**QUEEN'S
UNIVERSITY
BELFAST**

Evidential Holism

Morrison, J. (2017). Evidential Holism. *Philosophy Compass*. <https://doi.org/10.1111/phc3.12417>

Published in:
Philosophy Compass

Document Version:
Peer reviewed version

Queen's University Belfast - Research Portal:
[Link to publication record in Queen's University Belfast Research Portal](#)

Publisher rights

© 2017 John Wiley & Sons Inc. All rights reserved.

This work is made available online in accordance with the publisher's policies. Please refer to any applicable terms of use of the publisher.

General rights

Copyright for the publications made accessible via the Queen's University Belfast Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The Research Portal is Queen's institutional repository that provides access to Queen's research output. Every effort has been made to ensure that content in the Research Portal does not infringe any person's rights, or applicable UK laws. If you discover content in the Research Portal that you believe breaches copyright or violates any law, please contact openaccess@qub.ac.uk.

EVIDENTIAL HOLISM

JOE MORRISON j.morrison@qub.ac.uk

Evidential holism begins with something like the claim that “it is only jointly as a theory that scientific statements imply their observable consequences.” This is the holistic claim that Elliott Sober tells us is an “unexceptional observation” (1993: 35). But variations on this ‘unexceptional’ claim feature as a premise in a series of controversial arguments for the following conclusions:

- There is no analytic / synthetic distinction
- The meaning of a sentence cannot be understood without understanding the whole language of which it is a part
- All knowledge is empirical knowledge (there is no apriori knowledge)
- Abstract mathematical objects (such as numbers, sets, and functions) exist
- There is no fact / value distinction

This paper is a survey of what evidential holism is, how plausible it is, and what consequences it might have. §1 will distinguish a range of different holistic claims, §2 and §3 explore how well motivated they are and how they relate to one another, and §4 returns to the arguments listed above and uses the distinctions from the previous sections to identify holism’s role in each case.

From the outset, I should state that I think research about evidential holism itself is largely in decline; much has been written about it, but not particularly recently, and this piece is an attempt to pull a lot of those things together. If there is anything like a consensus view, then it is divided between philosophers of science and epistemologists who have examined evidential holism’s claims in detail and found that very few of them are properly defensible, and on the other side, people who’ve heard some of the slogans associated with holism and have wanted to endorse or exploit those ideas in

distinct areas of philosophy. So there is clearly some sort of need for further dialogue about whether holism can hold its value as a currency in the various transactions in which it is supposed to take part.

§1

When philosophers introduce evidential holism, they often use the following two things:

- the image of a ‘web of beliefs’, in which all our beliefs are connected to all of the others, and where the web contacts the world only at the edges while the beliefs towards the centre are supported just by their connections to other beliefs,
- perhaps the most frequently quoted snippet from one of WVO Quine’s most frequently cited papers, to wit: “any statement can be held true come what may... by the same token, no statement is immune to revision” (Quine 1953: 40).

The first image is very suggestive, and helps us to picture how there could be something holistic about knowledge, about how our beliefs rely on each other or prop up each other collectively. And the quotation is provocative; once we think of our beliefs as a web, we might imagine that we can take out (‘revise’) any one of the beliefs and yet still keep the rest of the web intact and functioning. That beliefs like ‘ $5+7=12$ ’ turn out to be revisable is a radical claim – philosophers have often wanted to point out that some beliefs – mathematical, logical, or *cogito*-like beliefs – have a special or privileged epistemic status; that because of their particular contents or the way we come to acquire them these beliefs turn out to be ‘certain’, ‘indubitable’ or ‘incorrigible’. But the web-of-belief picture encourages us to think that all beliefs get their epistemic status in the same way: if the only thing that contributes to the status of any belief is to do with its connections to the rest of the beliefs in the web, then any belief is capable of being revised, and, conversely, any belief can be treated as certain, indubitable or immune-to-revision, so long as the web holds together.

While both the image of the web and the famous quotation are certainly related to evidential holism, they are better understood as downstream consequences rather than statements of evidential holism itself. Although they are both powerful pictures, neither the weblike nature of beliefs nor the claim about revisability / immunity constitute a theory about the nature of evidence. If we are trying to work out what evidential holism is, it is better to start elsewhere and to come back to these claims.

But a complication is that evidential holism is referred to using lots of different terms, often interchangeably; such terms include ‘epistemic holism’, ‘evidential holism’, ‘confirmational holism’, and ‘the Duhem-Quine’ or ‘Quine-Duhem’ thesis, since both Pierre Duhem and Quine are credited with it. It is not clear that these are all equivalent – for example, some take ‘the Quine-Duhem thesis’ to be a series of independent insights which motivate evidential holism, while others take the Quine-Duhem thesis and evidential holism to be identical and make no distinction between them. Jerry Fodor and Ernest Lepore are keen to point out that there is “a galaxy of non-equivalent (but closely interrelated) doctrines” here, but they claim that the differences are immaterial:

Since we think that most of the epistemological doctrines that cluster together in the Q/D galaxy are probably right, we will not insist on the differences between them. Sooner or later, somebody will doubtless write a doctoral dissertation that sorts them all out. (1992: 40)

We can consider ourselves doing that kind of sorting out here. To start us off, a loose definition of evidential holism is that evidence only supports or counts against entire theories, rather than specific or individual claims. Put slightly more precisely and accurately, it says that for any evidential relation¹ to obtain, one of its *relata* will be a whole system or network of claims (see Esfeld 1998 for a fully general characterisation of what makes something ‘holistic’).

Here are range of holistic doctrines, and in the next section I'll start to explore how they might be related.

Prediction: It is only jointly as a theory [a collected body of statements²] that statements imply their observable consequences.

Confirmation: The unit of (dis)confirmation is collective bodies of statements; individual statements do not get (dis)confirmed by experiments, only bodies of statements do.

Falsification: The unit of falsification is collective bodies of statements; individual statements do not get falsified by experiments, only bodies of statements do.

Underdetermination of Falsification: A failed prediction only tells us that a whole theory is false, and it doesn't specify which part(s) of the theory is (are) falsified – it could, in principle, be any one of the constituent parts that is to blame.

Revision: Any belief can be revised on the basis of empirical experience.

Values – Our values / preferences / pragmatic considerations play a central role in determining whether a statement has been confirmed / falsified.

Each of these theses tells us about how something to do with evidence or justification is holistic.³ We'll explore whether they are equally supported or motivated, but apart from the question about whether to endorse some or all of them there are are a couple of other ways that evidential holism might admit of more or less radical varieties. One dimension is what I'll call the 'scope' of holism. Compare the following:

Global scope: evidential relations (prediction / confirmation / falsification etc) depend on every single statement that makes up current science: when we say ‘whole theory’ it means ‘everything in science’.

Moderate scope: evidential relations depend upon some holistic unit smaller than the whole of science; it might be a single scientific theory, or a few of them together. Similarly, if there is a viable distinction between empirical and non-empirical statements, the ‘scope’ of holism might be limited to the class of all empirical statements.⁴

Another dimension is what I’ll call the ‘differential distribution’ of evidence. Compare the following:

Non-distributivism: it doesn’t make sense to talk about the epistemic status of an individual belief or statement; epistemic support is not something which gets parcelled out or pumped into particular beliefs or statements, it’s only ever a property of entire webs of beliefs / networks of statements.

Distributivism: even if the primary unit or bearer of epistemic status is the entire web / network, from this we are able to derive a measure of the epistemic status of an individual statement / belief.

§2

Prediction is an important part of the epistemology of science – scientific hypotheses are tested by deriving predictions from them, and experiments are devised to see whether those predictions are born out. That’s to say: part of the way that our scientific hypotheses are epistemically justified is through this process of prediction and testing. And *Prediction* says that in science, prediction is holistic. Furthermore, *Prediction* seems to be committed to a *Non-distributivist* account of prediction, in the

sense that it is impossible to ‘distribute’ or pin down the source of a prediction to particular parts of the theoretical whole.

Global scope modifies *Prediction* by telling us that the scope of the phrase ‘jointly as a theory’ is completely global; it says that a prediction depends on the total network of beliefs or claims that are at work in a theory, not just a subpart of it. So for example, if a hypothesis relies on a scientific theory which includes some mathematical claims, and those mathematical claims are themselves part of wider mathematical theories, then a *Global scope* reading of *Prediction* says that any prediction which is derived from that hypothesis depends on all of those mathematical claims and theories as well.

While Pierre Duhem certainly seemed to endorse *Prediction*, it’s not obvious that he thought that its scope had to be totally global:

“[I]n order to deduce from this proposition the prediction of a phenomenon and institute the experiment which is to show whether this phenomenon is or is not produced ... [the scientist] does not confine himself to making use of the proposition in question; he makes use also of a whole group of theories accepted by him as beyond dispute.” (Duhem 1954: chapter 6, §2).

The first clause shows that he accepts *Prediction*: he thinks that individual hypotheses taken alone do not give rise to predictions all by themselves. But the last clause suggests that he doesn’t endorse the ‘global scope’ version of *Prediction*: predictions don’t depend on *all* of the claims implicated in a theory, but rather depend on a more restricted networks of claims, those which are ‘accepted ... as beyond dispute’. In contrast, WVO Quine’s holistic view was once expressed as being completely global: that “the whole of science” is implicated in making predictions (Quine 1953: 42). He later retreated to say that only “more modest chunks” or “sufficiently inclusive” bodies of statements are required for making predictions (Quine 1981: 71). So even though both Duhem and Quine endorsed holism about

prediction, it looks as though neither of them were ultimately committed to a global-scope version of it. The ‘scope’ challenge for ‘moderate’ holists is to be able to explain why holism about prediction is correct while simultaneously restricting its scope to some unit that’s smaller than the ‘whole of science’.

The relationships between *Prediction*, *Falsification* and *Confirmation* are illustrated in figure 1 below.

Figure 1

If we think of figure 1 as a decision tree, then we can see several points where we can attempt to resist taking on board more or stronger holistic claims. But if we’re starting from the top and working down, it’s clear that we need some motivation to accept the *Prediction* in the first place.

Is *Prediction* correct? Some accept *Prediction* as a truism about scientific prediction – we saw above that Sober calls it an “unexceptional observation”. While philosophers of science have had a lot to say about the relationship between prediction and other evidential concepts like confirmation or explanation, the nature of scientific prediction itself – its psychology, its epistemology, its inferential structure – has not received very much close attention.⁵ Duhem seems to have accepted *Prediction* as a strictly descriptive statement about the habits of working physicists; in a similar vein it could turn out to be a descriptive claim about our psychological dispositions to make certain kinds of inferences. Morrison (2010) argues that *Prediction* is a consequence of a widely accepted view about the nature of scientific observations: that they are theory-laden. Both kinds of motivation could be developed further in various directions – *Prediction* could be an empirical claim about the nature of perception, or it could be about the metaphysically or epistemically necessary conditions for making predictions.

In the background here is a question about what it means to ‘derive a prediction’: are a theory’s predictions some or all of its logical consequences, or is ‘predicting’ something else? Weir (1985: 230) argues that if the business of identifying a theory’s predictions is just a question of identifying all of its first-order deductive entailments, then *Prediction* is certainly false, since the first-order predicate calculus is compact. Very loosely speaking, compactness means that if the total theory T deductively entails a particular observable consequence O, then every finite subset of T will also entail O, and none of them will entail not-O. It follows that *Prediction* is false: the whole theory T is not required to yield a prediction O, since any proper subset of T will have the same entailment. To save *Prediction*, Weir suggests that ‘deriving a prediction’ must be something other than working out the logical entailments of a theory. As such, this argument undermines the most plausible way of connecting *Prediction* to *Falsification* (see the right-hand fork of figure 1).

In scientific reasoning, how does confirmation relate to prediction? Suppose we thought that the relationship is roughly symmetrical: whatever the unit is that gives rise to a prediction, it is the same unit that gets confirmed or disconfirmed by the experimental result. So, if an individual hypothesis H can predict some experimental result E , then E confirms (and $\sim E$ disconfirms) H alone. But if holism about prediction (*Prediction*) is true, then the unit of prediction is never an individual hypothesis and only ever a collected body of statements. So on the assumption that prediction and confirmation are symmetrical, from *Prediction* we get holism about confirmation (this is the left-hand fork of figure 1):

Confirmation: The unit of (dis)confirmation is collective bodies of statements; individual statements do not get (dis)confirmed by experiments, only bodies of statements do.

As above, a stronger version of holism will modify *Confirmation* to say that the scope of these ‘collective bodies’ is global, while a more moderate holism will contend that the unit of confirmation might be smaller than ‘the whole of science’. When philosophers talk about ‘confirmational holism’, it tends to be a *Global Scope* reading of *Confirmation* that they have in mind. We’ll explore some of the applications of ‘confirmational holism’ in §4, but for the meantime we can note that it depends on both *Prediction* and the idea that prediction and confirmation are symmetrically related.⁶

At first glance, *Confirmation* looks very similar to *Non-distributivism*, since both include the idea that whole theories are the units of justification. But it’s important to see past the overlapping content and tease out the differences. *Confirmation* tells us that the result of an experiment is a boost or a decrease in the degree of support for an entire theory. *Non-distributivism* says that it is never possible to divide up the degree of a support that an overall theory has into (perhaps differing) degrees of support for

its individual parts. But to say that experiments (dis)confirm entire theories, not individual parts thereof, is not the same as saying that individual claims can never be (dis)confirmed. *Confirmation* can be true while *Non-distributivism* is false if we accept that experiments alone will not determine the degree of support of an individual statement and will only serve to boost or weaken the support for the entire theory, but that we can then bring to bear a body of further considerations, beyond the experimental evidence alone, which allows us to distribute degrees of support differentially among the component parts of the theory. There are many ways that this subsequent distribution could be effected; prominent examples might include a ‘bootstrapping’ approach (Glymour 1975, 1980), employing the Bayesian calculus and its various measures of confirmation (Strevens 2001), or Deborah Mayo’s ‘error-statistical’ approach (see Mayo 1996, 1997). For all his notorious association with holism, Quine rejected *Non-distributivism* – he was quite happy about involving the ‘technology of probability and mathematical statistics’ with his evidentially holistic position to allow that scientists might argue ‘logically or probabilistically from other beliefs already held’ to talk about the epistemic status of individual statements (Quine 1990, p. 13). But it remains a live issue to show that rejecting *Non-distributivism* will not thereby water down the holistic content of *Confirmation* (see Christensen 1992, Stevens 2001, 2005, Fitelson and Waterman 2005, 2007).

I have suggested that *Confirmation* depends upon a combination of *Prediction* and the idea that prediction and confirmation are symmetrically related. This symmetry claim is at once *prima facie* plausible and highly controversial – it has been widely rejected by most theories of confirmation (see Earman 1983; Gemes 1994, 1998; Hempel 1965: 65, Horwich 1983; Howson and Urbach 1993; Schurz 1991, 1994). But we can still get to another kind of evidential holism even if we deny that prediction and confirmation are symmetrically related. Consider the concept that Karl Popper called ‘falsification’ (Popper 1959, 1963). Falsification is supposed to be an evidential relation which obtains between a hypothesis and an evidence statement: if a hypothesis

H predicts an outcome E, and not-E occurs, then H is refuted or falsified. While confirmation and disconfirmation are inductive relations, falsification is supposed to be a deductive relation: the evidence not-E entails the falsity of H. A holistic version of falsification is obtained by endorsing *Prediction*, along with the assumption that prediction is itself a deductive relation. If the concept ‘predicts’ is considered as form of first-order deductive entailment, such that to say H predicts E is equivalent to saying H entails E, then when this is combined with *Prediction*, we find the following holistic *Falsification* as its logical corollary:

Falsification: The unit of falsification is collective bodies of statements; individual statements do not get falsified by experiments, only bodies of statements do.

Just like *Confirmation*, *Falsification* can come in weaker or stronger varieties. *Global Scope* widens the scope of ‘collective bodies of statements’ to encompass perhaps the whole of science – such that in principle a failed prediction entails that the whole of science of has been falsified. *Non-distributivism* says that from the fact that an experiment has falsified an entire theory, we cannot make any inferences about the falsification of any particular component statement within that theory. As with our previous discussion in the context of *Prediction*, global falsification-holism looks extremely radical. But again, restricting the scope to a smaller sub-part of science requires an explanation of why some sort of holism about prediction/falsification is correct while simultaneously restricting its scope to some unit that’s smaller than the ‘whole of science’. And as with our previous discussion in the context of *Confirmation*, it looks as though some further argument is required in order to insist on a *Non-distributive* version of holistic falsificationism. That is, we should not move from the idea that experiments only falsify theories to the idea that no individual theoretical statements can be falsified, without first showing how these two ideas are connected.

Falsification is a consequence of coupling *Prediction* with the idea that prediction is a deductive relationship (plus the Popperian idea that falsification is in fact an evidential relation). But this is not straightforward, as we saw above – Weir’s argument suggests that if we construe prediction as a deductive relation, then *Prediction* is certainly false. Moreover, most Bayesians maintain that scientific prediction should not be understood as a deductive relation. Instead, we assign probabilities to a hypothesis and other theoretical claims, and we then use the Bayesian formula to calculate the relative probabilities of certain observations being made. As such, hypotheses and theories do not entail observations, rather they assign a degree of expectation on particular observations (Sober 1999: 13). On this view, although it might still be acceptable to maintain that ‘only whole theories predict observations’, it is not true that they do so deductively. As a result, *Falsification* cannot be logically equivalent to *Prediction* – if ‘predicts’ is not a deductive relation then it no longer has a contrapositive in falsification. (See Sober 2000: 438).

This section explored the links between three of the evidentially holistic claims, and ways that they can admit of stronger or more moderate readings. The next section looks at the other three theses, and in section 4 I’ll explore some of the applications of these claims in other areas of philosophy.

§3

We turn now to :

Underdetermination of Falsification: A failed prediction only tells us that a whole theory is false, and it doesn’t specify which part(s) of the theory is (are) falsified – it could, in principle, be any one of the constituent parts that is to blame.

Revision: Any belief can be revised on the basis of empirical experience.

Values – Our values / preferences / pragmatic considerations play a central role in determining whether a statement has been confirmed / falsified.

I'll introduce *Underdetermination* and *Revision* first, as they are closely related to each other and often treated as if they say the same thing, and I'll return to explore *Values* at the end of this section.

Underdetermination is widely regarded as a straightforward consequence of *Falsification*. Here's an argument for *Underdetermination*: Suppose that if a theory makes a false prediction, the result will falsify the entire whole theory (*Falsification*). The experimental outcome tells us that at least one of the constituent claims of the theory must be false, because when they are taken all together the theory makes wrong predictions about what we'll observe. But this observation taken alone fails to tell us which particular member of the set of theoretical claims has been falsified – any one (or several) of the many statements making up the theory could be entirely to blame or partially contributory to the fact that the theory has made a false prediction. So the observational evidence alone is said to 'underdetermine' – it fails to specify – which of the particular claims within the theory have been falsified. In this sense, *Underdetermination* represents a instance of a *Non-distributivist* reading of *Falsification*, since it says that from the fact that a whole theory has the evidential status 'falsified', we are unable to 'distribute' that status down to any specific component of the theory.

Alex Orenstein spells out what he takes to be Quine's reason for linking *Underdetermination* to *Revision*:

"No sentence can be singled out as being in principle incorrigible; for in the attempt to fit theory to observation, any one sentence may become a candidate for revision." (Orenstein 1977: 85/86).

The argument condensed in this sentence is as follows: since the observational evidence underdetermines which of the parts of a predictively unsuccessful theory *T* are to blame for its false predictions (*Underdetermination*), it follows that any one of the constituent statements making up *T* can be replaced with some other statement such that the new theory *T** then becomes predictively successful, and thus that any statement within *T* is in principle ‘revisable’. That is, any member of *T* will be revisable just in case by revising it we thereby make the modified *T** predict the observed outcome.

But this is a notoriously contentious line of argument – indeed, it looks like a non-sequitur: why should we suppose that because any one of the constituent statements of *T* might, in principle, be to blame for its predictive failures should it follow that we can in principle revise any one of them to make the theory predictively successful? *Underdetermination* is insufficient to secure *Revision*: for it to do so we would need a further argument to show that for any constituent part of a theory *T*, a revision of that particular component is always available which would save the theory from further recalcitrant observations. *Underdetermination* alone does not provide any such assurances. In a lengthy exegetical rally, Adolf Grünbaum (1960) identified and routed the same fallacious reasoning in Pierre Duhem’s holistic position.⁷ Many commentators proceeded to point out that Duhem made no such mistake, but that Quine did (Ariew 1984; Giannoni 1967; Laudan 1965; Quinn 1969; Wedeking 1969; Yoshida 1975). The general point here is that although any member of a falsified theory might count as ‘falsifiable-in-principle’, as *Underdetermination* asserts, this not obviously equivalent to ‘being empirically revisable’.

More significantly, the move from *Underdetermination* to *Revision* also involves a switch from talking about the component statements of a theory being falsifiable, to a claim about whether all of our beliefs are revisable. This temptation to slide from the

context of the epistemology of science to the context of general epistemology has proved irresistible to many philosophers in many traditions, who wish to use evidentially-holistic ideas to make wider statements about epistemology. Robert Klee argues there is straightforward error here: it is a mistake to confuse *Falsification*, which makes a ‘logico-metaphysical point’, with *Revision*, which makes ‘an epistemological point’ (Klee 1992: 490), unless (perhaps like Quine) we also have prior commitments for thinking that there is nothing more to epistemology than the study of the logic of scientific inference. Similarly, Michael Resnik and Nicolleta Orlandi think that this attempt to extract *Revision* from *Falsification* is to mistake a point about deduction for one about deliberation (Resnik and Orlandi 2003: 302). More general worries about *Revision* include the idea that it is self-defeating: Jerrold Katz urges that Quine must exempt *Revision* from revision itself on pain of paradox (Katz 1998: 74), although Resnik and Orlandi (2003) and Colyvan (2006) are critical of Katz’s claim that there is a ‘paradox of revisability’. Hartry Field suggests that Quine’s position faces deeper (but related) problems concerning the empirical revision of logic (Field 1998, 2000, 2005).

The last of the evidential holistic claims we will explore here is *Values*.

Values – Our values / preferences / pragmatic considerations play a central role in determining whether a statement has been confirmed / falsified.

Values can be seen as a consequence of evidential holism, but distinct arguments suggest it emerges independently from holism; the discussions in, e.g. Kuhn 1962 and Putnam 2002, suggest non-holistic routes to a similar conclusion, and more generally there are prominent positions in feminist philosophy of science and the history of science which argue that *Values* is true *tout court*.

One way that a version of *Values* might emerge from holism as follows: if a *Global Scope* and *Non-distributivist* understanding of *Confirmation* and *Falsification* are

correct, then it is hard (and perhaps impossible) to compare the epistemic status of two theories which have different commitments but make the same predictions: both will be equally well supported by observation and experiment, and since the evidence won't tell us which parts of the theories are responsible for getting the observable evidence right or wrong, we are epistemically unequipped for deciding which of the competing sets of commitments to endorse. The suggestion is that evidential holism may contribute to what is known as the 'underdetermination of theory by evidence', whereby our evidence is insufficient to select between competing theories. If that is the case, then we may have to appeal to something else to help us to decide between theories. Philosophers have argued that the way that these stalemate situations of underdetermination can, should be, and are settled is by appeal to something other than observational evidence: scientists use 'pragmatic' considerations. For example, scientists prefer whichever theory is most simple, dovetails the best with other things, has the greatest explanatory power, is most empirically adequate, is the most fecund, is most conservatism, has the broadest scope, is the easiest to reproduce.

But why should we think that there can be empirically equivalent theories? This possibility is also supposed to be a direct consequence of holism, specifically *Revision* (see Quine (1975)). Here's Lars Bergström expressing what he takes to be Quine's line of argument:

According to holism [*Revision*], no theoretical sentence is immune to revision. Therefore, we may replace any sentence in our total system of the world by its negation; this is possible as long as we make 'drastic enough adjustments elsewhere in the system' (Quine 1953: 42). The result of such adjustments, if they are successful, would be another system of the world with the same empirical content as our original. The two systems would even be logically incompatible, since one contains the negation of a sentence in the other. (Bergström 2000: 388)

Bergström (1993), Laudan (1990), Laudan and Leplin (1991), and Okasha (2002) critically explore the attempt to link holism and underdetermination together in this way. But we should note that the argument just considered proposes *Values* as a solution to a problem which might emerge from evidential holism, rather than construing *Values* as an evidentially-holistic thesis itself.

A way to see *Values* as a holistic claim about evidence is to consider it as a *Global Scope* reading of *Prediction*: when a scientist wants to test a hypothesis, *all* of the statements making up their world view are implicated – not just claims of the scientific theory and the mathematics that it relies upon, but *everything*, which includes the scientists' metaphysical beliefs (are there composite objects? is there such a thing as human nature?), their political opinions (can there be group agency? is it always necessary to pay attention to gender?), their prejudices and their convictions.

If this point generalises, then the history of science can be reinterpreted as crucially involving those evaluative commitments. A failed experiment might be taken to falsify the experimental hypothesis, but if *Falsification* (global scope) is correct, then in fact the experiment fails to specify which of the scientist's many commitments has been falsified (this is *Underdetermination*), and it may be that they choose or prefer to lay the blame on certain statements rather than others because of their preferences, biases or metaphysical outlook. Some key instances in the history of science have suggested that, on occasion, scientists have chosen to interpret a failed experiment as resulting from a fault in some of its theoretical claims, when in fact what needed to be revised was some other part of their world-view. Examples might include the role of Lysenkoism in the development of (Lamarckian) genetic biology in the 1940s and 1950s, or sexist biases in evolutionary accounts of female orgasm (Lloyd 2005). That these researchers were not always able to notice that their values or preferences were influencing their judgements about confirmation or falsification suggests that this might be going on much more generally; that given the latitude of freedom afforded by

Underdetermination, decisions about which statements to retain and which to reject might often be fixed, in part, by personal values, preferences, etc. Lynne Hankinson Nelson's (1990) and Helen Longino's (1990) are two thorough-going book-length analyses of these lines of argument (See also: Anderson 2004 and Intemann 2001).

§4

Having seen a spread of different evidentially-holistic claims and how they might be motivated and interrelated, I'll now explore how they've been applied in other areas of philosophy by returning to look at each of the arguments I listed at the outset.

- There is no analytic / synthetic distinction

Quine's 'Two Dogmas of Empiricism' (1953) is famous for presenting his argument against the existence of a hard and fast analytic/synthetic distinction. In the paper we find Quine searching for a viable notion of 'analyticity' which will be capable of playing the epistemic role that empiricists require in their explication of scientific and mathematical knowledge. Quine eliminates several options before addressing a verificationist characterisation of analyticity. The verificationist criterion of meaning identifies the meaning of a statement with the set of observations it predicts if true. This criterion serves to demarcate meaningful statements from meaningless statements, as meaningless statements entail no observations (make no predictions). It also provides an account of analytic sentences: verificationists hold that analytic sentences are not only capable of making predictions (and are thus meaningful), but entail every observable state of affairs; analytic sentences, being logical truths, imply every consistent state of affairs. So an analytic sentence cannot be revised on the grounds of experience. Thus a verificationist account of a sentence being 'analytic' says: a sentence is analytic iff it is true regardless of evidence. The contrasting case, syntheticity, is

defined as those sentences which make predictions such that they are capable of being revised on the grounds of an experience.

So if this is the target account of analyticity, what is the case for thinking that we cannot distinguish analytic sentences from synthetic ones? Quine's argument here hinges on *Underdetermination*. Quine claims that if observational evidence alone always underdetermines which particular statement has been falsified and should be abandoned, then it remains in principle possible to exempt any one particular statement from ever being falsified / revised. He suggests that any purportedly synthetic statement can permanently resist falsification; we just need to lay the blame on other parts of the theory and "make drastic enough changes in the rest of the system. Even a statement very close to the periphery can be held true in the face of recalcitrant experience by pleading hallucination or by amending certain logical laws." (1953: 43). But now the contrast between analytic and synthetic statements cannot be drawn, since both 'synthetic' and 'analytic' statements can permanently resist empirical revision. If any statement can be one which can be retained regardless of what we observe, then there is no longer a dichotomous distinction between analytic statements and synthetic statements.

An importantly distinct way that Quine gets to the same conclusion does not involve showing that any statement whatsoever can qualify as 'analytic' (held true regardless of the evidence), but instead attempts to show that any statement can qualify as 'synthetic' (revised on the grounds of an experience). If *Revision* is true then any statement whatsoever could be revised on the basis of empirical experience, including purportedly analytic sentences like 'all sisters are siblings'.

The argument to show that, in principle, any statement can be treated as if it is analytic (revision-resistant) relies on *Underdetermination* and *Global Scope*. And the argument to show that, in principle, any statement can be treated as if it is synthetic (empirically

revisable) relies on *Revision* and *Global Scope*. Without global scope, it would no longer be correct for either of these arguments to maintain that *any* statement can be treated as if it is analytic/synthetic, since a more moderate holism might deny that *all* of our beliefs enter into these evidential relations. Both arguments also rely on a verificationist construal of the analytic/synthetic distinction, which makes the meaning of a statement strictly an epistemic or evidential matter. Figure 2 illustrates the main components of these arguments:

Figure 2.

The next application of evidential holism is for the following conclusion:

- The meaning of a sentence cannot be understood without understanding the whole language of which it is a part

Arguments for ‘semantic holism’ also appeal to verificationism and the dissolution of the analytic/synthetic distinction. Semantic holism is a claim about the metaphysically necessary conditions for a sentence to have the meaning that it does, and it states that the meaning of a particular sentence is constituted by its relations to all the other sentences in the theory of which it is a part. (See Block 1994, 1998; Fodor and Lepore 1992, 1993a). In contrast, semantic atomism maintains that meaning of a sentence is not constituted by any of its relations to any other sentences. In the middle-ground between atomism and holism is semantic molecularism, which holds that the meaning of a sentence is constituted by its relations (usually inferential relations) to some but not all of the other sentences in the theory of which it is a part. Semantic molecularists (such as Michael Devitt (1996) and Georges Rey (1993)) claim that this is because the inferential relationships between ‘cat’ and ‘animal’ partly constitute their meanings, and they take our inferential abilities (such as the ability to infer ‘there is an animal on the mat’ from ‘there is a cat on the mat’) to be explained by and to provide evidence for the claim that meaning is inferential role. But molecularists deny that the meaning of any one token sentence is dependent on its relations to all the other sentences in the language.

The standard ‘Quinean’ form of argument from evidential holism to semantic holism comes from combining *Prediction* and a verificationist criterion of meaning (see Okasha 2000). Verificationism holds that the meaning of a statement is the set of observations it predicts if true, but if *Prediction* is correct then individual statements never make predictions by themselves, so either semantic molecularism or holism results. The difference will depend on whether a *Global Scope* reading of *Prediction* is endorsed or not. This is illustrated in figure 3.

Figure 3.

There is another line of argument for semantic holism which also involves evidential holism, only somewhat more indirectly. This argument relies on Quine's case against there being an analytic/synthetic distinction, a case which we have seen to rely on *Underdetermination* and *Revision*. The argument for holism takes the form of a dilemma: the first horn of the dilemma requires semantic molecularists to specify why some statements contribute to the meaning of others, but others do not; molecularists must give 'principled grounds' which specify when the connections between statements are meaning-constituting and when they are not. But if there is no analytic/synthetic distinction then molecularists cannot endorse any meaning-constituting relations. The second horn of the dilemma is that molecularists must deny that there can be a criterion of meaning-constitutivity, but this is supposed to generate a 'slippery slope' to semantic holism (see Berg 1993, Devitt 1993, 1996, Fodor and Lepore 1992, Rey 1993, 1997). A verificationist principle of meaning seems to be a constituent part of all of these arguments, and verificationism is not without its critics.

We turn now to:

- All knowledge is empirical knowledge (there is no apriori knowledge)

The argument against the existence of apriori knowledge is as follows (versions of which can be found in Bealer 1992; BonJour 1998; Devitt 1998, 2005; Field 1998, 2000 and Godden 2006):

1. A necessary condition for a statement to be known apriori is that it is empirically indefeasible.
2. *Revision* – any belief can be revised on the basis of empirical experience.
3. Therefore no propositions are known apriori.

Not everyone agrees with the first premise, and this issue has been much discussed. Some maintain that respectable analyses of apriori knowledge should not make empirical indefeasibility a necessary condition (e.g. Rey 1998, Jenkins 2008) – to say that a belief is justified ‘independently of experience’ is just to say that it can be justified even though it has not received any empirical support, and that the further demand that its justificatory status should be impervious to empirical updating is a mistake. Tyler Burge (1993, 1998) thinks that it is important for accounts of apriori justification to track a property which is not equivalent to empirical indefeasibility. Laurence BonJour thinks that the important feature of a priori knowledge is the positive claim about how the beliefs do get to be justified, e.g. “justification that derives from pure thought or reason alone with no positive dependence on experience” (BonJour 1998). So both Burge and BonJour think that the set of empirically indefeasible beliefs and the set of apriori justified beliefs need not be coextensive. Others disagree: Hartry Field has argued that if the justification of a belief depends on whether we have not had any defeating experiences, then that belief’s justification does depend on experience; as such, a belief which is empirically defeasible is not one which can be justified ‘independently of experience’ (see Field 1996, 1998, 2000, 2005). For

our purposes, we can note that the evidentially holistic component of the argument is *Revision* which is supposed to be a consequence of *Underdetermination*.

- Abstract mathematical objects (such as numbers, sets, and functions) exist

The evidentially holistic argument for the conclusion that abstract mathematical objects exist is often called the ‘holistic indispensability argument’. The most widely recognised version of the ‘holistic’ indispensability argument invokes both *Prediction* thesis and *Confirmation*, along with the provision that the ‘scope’ is wide enough to include mathematical statements within the ‘collective bodies of statements’ that scientists are epistemically dependent upon. The argument usually proceeds as follows:

1. In order to make testable experimental predictions, scientists depend upon a whole body of theory (including mathematical statements) [*Prediction*].
2. Those tests confirm the whole body of statements, so the mathematical claims are confirmed along with the non-mathematical claims. [*Confirmation*].
3. Ontological commitment: sentences which say ‘there is an X’ are ontologically committed to the existence of Xs.
4. Scientific statements are ontologically committed to the existence of the entities that they quantify over (such as electrons, muons, genes, species), and mathematical statements are ontologically committed to the existence of mathematical objects (such as numbers, sets, functions etc)
5. So the existence of mathematical objects is as well confirmed as that of the other entities that science tells us about.

It’s widely believed that this is (to a rough approximation) an argument that Quine endorsed (see Baker 2005, Bangu 2012, Colyvan 2001, Field 1980, 1989, Liggins 2008, Maddy 1990, 1992, 2005, Resnik 1995, Sober 1993). The holistic premises of

the argument have been directly challenged (Sober 1993, Maddy 2005, Morrison 2012), and the debate has moved on to other versions of the indispensability argument which do not appeal to evidential holism (but for a resistant view which attempts to reveal a continued underlying commitment to holism, see Marcus 2014, and the reply by Knowles and Liggins 2015).

- There is no fact / value distinction

In so far as science is presumed to deal with something like ‘facts’ in isolation from anything like ‘values’, the import of the holistic thesis *Values* is to suggest that the boundary between the two is nowhere near as easy to draw. What may seem to be a scientific fact could turn out to implicitly have involved many of the scientists’ ‘values’. Apart from *Values*, the argument hinges on a) the extent to which scientific methods and inferences are taken as exemplary instances of how we come to grasp ‘facts’, and b) the concern that the intrusion of ‘values’ into scientific method “will impede the free acquisition of knowledge and the correct representation of (independent) reality” (Lloyd 1995: 354). Neither claim is uncontroversial (see Douglas 2000, 2004, Longino 1990, and Putnam 2002).

Concluding remarks

In §1 I separated out several distinct ways that evidence can be holistic, and in §2 and §3 I’ve explored how they might be motivated and how they might relate to each other. §4 introduced some prominent arguments in epistemology, philosophy of science, metaphysics, the philosophy of mind and language and beyond which have relied on evidential holism, and in each case I’ve identified which of the holistic theses from §1 are involved. The arguments that appear in §4 are not the only applications of evidential holism. For example, this survey could also have looked at the way that evidential holism appears in arguments for moral particularism, how it relates to

coherentism in epistemology, or any of the many instances where the image of Quine's 'web of beliefs' has been used to prop up relativistic conclusions. But because §2 and §3 have done some sorting out of the galaxy of claims that might constitute evidential holism, it should be possible for the reader to identify which kinds of holistic claims are being appealed to in further cases, and thereby to make an evaluation about whether those arguments are relying on well-motivated and defensible holistic theses, or, alternatively, an implausibly strong or undefended type of evidential holism.

Works cited

- Anderson, E. (2004), "Use of Value Judgments in Science". *Hypatia*, 19 (1), pp. 1-24.
- Ariew, R. (1984), "The Duhem Thesis", *British Journal for the Philosophy of Science*, 35 (4), pp.313-325.
- Bangu, S. (2012) *The Applicability of Mathematics in Science: Indispensability and Ontology*, Palgrave Macmillan
- Bealer, G. (1992), "The Incoherence of Empiricism", *Proceedings of the Aristotelian Society*, 66, pp.99-138.
- Berg, J. (1993), "Inferential Roles, Quine, and Mad Holism", in Fodor and Lepore 1993, pp.283-302.
- Bergström, L. (1993), "Quine, Underdetermination and Skepticism", *Journal of Philosophy*, 90 (7), pp.331-358.
- Bergström, L (2000) "Quine" in *A Companion to the Philosophy of Science* ed W.H. Newton-Smith 385-290 Blackwell
- Block, N. (1994), "An Argument For Holism", *Proceedings and Addresses of the American Philosophical Association*, 94, pp.151-69.
- Block, N. (1998), "Holism: Mental and Semantic" in *The Routledge Encyclopaedia of Philosophy*, Craig (Ed.), London: Routledge.

- BonJour, L. (1998), *In Defense of Pure Reason - A Rationalist Account of A Priori Justification*, Cambridge: Cambridge University Press.
- Burge, T. (1993), "Content Preservation", *Philosophical Review*, 102 (4), pp.457-488.
- Burge, T. (1998), "Computer Proof, Apriori Knowledge and Other Minds", *Philosophical Perspectives*, 12, pp.1-37.
- Christensen, D. (1992) "Confirmational Holism and Bayesian Epistemology," *Philosophy of Science*, 59, pp.540-557.
- Colyvan, M. (2001), *The Indispensability of Mathematics*, Oxford: Oxford University Press.
- Colyvan, M. (2006), "Naturalism and the Paradox of Revisability", *Pacific Philosophical Quarterly*, 87 (1), pp.1-11.
- Devitt, M. (1993), "A Critique of the Case for Semantic Holism", in Fodor and Lepore 1993 pp.17-60.
- Devitt, M. (1996), *Coming To Our Senses: A Naturalistic Case for Semantic Localism*, Cambridge: Cambridge University Press.
- Devitt, M. (1998), "Naturalism and the A Priori", *Philosophical Studies*, 92, pp.45-65.
- Devitt, M. (2005), "There Is No A Priori", in Steup and Sosa (Eds.), *Contemporary Debates in Epistemology* Vol. 3, Oxford: Blackwell Publishing, pp.105-115.
- Douglas, H. (2000) "Inductive Risk and Values in Science", *Philosophy of Science* 67, pp.559-579
- Douglas, H. (2004) "The Irreducible Complexity of Objectivity", *Synthese*, Vol. 138, No. 3, pp.453-473
- Duhem, P. (1954), *The Aim and Structure of Physical Theory* (Wiener, Trans.), Princeton: Princeton University Press.
- Earman, J. (Ed.), (1983), *Testing Scientific Theories*, Minneapolis: University of Minnesota Press.
- Esfeld (1998), "Holism and Analytic Philosophy" *Mind*, Vol. 107, No. 426, pp. 365-380
- Field, H. (1980), *Science Without Numbers*, Oxford: Basil Blackwell.

- Field, H. (1989), "Realism, Mathematics and Modality", in Field (Ed.), *Realism, Mathematics and Modality* (Oxford: Basil Blackwell), pp.227-281.
- Field, H.(1996), "The A Prioricity of Logic", *Proceedings of the Aristotelian Society*, 96, pp.359-379.
- Field, H. (1998), "Epistemological Nonfactualism and the A Prioricity of Logic", *Philosophical Studies*, 92 (1-2), pp.1-24.
- Field, H. (2000), "Apriority As An Evaluative Notion", in Boghossian and Peacocke (Eds.), *New Essays on the A Priori*, Oxford: Oxford University Press, pp. 117-149.
- Field, H. (2005), "Recent Debates About The A Priori", in Gendler and Hawthorne (Eds.), *Oxford Studies in Epistemology* Vol. 1, Oxford: Clarendon Press, pp. 69-88.
- Fitelson, B., & Waterman, A. (2005), "Bayesian Confirmation and Auxiliary Hypotheses Revisited: A Reply to Strevens", *British Journal for the Philosophy of Science*, 56, pp.293-302.
- Fitelson, B., & Waterman, A. (2007), "Comparative Bayesian Confirmation and the Quine-Duhem Problem: A Rejoinder to Strevens", *British Journal for the Philosophy of Science*, 58 (2), pp.333-338.
- Fodor, J., & Lepore, E. (1992), *Holism: a Shopper's Guide*, Oxford, UK: Blackwell Publishers.
- Fodor, J., & Lepore, E.(1993), Eds. *Holism: A Consumer Update*, Amsterdam: Editions Rodolphi.
- Gemes, K. (1994), "Schurz on Hypothetic-Deductivism", *Erkenntnis*, 41, pp.171-181.
- Gemes, K. (1998), "Hypothetico-Deductivism: The Current State of Play", *Erkenntnis*, 49, pp.1-20.
- Giannoni, C. (1967), "Quine, Grünbaum and the Duhemian Thesis", *Nous*, 1, pp. 283-298.
- Gibson, R. F. (1982), *The Philosophy of W. V. Quine: An Expository Essay*, Tampa: University Presses of Florida.

- Glymour, C. (1975), "Relevant evidence", *The Journal of Philosophy*, 72.14, pp. 403-426.
- Glymour, C. (1980), *Theory and Evidence*, Princeton, New Jersey: Princeton University Press.
- Godden, D. (2006), "Departmental boundaries within the corporate body of theory: Quine on the holistic foundations of logic." *Dialogue* 45.03, pp.505-528.
- Grünbaum, A. (1960), "The Duhemian Argument", *Philosophy of Science*, 27 (1), pp. 75-87.
- Hankinson Nelson, L. (1990), *Who Knows: From Quine to a Feminist Empiricism*, Philadelphia: Temple University Press.
- Hempel, C. G. (1965), *Aspects of Scientific Explanation*, New York: The Free Press.
- Horwich, P. (1983), "Explanations of Irrelevance", in Earman 1983, pp.55-65.
- Howson, C., & Urbach, P. (1993), *Scientific Reasoning: The Bayesian Approach* (2nd ed.), Chicago: Open Court.
- Intemann, K. (2001), "Science and values: Are Value Judgments Always Irrelevant to the Justification of Scientific Claims?", *Philosophy of Science*, 68 (3), pp. S506-S518.
- Jenkins, C. S. (2008). "A Priori Knowledge: Debates and Developments", *Philosophy Compass* 3, pp. 436-50.
- Katz, J. J. (1998), *Realistic Rationalism*, Cambridge, Massachusetts: The MIT Press.
- Klee, R. (1992), "In Defense of the Quine-Duhem Thesis: A Reply to Greenwood", *Philosophy of Science*, 59, pp.487-491.
- Knowles, R. and Liggins, D. (2015). "Good weasel hunting." *Synthese* 192.10 (2015): 3397-3412.
- Kuhn, T. (1962), *The Structure of Scientific Revolutions*, Chicago, Illinois: The University of Chicago Press.

- Laudan, L. (1965), "On the Impossibility of Crucial Falsifying Experiments: Grünbaum on 'The Duhemian Argument'", *Philosophy of Science*, 32, pp. 295-299.
- Laudan, L. (1990) "Demystifying Underdetermination", in C. Savage (Ed.) *Scientific Theories, Minnesota Studies in the Philosophy of Science*, Volume 14, University of Minnesota Press (Minneapolis) pp267 – 297.
- Laudan, L. and Leplin, J. (1991) "Empirical Equivalence and Underdetermination", *The Journal of Philosophy*, Vol. 88, No. 9, pp. 449-472.
- Liggins, D. (2008). Quine, Putnam, and the 'Quine-Putnam' indispensability argument. *Erkenntnis*, 68 (1), 113–127
- Lloyd, E. (1995) 'Objectivity and the Double Standard for Feminist Epistemologies', *Synthese* 104, 351-381
- Lloyd, E. (2005) *The Case of the Female Orgasm: Bias in the Science of Evolution*, Harvard University Press (Cambridge, Massachusetts).
- Longino, H 1(1990) *Science as Social Knowledge*, Princeton University Press, Princeton.
- Maddy, P. (1990), *Realism in Mathematics*, Oxford: OUP.
- Maddy, P. (1992), "Indispensability and Practice", *The Journal of Philosophy*, 89, pp. 275-289.
- Maddy, P. (2005), "Three Forms of Naturalism" in Shapiro (Ed.), *Oxford Handbook of Philosophy of Mathematics and Logic*, Oxford: OUP, pp.437-459).
- Mayo, D. G. (1996) *Error and the Growth of Experimental Knowledge*, Chicago: University of Chicago Press.
- Mayo, D. G. (1997) "Duhem's Problem, the Bayesian Way, and Error Statistics, or "What's Belief Got to Do with It?"", *Philosophy of Science*, Vol. 64, No. 2, pp. 222-244
- Morrison, J. (2010) "Just How Controversial Is Evidential Holism?" *Synthese*, vol. 173, issue 3, pp. 335-352.
- Morrison, J. (2012) "Evidential Holism and Indispensability Arguments", *Erkenntnis*, vol. 76, number 2, pp. 263-278.

- Okasha, S. (2000), "Holism about meaning and about evidence: in defence of W.V. Quine", *Erkenntnis*, 52, pp.39-61.
- Okasha, S. (2002), "Underdetermination, Holism and the Theory/Data Distinction", *Philosophical Quarterly*, 52 (208), pp.303-319.
- Orenstein, A. (1977), *Willard Van Orman Quine*, Boston: Twayne.
- Popper, K. R. (1959), *The Logic of Scientific Discovery*, Hutchinson (London)
- Popper, K. R. (1963), *Conjectures and Refutations*, London: Routledge & Kegan Paul.
- Putnam, H. (2002), *The Collapse of the Fact / Value Dichotomy*, Cambridge, Mass; London: Harvard University Press.
- Quine, W. V. (1953), "Two Dogmas of Empiricism", Reprinted in *From a logical point of view*, 2nd edition, Cambridge, MA: Harvard UP, pp. 20–46.
- Quine, W. V. (1975), "On Empirically Equivalent Systems of the World", *Erkenntnis*, 9, pp.313-328.
- Quine, W. V. (1981), "Five Milestones of Empiricism", in Quine (Ed.), *Theories and Things*, Cambridge, Mass.: Harvard University Press, pp.67-72.
- Quinn, P. L. (1969), "What Duhem Really Meant", in *Boston Studies in the Philosophy of Science*, vol. XIV, Dordrecht: Reidel, pp.33-56.
- Resnik, M. D. (1995), "Scientific vs. mathematical realism: the indispensability argument", *Philosophia Mathematica*, 3 (2), pp.166-174.
- Resnik, M. D., & Orlandi, N. (2003), "Holistic Realism: A Response to Katz on Holism and Intuition", *The Philosophical Forum*, 34 (3 & 4), pp.301-316.
- Rey, G. (1993), "The Unavailability of What we Mean. A Reply to Quine, Fodor and Lepore" in Fodor and Lepore 1993, pp.61-102.
- Rey, G. (1997), "Devitt's Naturalism: A Priori Resistance to the A Priori?", in Jutronic (Ed.), *The Maribor Papers in Naturalized Semantics* (Maribor: Pegagoska fakulteta Maribor), pp.141-154.
- Schurz, G. (1991), "Relevant Deduction", *Erkenntnis*, 35, pp.391-437.
- Schurz, G. (1994), "Relevant Deduction and Hypothetico-Deductivism: A Reply to Gemes", *Erkenntnis*, 41, pp.183-188.

- Sober, E. (1993), “Mathematics and Indispensability”, *Philosophical Review*, 102 (1), pp.35-57.
- Sober, E. (1999), “Testability”, *Proceedings and Addresses of the American Philosophical Association*, 73, pp.47-76.
- Sober, E. (2000), “Simplicity” in Newton-Smith (Ed.), *A Companion To The Philosophy Of Science* Vol. 18, Oxford: Blackwell Publishers, pp.433-441.
- Strevens, M. (2001), “The Bayesian Treatment of Auxiliary Hypotheses”, *British Journal for the Philosophy of Science*, 52, pp.515-537.
- Strevens, M. (2005), “The Bayesian Treatment of Auxiliary Hypotheses: A Reply to Fitelson and Waterman”, *British Journal for the Philosophy of Science*, 56, pp. 913-918.
- Thompson, N. (forthcoming) “Metaphysical Interdependence, Epistemic Coherentism and Holistic Explanation”
- Wedeking, G. (1969), “Duhem, Quine and Grünbaum on Falsification”, *Philosophy of Science*, 36 (4), pp.375-380.
- Weir, A. (1985), “Against Holism”, *Philosophical Quarterly*, Vol. 35, No. 140, pp. 225-244.
- Yoshida, R. M. (1975), “Five Duhemian Theses”, *Philosophy of Science*, 42 (1), pp. 29-45.

¹ Evidential relations may include ‘confirms’, ‘disconfirms’, ‘falsifies’, ‘predicts’, ‘explains’.

² I take it that evidential holism is properly neutral with respect to the competing ‘syntactic’ and ‘semantic’ views of theories. Roughly speaking, on the syntactic view a theory is a bit like a large conjunction of sentences, while on the semantic or ‘model-theoretic’ view a theory is not linguistic, but is rather a family of models which satisfy theoretical laws; a scientific theory is not a body of axiomatic theorems or laws, but instead is what those formulations refer to when they’re given a formal semantic interpretation. Some of the problems with the syntactic view are highlighted by some consequences of holism, while the semantic view of theories attempts to encompass those holistic insights. For example, since the semantic view says that there’s nothing epistemologically significant about an observable/unobservable distinction, or an observation/theoretical distinction for sentences, it’s seemingly consistent with some of holism’s definitional claims.

³ This paper does not include discussion of ‘explanatory holism’, which might be a claim something like:

Explanation: The unit of explanation is collective bodies of statements; individual statements do not explain things, only bodies of statements do.

I’ve omitted it for a couple of reasons. One, it doesn’t feature very prominently in subsequent discussions in the philosophy of science / epistemology, and hasn’t been applied in any significant role in other areas of philosophical argument. Two, it could be quite easily confused with a different kind of ‘explanatory holism’: holism about ‘metaphysical explanation’ or ‘grounding relations’. Metaphysical explanation is not primarily an evidential relation (although it might well underwrite the evidentially-explanatory relation), so the two shouldn’t be confused. See Thompson (forthcoming).

⁴ Thanks to an anonymous reviewer for *Philosophical Compass* for suggesting that I mention this point explicitly.

⁵ I suspect that this is largely due to two highly questionable but prevalent assumptions: (i) that ‘prediction’ is just a trivial form of either deductive or probabilistic inference, (ii) that any attempt to spell out criteria for how or whether a theory or statement ‘predicts’ some empirical consequences is tantamount to an attempt to give a ‘criterion of empirical significance’, which is taken to be a fool’s errand.

⁶ While the text might seem to imply that the only motivation for *Confirmation* is a combination of *Prediction* and something like the symmetry principle, it remains possible that that *Confirmation* could be established on the basis of other things (see Morrison 2010). For example, an argument could be made from some sort of semantic holism to the truth of *Confirmation* – this seems to be the kind of motivation for *Confirmation* that Fodor and Lepore ascribe to Quine (Fodor and Lepore 1992, p. ch1), but see (Okasha 2000).

⁷ The commentators listed in the text refer to this issue as ‘the problem of the missing existential proof’. Suppose T is a set composed of statements {F, G, H ...}, and T predicts E, but E is false. It follows that T is falsified, and thus that the set {F, G, H ...} is falsified, while the falsification of any particular member of that set remains undetermined by ~E alone. But to get to *Revision* we need a proof to show that for any member of T, there always exists a revision of that statement such that exchanging, say, F for F* will always result in a new set which will be predictively successful. The problem is that no such proof is available, and none seems possible.