP2. Therefore, for any two logically expansive argument patterns AP1 and AP2 whose instances each have some conclusion that the other lacks, each will have countably infinitely many conclusions that the other lacks. We cannot say that one has more of factor (a) than another. Therefore, factor (a) will be almost entirely insignificant in comparing the argument pattern unity of theories, only coming into play when the conclusions of instances of one set of argument patterns are a proper subset of those of another. In short, the argument pattern account would make unification almost entirely insensitive to how much of our evidence was unified. This is dramatically implausible.

The second horn of the dilemma is that a theory’s unity isn’t determined by argument patterns with logically expansive sets. The difficulty with this is that sets with logically expansive argument patterns will simply have a higher unity than those without, given Kitcher’s central ideas. Consider any set of (non-logically expansive) argument patterns SAP. There is a more unified set of argument patterns SAP* consisting of logically expansive versions of each of the members of SAP. The members of SAP* will be equally stringent, numerous and similar to one another as those of SAP;¹⁰ i.e. SAP* will be at least as good on factor b and c as SAP. Yet SAP* will typically have infinitely many more conclusions of accepted arguments that are instances of its members than will SAP; i.e. SAP* will do far better on factor a. On the very plausible assumption of strong Pareto (Sect. 1), it follows that SAP* will be more unified than SAP. Therefore, the stipulation that a theory’s unity is determined by SAP rather than SAP* denies that a theory’s unity is determined by how much it can help us unify our evidence. For, on the argument pattern account, any theory can help us unify best by licensing the logically expansive set SAP*. So it’s a mystery why we should measure the unity of that theory by the fact that it licenses the less unifying (and non-logically expansive) set

¹⁰ See Roland (2008, pp. 501–502) for an argument that they will be more similar in the relevant sense. This stronger claim isn’t necessary for our purposes.

SAP. This section doesn't show that we cannot accept Kitcher's account of unity, his resulting account of explanation, or his broader naturalistic project. Each of these might be preserved. However, it does show that Kitcher's account of unity cannot be extended to give an account of unity as an epistemic virtue of theories.

6 How to Easily Satisfy the Equivalence Conditions

The above arguments show that the leading accounts of unification violate TEU and/or EEU. The difficulty arises because these accounts make how much a theory unifies the evidence depend on properties of inferences that vary between equivalent theories and/or bodies of evidence. Other accounts of unity depend on properties of the inferential connections that don't vary in this way. But as shown in Sect. 3, these accounts are radically implausible. Perhaps there is some hitherto undeveloped account that makes unity a function of properties of the inferential connections that don't vary between equivalent theories or bodies of evidence. It is difficult to see what those properties could be, or what central idea of unity could make such an account plausible. The task of finding an inferential account of the epistemic virtue of unity is either exceedingly difficult, or hopeless.

There is an alternative approach to unity that easily satisfies the equivalence conditions. This approach abandons the inferential approach's idea that unity is a function of the properties of inferences involving theory and evidence. Instead, it adopts the constraint that unity depends wholly on events and properties that the theory and evidence claim exist in the world. Call this the 'worldly' approach. The central idea of unity at work in accounts following the worldly approach is quite different to the central ideas employed by the Bayesian or argument pattern accounts. According to the argument pattern account, a highly unifying theory is one that systematizes our inferences by showing how to derive much of our body of evidence in a similar way. But according to worldly accounts, a highly unifying theory is one that claims the worldly events and properties are systematized. More evocatively, we may say that worldly accounts make unity a matter of postulated neatness of the world, rather than neatness of inferences.

One such worldly account of unity is suggested in passing by Salmon (1984 276), and sketched by Skipper (1999) and Skipper and Millstein (2005). On this account, the degree to which a theory unifies the evidence is measured by features of the causal mechanism it postulates to causally explain the evidence.¹¹ Roughly, the idea seems to be:

The causal account of unity: the degree to which a theory unifies the evidence is a function of the number of events in our evidence that it implies a causal mechanism for, and the number and similarity of events which cause

¹¹ Simply mentioning causation in an account of unification doesn't imply that the account is worldly on the definition given here. Roland (2008) introduces causation to solve the 'obsessive unifier' problem with Kitcher's account of unity. Roughly, Roland's idea is that the best way to exclude irrelevant information from argument patterns is to exclude causally irrelevant factors. Roland's proposal differs from ours in an important way: on his account, unity is only partly influenced by the actual causal facts; whereas on our account, unity is wholly determined by the postulated causal facts. Our account is worldly; Roland's is not.

that evidence by the theory's causal mechanism, where these are all properties of the world postulated by the theory. Unity is proportional to the first factor, and inversely proportional to the second.

The details of the account remain sketchy. A complete account would need to specify a way of individuating events and causal mechanisms and a way of determining similarity of evidence, which didn't vary between logically equivalent theories and bodies of evidence. This work has been started in a different context (e.g. Davidson 1967; Kim 1976), but the problems have not been solved to everyone's satisfaction. Even if these details were filled in a satisfying way, the account faces other *prima facie* difficulties. For example, causation is rare or entirely absent in physics yet unity remains important to theory choice in physics. And it isn't the only worldly account available, or even the only causal worldly account available. An alternative causal account might be found by adapting Strevens's (2004) causal account of explanation, which includes some ideas taken from accounts of unity. In a subsequent paper, I fully develop an account that develops Skipper and Strevens' work, and solves these problems.

However, we can already see that any worldly account like Skipper's easily satisfies TEU and EEU. For equivalent bodies of theory always refer to the same worldly objects and relations, and so do equivalent bodies of evidence. If we make unity depend solely on worldly properties postulated by theory and evidence, there is no threat of having equivalent bodies of theory differently unify the same body of evidence, and there is no threat of having equivalent bodies of evidence be differently unified by the same theory. There's a satisfying link between this easy success of worldly accounts of unity and Hempel's arguments for his equivalence condition. Hempel's equivalence condition holds because confirmation is primarily a relation between what Hempel calls the "content" of our bodies of theory and bodies of evidence, which remains unchanged no matter which equivalent body of theory or evidence we choose to express it in. The worldly accounts of unity are those that make unity depend on the content of our theory and evidence, i.e. the worldly relations and events those theories postulate. For this reason, the worldly accounts have no difficulty in taking unity to be an epistemic virtue.

7 Conclusion

The above argument shows that the leading inferential accounts of unity cannot give an account of the epistemic virtue of unity, and it does so in a way that differs from other criticisms of those accounts. Many criticisms of these accounts claim that they fail to match our intuitions about which theories unify. This paper offers a different sort of objection. By taking into account the role that unity plays in theory choice, we see that the unity must satisfy the constraints of elementary principles of confirmation. There is no apparent way to satisfy these constraints so long as we maintain an inferential approach and hold that unity is an epistemic virtue. By shifting our focus from inferences to worldly properties and events, we can easily solve these problems.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Andreasson, D. L. (2003). Plate tectonics; the general theory: Complex earth is simpler than you think. *Geological Society of America Special Papers*, 2006(413), 29–38.
- Camp, V. E. & Ross, M. E. (2004). Radiating volcanic migrations: An example from the Pacific Northwest, U.S.A. *Journal of Geophysical Research*, 109, 2156–2202.
- Davidson, D. (1967). The logical form of action sentences. In N. Rescher (Ed.), *The logic of decision and action* (pp. 81–95). Pittsburgh: University of Pittsburgh Press.
- Day, T., & Kincaid, H. (1994). Putting inference to the best explanation in its place. *Synthese*, 98(2), 271–295.
- Fitelson, B. (2008). Goodman’s “new riddle”. *Journal of Philosophical Logic*, 376, 613–643.
- Forster, M. R. (1995). Bayes and bust: Simplicity as a problem for a probabilist’s approach to confirmation. *The British Journal for the Philosophy of Science*, 463, 399–424.
- Foulger, G. R. (2010). *Plates vs plumes: A geological controversy*. Hoboken: Wiley-Blackwell.
- French, S. (2013). The structure of theories. In M. Curd & S. Psillos (Eds.), *The Routledge companion to philosophy of science*. Abingdon: Routledge.
- Friedman, M. (1974). Explanation and scientific understanding. *The Journal of Philosophy*, 711, 5–19.
- Garber, D. (1983). Old evidence and logical omniscience in Bayesian confirmation theory. *Testing Scientific Theories*, 10, 99–133.
- Glymour, C. (1980). *Why i am not a Bayesian. In his theory and evidence*. Princeton: Princeton University Press.
- Hanen, M. (1967). Goodman, Wallace, and the equivalence condition. *The Journal of Philosophy*, 649, 271–280.
- Hempel, C. G. (1945a). Studies in the logic of confirmation. *Mind*, 54213, 1–26.
- Hempel, C. G. (1945b). Studies in the logic of confirmation (I). *Mind*, 54(213), 1–26.
- Humphreys, P. (1993). Greater unification equals greater understanding? *Analysis*, 533, 183–188.
- Jones, T. (1995). Reductionism and the unification theory of explanation. *Philosophy of Science*, 62, 21–30.
- Kim, J. (1976). Events as property exemplifications. In M. A. S. Feldman, E. Sosa & E. Wierenga (Eds.), *Action theory* (pp. 159–177). Dordrecht: Reidel.
- Kitcher, P. (1981). Explanatory unification. *Philosophy of Science*, 48, 507–531.
- Kitcher, P. (1989). Explanatory unification and the causal structure of the world. In P. Kitcher & W. C. Salmon (Eds.), *Scientific explanation: Volume 13 of Minnesota studies in the philosophy of science* (pp. 410–505). Minneapolis: University of Minnesota Press.
- Lange, M. (2004). Bayesianism and unification: A reply to Wayne Myrvold. *Philosophy of Science*, 71(2), 205–215.
- Lipton, P. (2004). *Inference to the best explanation* (2nd ed.). London: Routledge.
- Mäki, U., & Marchionni, C. (2009). On the structure of explanatory unification: The case of geographical economics. *Studies in History and Philosophy of Science Part A*, 402, 185–195.
- McGrew, T. (2003). Confirmation, heuristics, and explanatory reasoning. *The British Journal for the Philosophy of Science*, 544, 553–567.
- Miller, D. (1972). The truth-likeness of truthlikeness. *Analysis*, 332, 50–55.
- Morgenbesser, S. (1962). Goodman on the ravens. *The Journal of Philosophy*, 5918, 493–495.
- Myrvold, W. C. (2003). A Bayesian account of the virtue of unification. *Philosophy of Science*, 702, 399–423.
- Novack, G. (2007). Does evidential variety depend on how the evidence is described? *Philosophy of Science*, 74(5), 701–711.
- Okasha, S. (2000). Van Fraassen’s critique of inference to the best explanation. *Studies in History and Philosophy of Science Part A*, 31(4), 691–710.

- Roland, J. W. (2008). Kitcher and the obsessive unifier. *Philosophy and Phenomenological Research*, 772, 493–506.
- Schupbach, J. N. (2005). On a Bayesian analysis of the virtue of unification. *Philosophy of Science*, 724, 594–607.
- Schurz, G. (1999). Explanation as unification. *Synthese*, 1201, 95–114.
- Skipper, R. A., Jr. (1999). Selection and the extent of explanatory unification. *Philosophy of Science*, 66, S196–S209.
- Skipper, R. A., Jr., & Millstein, R. L. (2005). Thinking about evolutionary mechanisms: Natural selection. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 362, 327–347.
- Steel, D. (1998). Warfare and western manufactures: A case study of explanation in anthropology. *Philosophy of Science*, 65, 649–671.
- Strevens, M. (2004). The causal and unification approaches to explanation unified—Causally. *Noûs*, 38(1), 154–176.
- Swinburne, R. G. (1971). The paradoxes of confirmation: A survey. *American Philosophical Quarterly*, 84, 318–330.
- Tentori, K., Crupi, V., Bonini, N., & Osherson, D. (2007). Comparison of confirmation measures. *Cognition*, 103(1), 107–119.
- Thagard, P. R. (1978). The best explanation: Criteria for theory choice. *The Journal of Philosophy*, 75, 76–92.
- Tichý, P. (1974). On Popper's definitions of verisimilitude. *The British Journal for the Philosophy of Science*, 252, 155–160.
- Woodward, J. (2005). *Making things happen: A theory of causal explanation*. New York: Oxford University Press.