

**Objectivity in Confirmation:
Post Hoc Monsters and Novel Predictions**

Ioannis Votsis

votsis@phil.hhu.de

1. Introduction

The study of confirmation is the study of the conditions under which a piece of evidence supports, or ought to support, a hypothesis as well as of the level of that support. There are two major kinds of confirmation theories, objective and subjective. Objective theories hold that confirmation questions are settled by purely objective considerations. Subjective ones hold that at least some non-objective considerations come into play. With some exceptions (see, for example, Williamson 2010), most confirmation theorists nowadays opt for subjective theories. The pessimism over objective theories is most probably due to the fact that it has proved very hard, some may even say impossible, to find reasonable principles that decide questions about confirmation in purely objective terms. The aim of this paper is to reverse some of that pessimism by putting in place some cornerstones in the foundations for an objective theory of confirmation. This is achieved by considering lessons not from the failures of subjective theories, of which, no doubt, there are many, but rather from the failures of a certain kind of mini-theory of confirmation, namely predictivism, that is typically conceived of as objective.

2. The Completion Challenge

Imagine a scientist S who endorses a hypothesis H but is confronted with incontrovertible evidence E that contradicts H . Unless S is a defeatist, two options seem to be available. S can attempt to either modify H or else construct an entirely new hypothesis. Whatever the chosen option, S will have to ensure that the hypothesis endorsed stands in the right kind of inferential-semantic relations, ideally the entailment of true propositions, to the wayward and other established empirical results. More subtle options become available if the example is modified a little. Suppose that $\neg E$ is not entailed by H alone but by a cluster of claims which includes H . Duhem ([1914] 1991) famously argued that this kind of situation is not the exception but the unexceptional rule. Hypotheses, according to him, cannot be tested in isolation, for they have no testable consequences on their own. To derive such consequences additional assumptions, the so-called ‘auxiliaries’, are needed, e.g. assumptions about initial and boundary conditions. Thus when evidence contradicts a given consequence, it may not be immediately obvious where the blame lies. In the case under consideration, all we know is that at least one of the claims in the cluster is the culprit. This opens up the available options. S can attempt to modify or replace (anew) one or more of the existing claims in the cluster. Hereafter, and for simplicity, I will often talk of a hypothesis being confirmed or disconfirmed instead of a cluster of hypothesis plus auxiliaries although the latter is intended.

A question that emerges at this point is whether E confirms the new hypothesis or the modified one regardless of how S went about making sure that E stands in the right kind of inferential-semantic relations to the chosen hypothesis. Those who answer this question in the affirmative endorse the view that standing in the right kind of inferential-semantic relations, e.g. the logical entailment of true propositions, is sufficient for confirmation.¹ Various confirmation theorists nowadays consider this view false. The contrary view is motivated by the ease with which one can arrange a hypothesis to stand in the right kind of inferential-semantic relations to the evidence. Thus Worrall complains:

If having the right empirical consequences is the only criterion [for confirmation], then since any core idea can be incorporated into a theoretical system that has the right consequences, there is no empirically-based rational preference for any such core idea over any other (2002, p. 193).

¹ We assume that standing in the right kind of inferential relations is at least necessary for confirmation.

Laudan and Leplin (1991) even go as far as to claim that “[n]o philosopher of science is willing to grant evidential status to a result e with respect to a hypothesis H just because e is a consequence of H ” (p. 466). Echoing these complaints but targeting not just logical entailment but also probabilistic relations, Hitchcock and Sober note: “*Accommodation is easy*. It is always possible, after the fact, to come up with some hypothesis or other that accommodates a given body of data” (2004, p. 6) [original emphasis]. Various types of logical and/or mathematical manipulations can be recruited to modify auxiliaries or hypotheses post hoc so as to bring about the appropriate inferential-semantic connection to the evidence. One well-known type of post hoc manipulation can be found in the domain of curve fitting. To account for a wayward datum, one can always modify or create a polynomial which includes a term that guarantees the curve passes through or sufficiently near the datum in addition to it passing through or sufficiently near existing data.

Because accommodation, either by logical entailment or by probabilistic inference of an appropriate strength, can be obtained relatively cheaply, various confirmation theorists agree that it cannot be the whole story about confirmation. Moreover, they agree that a big part of the story concerns what a theory of confirmation ought to say about post hoc constructed or modified hypotheses. That this is not a mere philosophical quibble can be illustrated by considering the various undesirable hypotheses put forward as a consequence of post hoc manipulations, e.g. Ptolemaic astronomy. We can call such hypotheses ‘post hoc monsters’.² A central challenge facing confirmation theorists then appears to be the discovery of what is needed in addition to the right kind of inferential-semantic relations in order to attain a complete account of confirmation that at the same time tackles the vexing issue of whether or not, and if so to what extent, post hoc monsters can be confirmed. We can call this the ‘completion challenge’.

3. Predictivism

The requirement that, in addition to standing in the right kind of inferential-semantic relations to the evidence, hypotheses make predictions is a way to meet the completion challenge, either in part or in full. According to this movement, accommodated evidence and, for some of its theorists, even evidence that could have been accommodated, i.e. evidence that could have been used in post hoc constructions, is somehow inferior to predicted evidence, or, as such predictions are sometimes called, ‘novel predictions’. For obvious reasons this movement has come to be called ‘predictivism’ and it is contrasted to the ‘accommodationism’ movement which denies any such inferiority. Predictivist views are typically not intended as complete theories of confirmation but instead as illuminating one aspect of it. When this happens the predictivist requirement offers only a partial way to meet the completion challenge. It is also worth noting that there are weak and strong versions of predictivism. Let us call a predictivism ‘weak’ when it holds that predictions, as opposed to mere accommodations, confer greater support to the hypotheses that issue them. By contrast, let us call a predictivism ‘strong’ when it holds that predictions are the *only* source of support.³ Hereafter, I will express the fact that predictivist views can be given both a strong and a weak formulation by presenting them side by side, with the weak formulation appearing in parentheses.

Most predictivist theories are objective, though tacitly so. The objective character of the two most prominent kinds of predictivist views, i.e. temporal novelty and use-novelty, will become clear in the sections that follow. Before we continue to those sections, it is worth considering, if only briefly, an

² I offer an explication of the notion of ‘monstrous hypotheses’ in Section 10. Until then the readers will have to rely on the examples given and their own intuitions.

³ William Whewell (1847) is often cited as the father of predictivism. Hitchcock and Sober (2004, §2) as well as Harker (2008) and Barnes (2008) provide a taxonomy of predictivism along a number of axes. My use of the terms ‘strong’ and ‘weak’ differs from that of the cited authors. For them, weak predictivism is the view that predictions have a special confirmational role only because they indicate the presence of some other epistemic virtue in the hypothesis that issues them, e.g. simplicity. Strong predictivism, for them, is the view that predictions have this special role only because prediction is intrinsically superior to accommodation.

alternative kind of predictivist view, one that appears to be subjective. I have here in mind the ‘endorsement novelty’ view found in Barnes (2008). According to this weak predictivist view, “when true evidence N confirms T, endorsed by X, T is more strongly confirmed (for some evaluator) when N is endorsement-novel relative to X than when it is not” (p. 37). Roughly speaking, a piece of evidence is endorsement-novel relative to X if X endorses T (to a sufficient degree so as to not arouse scepticism about her commitment) without recourse to observations relating to N’s truth. Given Barnes’ suggestion – see his contribution to this issue – that the endorser determines the precise level of their endorsement on partly subjective grounds, it seems reasonable to construe his view as a subjective theory. That’s about as much as I will say about this view.⁴

4. Temporal Novelty

One particular brand of predictivism is temporal novelty. According to this view, support emanates solely (or more plentifully) from phenomena that become known after the hypothesis, plus any auxiliaries, that predicts them was formulated or modified.⁵ That is, it emanates from phenomena that are temporally novel. This puts the view squarely in the objective theory of confirmation camp, as whether or not a phenomenon is novel depends on historical facts. The rationale behind the view is simple. A hypothesis, plus any auxiliaries, could not possibly have been shaped to accommodate phenomena that were unknown prior to its formulation or modification. This obviously rules out the dreaded post hoc manoeuvres.

Using the temporal novelty account we can now make quick work of several undesirable hypotheses. Take the creationist hypothesis that the world was created in 4004BC. Since creationism (and its auxiliaries) never seems to predict any phenomena but merely accommodates them, it does not, according to the temporal novelty account, earn any support from them. The theory of evolution, by contrast, successfully predicts a diverse range of phenomena. Among other things, it predicts the existence of transitional organisms, i.e. organisms that exhibit features of both older and newer species. Together with auxiliaries concerning the process of fossilisation, the theory of evolution thus predicts the existence of fossilised remains corresponding to the transitional organisms. Many such fossils have been discovered since the theory was first formulated. The pertinent phenomena thus count as temporally novel. Beyond the creationist-evolutionist dispute, temporal novelty offers a prima facie plausible explanation to the confirmational boost given by a number of well-known predictions to their respective theories, e.g. Fresnel’s wave theory of light and the Poisson spot, Newtonian physics and the discovery of the planet Neptune, Mendeleev’s periodic table and the discovery of the elements gallium, germanium and scandium, etc.

In spite of some advantages, the account has been the subject of much criticism.⁶ One objection concerns the account’s reliance on contingent considerations. As Worrall puts it:

Why on earth *should* the apparently purely contingent historical issue of whether or not a theory was first developed before some particular piece of evidence became available matter at all in an account of the *rational* support that that evidence lends to theory? (2002, p. 194) [original emphasis].⁷

What Worrall fails to see is that, presented thus, the objection begs the question against the temporal novelty advocates. For, clearly, such advocates believe that a certain kind of contingent consideration is confirmationally relevant. Their belief is not entirely without reason. Recall that the rationale for temporal novelty is that phenomena that were unknown prior to the formulation or

⁴ For a critique see Glymour (2008).

⁵ Sometimes the view is loosened to include as novel phenomena those that were not *widely* known to the scientific community.

⁶ Temporal novelty has few supporters. Duhem ([1914] 1991), the early Lakatos and Maher (1988) are sometimes identified as supporters, though not without controversy.

⁷ See also Musgrave (1974) for a discussion of the problems afflicting the temporal novelty account.

modification of a hypothesis could not possibly have been used to post hocly give it a desirable shape.

Happily, Worrall does not fail to see another, this time genuine, objection. The temporal novelty view rules out the confirmation of post hoc monsters but in so doing it rules out too much. Any evidence gathered before the postulation or modification of a hypothesis is automatically dismissed as incapable of providing support (or as much support) to that hypothesis. This holds even when the hypothesis in question is, by anybody's count, not a monster! The general theory of relativity is arguably far from a monstrous theory. One of the most famous pieces of evidence in its favour, the precession of Mercury's perihelion, was known long before the postulation of the theory. In fact, temporal novelty 'goes one worse' as it promotes the view that even in those cases where already known evidence had not actually been employed to post hocly shape a hypothesis, the mere counterfactual possibility of it being so employed is sufficient to demote or eradicate its confirmational value.

5. Use-Noveltly

At least some data that were known prior to the formulation or modification of a given hypothesis, it seems, must be capable of providing support (or as much support) for that hypothesis.⁸ What we need then is a criterion that tells us when data possess this ability. The use novelty approach to predictivism, a.k.a. the 'heuristic' or the 'no double counting' approach, was formulated with this aim in mind. In generic terms, a use novelty account of confirmation holds that data used in the construction or modification of a hypothesis cannot support that hypothesis (or, at best, support it less than non-use-constructed data). Loosely speaking, novelty here is understood in terms of the unexpectedness that a hypothesis stands in the right kind of inferential-semantic relations to a set of data that was not used in its construction. The rationale behind the use novelty approach is that a hypothesis could not possibly have been shaped to accommodate (known or unknown) data if that data were not used in its construction. This obviously rules out the aforementioned post hoc manoeuvres.

One prominent case of use novelty is the precession of Mercury's perihelion and the general theory of relativity.⁹ Several other cases can be found in the literature. For example, Musgrave (1974, p. 11) lists Galileo's and Kepler's laws, facts about tides and the precession of the equinoxes in relation to Newton's theory, the Michelson-Morley experiment in relation to the special theory of relativity and Balmer's empirical formulas in relation to Bohr's theory of the hydrogen atom. Since phenomena unknown prior to the formulation or modification of a theory could not have been used in its construction this means that all temporally novel predictions are also use-novel ones. For this reason historical examples of temporally novel predictions are also examples of use-novel predictions.

A number of use-novelty views have been proposed over the years. I will concentrate on one of them, returning briefly to the rest in Section 8. Worrall's view, in particular as it is developed in his (2002; 2005; 2006), presents the most sophisticated and promising attempt to codify use-novelty.¹⁰ His strong version of predictivism attempts to tweak the notion of 'novel prediction' so as to make it insensitive to temporal issues. To achieve the desired insensitivity Worrall turns to logic, conceiving of confirmation relations in deductive terms. To be precise, he applies this conception only to cases

⁸ Up to now, I have used the concepts *evidence*, *phenomena* and *observations* interchangeably. I'm here adding the concept *data* to that list. I am fully aware of the controversy of using these terms interchangeably. For the sake of expedience, and following others in the novel predictions debate, I shall assume that my use is unproblematic, though strictly speaking I endorse the view that there are differences between some of these concepts, e.g. between phenomena and data. For more on this topic see Votsis (2011).

⁹ This example is not uncontroversial. Earman and Glymour (1978) cite two letters from Einstein, one to Sommerfeld and one to Lorentz, where he asserts that adequately accounting for the precession of Mercury's perihelion is a consideration that he used in selecting between different versions of his theory.

¹⁰ For an earlier version of his view see his (1985).

of deterministic theories in science, intentionally refraining from telling us how confirmation works in cases of indeterministic theories. At least part of his reason for doing so is his unhappiness with the existing probabilistic theories of confirmation (2006, p. 33).¹¹ He does not, however, indicate whether adequate probabilistic theories could ever be formulated.

To convey the particulars of Worrall's view requires a little stage-setting. Hereafter I follow his own conventions. T stands for a given general theory, theoretical framework or paradigm and T' stands for a specific theory, theoretical framework or paradigm developed out of T . For brevity, I will skip reference to frameworks and paradigms and refer only to theories. Suppose we have a general theory T with one or more free parameters, a specific theory T' developed out of T by fixing the free parameters and a datum d which may have been used to fix those free parameters. Suppose further that T does not entail d . We can then summarise Worrall's view in the following two principles, the first corresponding to support for specific theories and the second to support for general ones:

(S): A datum d provides confirmational support for T' if and only if (i) d is entailed by T' and *either* (ii) d is not used to fix free parameters in T (so as to yield T'), i.e. d is independent of any data used in the construction of T' , in which case the support is unconditional *or else* (iii) d is used to fix free parameters in T (so as to yield T') in which case the support is conditional upon accepting T .

(G): A datum d provides unconditional confirmational support for T if and only if (a) d is entailed by T and *either* (b) d is not used to fix free parameters in T (so as to yield T'), i.e. d is independent of any data used in the construction of T' , *or* (c) d falls naturally out of T .¹²

Worrall's view also appears to qualify as an objective theory of confirmation in that whether or not a datum provides (conditional or unconditional) support for T' or T is determined not by personal preferences but by facts about logic, facts about the fixing of free parameters and facts about a presumed natural relation between d and T .

Let us get a bit more grip on these principles, starting with (S). The first two conditions are fairly straightforward. Condition (i) expresses the idea that the logical entailment of data by a hypothesis is necessary for support. This idea makes sense in the context of Worrall's theory as he restricts his attention to non-probabilistic cases.¹³ Condition (ii) expresses the central idea behind use-novelty. Condition (iii) appears at first sight counterintuitive as it says something that one wouldn't expect from a use-novelty account. How could data *used* in the construction of a (specific) theory also support it? The trick is that the support provided is of a different kind than that referred to in (ii).

¹¹ Worrall seems to make two false assumptions here: (i) that indeterministic hypotheses cannot stand in any deductive relations to data and (ii) that deterministic hypotheses cannot stand in any inductive relations to data. All that he should be saying is that his theory only focuses on cases where a datum d (or not- d) is deductively entailed by a hypothesis.

¹² For an informal exposition of the conditions under which theories earn confirmation see Worrall (2002, pp. 203-204; 2005, pp. 817-818; 2006, pp. 50-51). For a principled exposition see his (2006, p. 56). The latter (but not the former) weirdly neglects the conditions under which general theories earn, according to him, so-called 'spill-over' support – those are listed as (b) and (c) in my principled exposition. Note also two further differences between my own and Worrall's principled exposition. First, his two principles concern conditional and unconditional support respectively, whereas mine carve up confirmation relations in terms of specific and general theories. Second, his two principles are formulated as material conditionals, mine as material bi-conditionals. Given that Worrall restricts his notions of confirmation to non-probabilistic cases, the other direction of the material implication also holds. That he approves my bi-conditional formulation has been established in e-mail correspondence I had with him.

¹³ If it is true that some (non-trivial) consequences of a hypothesis do not provide confirmational support to it then conditions (i) and (ii) are obviously insufficient. Also if support can come from non-consequences then condition (i) is not necessary. See Laudan and Leplin (1991) for arguments supporting the satisfaction of both antecedents. I argue against the satisfaction of the first antecedent in Section 9.

Given that a general framework, or research programme, is already accepted, then the data give – in the case of a genuine deduction – not just *some* support for the specific theory, but conclusive support (2002, pp. 201-203) [original emphasis] – see also (2006, p. 51).

The key to understanding condition (iii) is deductivism concerning evidential relations. There's no better way to transmit evidential warrant than by a deductively valid argument, for the truth of the premises guarantees the truth of the conclusion (2006, p. 43). Thus, provided we have grounds to believe in the truth of the premises, which in the cases discussed by Worrall include T and d , we have grounds to believe in the truth of the conclusion, i.e. T' . Since this support is conditional on the truth of the premises it is obviously conditional on the truth of T .

Despite calling it 'conclusive', Worrall ultimately downplays conditional support and claims that "[r]eal support for T' must be sought through *independent* evidence" (2002, p. 195) [original emphasis]. This is unsurprising given that, by the lights of his own theory, post hoc manoeuvres (which he construes as cases of parameter fixing) meet the requirements of conditional support. Worrall cites Velikovsky's theory as a case in point. In its general form the theory holds, among other things, that Venus was once a comet that broke off from Jupiter and eventually settled in its current orbit as the second planet from the sun. Velikovsky professed that this theory (plus auxiliaries) explains alleged cataclysmic phenomena of the past, e.g. the parting of the Red Sea, in terms of the proximity of the comet's orbit to that of the Earth's. Now if the general theory is correct, cataclysmic events must have taken place throughout the world during those close encounters. Velikovsky attempted to find records of such cataclysms but came up short. To compensate for this failure he postulated that events like these were so traumatic that many cultures did not record them. In its specific form Velikovsky's theory thus says, among other things, that cataclysmic events took place because of the comet's close encounters with our planet and even though all cultures witnessed these events, those that didn't record them were too traumatised. Worrall judges that this theory is supported by the cataclysmic record evidence *conditional* on our accepting the general version of Velikovsky's theory. Indeed, Worrall would add that unconditional confirmation is missing in this case since Velikovsky's specific theory never enjoyed any independent support.

Moving on to (G), we can skip discussion of condition (a) since it has the same motivation as condition (i) in principle (S) and instead focus on the other two conditions, i.e. (b) and (c). Each of these conditions represents a different type of unconditional spill-over support that a general theory may earn. Take condition (b) first. Worrall cites the specific version of Fresnel's wave theory of light as a successful example of the first type of spill-over support. The general wave theory posits that light consists of waves that are transmitted through an all-pervading mechanical medium, the ether. Waves from different sources have different wavelengths. The exact wavelength of a given source is not determined by the general theory itself but by performing experiments. For example, one can perform a two-slit experiment with a sodium arc as the light source and measure the distance between fringes in the observed diffraction pattern. Since the general theory posits a bijective mapping between fringe distance and wavelength, the wavelength of sodium can be determined. This allows the construction of a specific version of the wave theory, one that focuses on light from a sodium source. The specific theory can then be employed to successfully predict independent phenomena. In the case at hand, for example, it can predict fringe separations in the one-slit diffraction experiment with light coming from a sodium source (2006, p. 47).

Now take condition (c). Although Worrall sometimes talks of phenomena "fall[ing] naturally out of the core idea", of what we might presume is the core idea of a general theory, what he really means is that such phenomena fall naturally out of the general theory plus one or more auxiliaries (2002, p. 203). Indeed, most references to the notion of naturalness are in terms of natural auxiliaries. What are the identifying marks of natural auxiliaries? Alas, Worrall says very little about them. It is easier to jump straight to an example. Worrall (2005, p. 818; 2006, pp. 48-49) cites planetary stations and

retrogressions as falling naturally out of Copernicus' general heliocentric theory plus certain auxiliaries. 'Stations' refer to the apparent stops that planets come to in their journey across the night sky while 'retrogressions' refer to their apparent backward motion. The pertinent natural auxiliary is roughly the following: To an observer on Earth viewing the motion of a planet (Mars, Jupiter, Saturn and Uranus) against the largely stationary background of the stars, the planet *appears* to slow down and even stop (station) before turning backwards (retrogression) when the Earth in its smaller orbit around the Sun overtakes that planet. Although Worrall does not explain why an auxiliary like this is natural for Copernicus' general heliocentric theory, we can provide the following motivation: Because objects on Earth traversing elliptical paths of different radii around the same two foci yield the same station and retrogression phenomena as their cousins in space, no special manoeuvres are required in marrying Copernicus' theory with the said auxiliary.

Before we proceed to a critical evaluation of Worrall's view, we must dismiss a common misconception about it. Fitting data to theories is a widespread and respected practice in science. But fitting data is nothing other than using it to post hocly construct or modify theories or models. Sober and Hitchcock (2004) are keen to point out that "[a]ccommodation is not always bad", and in fact that "[f]it with existing data is a good thing" (p. 6) [original emphasis]. In their view the only time that accommodation goes astray is when the methods employed to accommodate the data fail to guard against over-fitting. Similarly, Lange (2001) argues that accommodation is typically bad when the resulting hypotheses are in fact coincidental truths or arbitrary conjunctions. Given what has already been said, it is tempting to assume that Worrall is squarely against the practice of data fitting or accommodation. That couldn't be further from the truth. What Worrall (2002, p. 198) objects to is the attribution of confirmational weight to data that has already been used to construct or modify a theory. That is, he endorses a variant of the so-called 'no double counting' rule, i.e. the rule that once a datum has been used to construct a theory, it cannot also be used to confirm that theory. In Worrall's variant (2006, p. 57) such a datum confirms a specific theory but in a conditional way, a kind of confirmation that, as we saw, he does not consider real.

6. The Complete Data Set Counterexample

In this section, I would like to briefly discuss an important counter-example to Worrall's theory and use-novelty more generally. Mayo (1996, p. 271) asks us to imagine a logic class which contains students who took the SAT examination. Suppose we want to find out the average SAT score of those students. Suppose further that we have access to all the scores. The best way to approach the matter is to add up all the individual scores and divide the resulting number by the number of students in the class. The result, suppose it is 1121, would then allow us to formulate the following true hypothesis: The average SAT score of this logic class is 1121. Surely, Mayo concludes, the data, i.e. the individual scores, support the hypothesis in question even though they have been used to construct it. In reply, Worrall complains that Mayo's objection is not a case of genuine confirmation. In more detail, he first complains that the assertion is not a genuine hypothesis because "the relationship between the individual scores and the average score is analytic". The implication is that genuine hypotheses do not stand in such a relationship to the data (2006, p. 59). He then complains that "a test of a theory surely must have a possible outcome that is inconsistent with the theory" (2006, p. 58). This is presumably absent in the SAT score case since nothing in the relevant construction process "could possibly refute the 'theory' that we end up with" (p. 58).

Both replies are unwarranted. Take the first. The relationship between any true hypothesis and the set containing *all and only* its consequences is bound to be analytic. After all, the two have exactly the same content. So a non-analytic relationship cannot be the mark of a genuine hypothesis. Now take the second. Does the SAT score data support the SAT score hypothesis? To see why this question must be answered affirmatively, consider the following twist to the example. Suppose that the same hypothesis were posited without recourse to the data, e.g. via some informed judgment about how students in logic classes typically perform in such tests. How would we go about supporting or refuting it? Well, obviously, by finding out all the individual scores and correctly

calculating their average. And since that calculation is in perfect agreement with the hypothesis we must concede that the data (fully) supports it. Moreover, if, as will be argued in the Section 9, support is an invariant relation that holds only between a data set and a hypothesis, then the SAT score data (fully) supports the SAT score hypothesis even in Mayo's original example, i.e. where the data were used to construct the hypothesis. The point generalises to any true hypothesis and its corresponding complete data set. There is thus nothing improper about claiming that a complete data set supports a true hypothesis and thus nothing improper about Mayo's counter-example to Worrall's theory.

7. Contingency and Conflicting Confirmational Assessments

Unlike Zahar (1973), Worrall does not urge us to peer into the notebooks of scientists in order to assess whether data has been used to construct a theory or auxiliary. He is wholeheartedly bent on eliminating psychologically contingent considerations from his account: "My account gives no role to any such psychological factor" (2005, p. 819). And it is not only psychological factors that he dismisses. As we saw in Section 4, he laughs off the idea that any "contingent historical issue" whatsoever should have an effect on confirmation (2002, p. 194). That Worrall sees his own confirmation theory as steering well clear of this idea is reinforced a few years later: "When properly understood, however, the 'heuristic' view I advocate does not have this historical character" (2006, pp. 55-56). But then how is confirmation to be decided according to him? The answer is on purely logical grounds. Referring to his theory, he says:

Although presented as a version of the 'heuristic approach', it is at root a *logical* theory of confirmation - the important logical relations being between (i) the evidence at issue e , (ii) the general theoretical framework involved T and (iii) the specific theory T' (2005, p. 819) [original emphasis].

A few pages later he reiterates this point: "it is the sort of logical connections between evidence, general and specific theories highlighted in my approach that really do the work" (2005 p. 823). And again in another publication: "The main conclusion of this paper is that there are two types of confirmation [conditional and unconditional] —both of them (three-place) 'logical'." (2006, p. 56).

In spite of Worrall's best efforts, contingent considerations are still permitted to play a confirmational role in his account. Suppose a data set O is entailed by a specific theory T' . Suppose further that $O = O_1 \cup O_2$ but also that $O_1 \cap O_2 = \emptyset$. Suppose moreover that O_1 and O_2 possess a number of qualities we admire in data sets, e.g. diversity, accuracy, informativeness, etc., and that they possess these qualities in equal measure, i.e. they are indistinguishable with respect to these qualities.¹⁴ Finally, suppose that T' can be constructed from either O_1 or O_2 by fixing exactly the same free parameter(s). If T' is constructed from O_1 and T' is utilised to predict O_2 , then T' gets unconditional support only from O_2 . If, however, T' is constructed from O_2 and T' is utilised to predict O_1 , then T' gets unconditional support only from O_1 . Worrall's account commits a double crime here. It issues conflicting confirmational judgments. And, contrary to its own pronouncements, its judgments depend on contingent considerations. Crucially, neither set can be dismissed as inferior and therefore as less deserving to provide unconditional support since, by supposition, O_1 and O_2 are qualitatively on par. And crucially, nothing in Worrall's allegedly pure logical account forbids the existence of disjoint data sets like O_1 and O_2 that are individually sufficient to fix exactly the same free parameter(s) in the construction of T' .

¹⁴ Referee 1 objected that we do not find data sets that are qualitatively indistinguishable in actual cases of science and, as such, the challenge my example poses amounts to general scepticism. Firstly, it must be pointed out that most confirmation theorists, and potentially also Worrall, aim to capture not only actual cases from the history of science but also hypothetical ones. Secondly, all it takes is one case of qualitative indistinguishability for Worrall's account to fail. Thirdly, actual science is littered with such cases. Just think of two sets of data gathered using the same instrument but each recording a distinct orbit of a planet around its parent star. The two sets are likely to be qualitatively indistinguishable.

Worrall (2006, pp. 51-56), inspired by Musgrave (1974, pp. 13-14), considers a family of objections to his view that are almost identical to the objection just advanced. The starting point for all these objections is two scientists *A* and *B* who each arrive at the same specific theory via different routes. Three such route-variations are considered. We only need entertain one here, the one that's most similar to the above objection. Suppose scientist *A* produces T' from some evidence e_1 and then goes on to predict e_2 from it, whereas scientist *B* does the reverse. This would mean that from scientist *A*'s perspective e_1 provides conditional support for T' and e_2 unconditional support for T' and T , but from scientist *B*'s perspective the opposite is true. Worrall's reply to this objection is that even though the two judgments are "strictly different" they are equivalent in what matters. In his own words:

each of *A* and *B* has shown that the general theory needs to fill in one parameter value on the basis of one piece of data, thus producing a specific theory that gains genuine empirical success from the other piece of data... So each scientist shows that there is, so to speak, one unit of genuine, unconditional, general-theory-involving data and hence delivers the judgment that that general theory is ahead in terms of empirical support of any theory that merely accommodates both pieces of data (p. 55).

This reply is tantamount to burying one's head in the sand for the simple reason that Worrall's account, as it is encoded in principles (S) and (G), still issues *conflicting* confirmational assessments regarding e_1 and e_2 .¹⁵ That's surely an undesirable feature for a confirmational theory to possess. Worrall seems to unwittingly subscribe to the undesirability of this feature, in the process shooting himself in the foot, when he asserts: "It is clearly a desideratum on any account of confirmation that it underwrite the judgement 'same evidence, same theory, same confirmation' and my account underwrites exactly this judgment" (p. 55). In the above two examples the *same* evidence does not provide the *same* theory with the *same* confirmation.

8. A Pandemic for Incidental Predictivists

Worrall's view is the only brand of predictivism that attempts, and as I just argued fails, to exclude contingent considerations from confirmational matters. All other brands of predictivism deliberately include such considerations in their calculations. Let us call these predictivist views 'incidental'. As we saw earlier, the inclusion of contingent considerations is not in and of itself objectionable, for such an objection would beg the question against incidental predictivists. Even so, I will now demonstrate that, much like Pandora's box, contingent considerations appear innocuous on the outside but conceal a terrible pestilence within. Indeed, if I am right, this is a pestilence that takes pandemic proportions with no incidental predictivism, at least not one found in the literature, left standing.

Incidental predictivism comes in various forms. What they all have in common is the inclusion of a condition whose satisfaction depends on specific contingent considerations, that is, considerations like whether or not a datum was known at the time a theory was formulated. Let us denote any such contingent condition with the letter 'X'. Beyond this condition, incidental predictivists, like non-incidental ones and confirmation theorists more generally, require the satisfaction of two further conditions. The first is a condition whose satisfaction depends on specific inferential considerations, i.e. considerations like whether or not a datum is logically entailed by a hypothesis. Let us denote any such inferential condition with the letter 'Y'. The second is a condition whose satisfaction depends on specific semantic considerations, i.e. considerations like whether or not a datum is true. Let us denote any such semantic condition with the letter 'Z'. A version of incidental predictivism then holds that a hypothesis is supported (or more supported) by a set of data if (and perhaps also only if) its specific construals of conditions *X*, *Y* and *Z* are satisfied.

¹⁵ Worrall may attempt to modify these principles accordingly but until he does we cannot begin to evaluate them. Note that whatever form such modification takes Worrall will have to ensure that his notion of 'use' does not allow contingent considerations to come into play, otherwise the same problems will surface.

Now consider the following variant of the two-set counterexample. First, suppose that disjoint data sets O_1 and O_2 each satisfy condition Y_s by standing in certain inferential relations to a hypothesis H . Second, suppose that each data set possesses a number of specific semantic qualities that amount to the satisfaction of condition Z_s . To simplify things, take the inferential relation to be entailment and the semantic quality to be truth since both are always included as limit cases in their respective conditions. Third, suppose that condition X_s can be satisfied by both sets. Fourth, suppose that we have two scientists, F and G , each of whom belongs to a distinct isolated scientific community but both of whom support the specific form of incidental predictivism that emerges out of conditions X_s , Y_s and Z_s . Finally, suppose that for scientist F condition X_s is satisfied by O_1 but not by O_2 but for scientist G condition X_s is satisfied by O_2 but not by O_1 . F endorses the claim:

(C1): O_1 supports hypothesis H but not O_2 (or O_1 supports hypothesis H more than O_2).

Whereas G endorses the claim:

(C2): O_2 supports hypothesis H but not O_1 (or O_2 supports hypothesis H more than O_1).

The two claims are clearly inconsistent. The inconsistency arises as a result of the fact that condition X_s can be satisfied by either set. And since condition X (but also Y and Z) can take any number of forms this makes the issuing of conflicting confirmational judgments a highly prevalent feature of incidental predictivism.

Why highly prevalent but not universal? Because there is one kind of form condition X can take, as far as I can see one of only two (see Footnote 16 for the other), which saves incidental predictivism from the above objection. The kind of form I have in mind requires condition X to take into account historical details from all scientists in the universe who find themselves in the above circumstances. Alas for incidental predictivism, this move exchanges one damning objection for another. Confirmational assessments now become impossible in practice since we do not have, and cannot realistically be expected to have, epistemic access to the history of all such scientists in the universe. Take Zahar's (1973) version of predictivism as an example. According to his conception of condition X , if a datum is to confirm a hypothesis it must not have been explanatorily targeted by the scientist who constructs that hypothesis. As it stands, his view falls prey to the objection in the previous paragraph – see Musgrave (1974, pp. 13-14) for a similar objection. To avoid that objection, Zahar may modify his version of condition X as follows: if a datum is to confirm a hypothesis it must not have been explanatorily targeted by *any* scientist in the universe who constructs that hypothesis. It should be obvious that checking whether this condition is ever met is in practice impossible.

All existing versions of incidental predictivism are subject to the first (or, if modified to avoid it, the second) objection regardless of the form their condition X takes. Aside from Zahar's version of condition, this includes the conditions of 'being temporally novel' (both weak and strong versions), Musgrave's (1974) (and, according to Musgrave, Lakatos' 1968) 'not having been predicted by a theory's existing rivals', Leplin's (1997) hybrid condition 'not having been developed on the basis of the observational results and not having been predicted or explained by a theory's existing rivals' and potentially Barnes' (2008) 'being endorsement-novel relative to a scientist'.¹⁶ In light of the problems raised in this section, it is reasonable to conclude that incidental predictivism is not likely to serve as a basis for an adequate theory of confirmation.

¹⁶ The other way to escape the first objection is to be a subjective confirmation theorist. Someone like Barnes may argue that different people having conflicting judgments is nothing to worry a subjectivist about. Although this reply does not fall prey to the second objection it does raise the usual gamut of concerns associated with subjective confirmation theories, e.g. how to establish rational inter-subjective agreement.

9. Cornerstones for an Objective Theory of Confirmation

It must seem like ages since we last spoke about objectivity and subjectivity in confirmation. Yet a highly pertinent discussion has already been taking place right under our noses. The failures of predictivism exposed in the last three sections are highly instructive in our search to lay some foundations for an objective theory of confirmation. In this section and the next, I would like to place some cornerstones in these foundations in the form of four desiderata that an adequate objective theory of confirmation would need to satisfy. Though necessary, these desiderata are not intended to be sufficient.

We have already learned that incidental versions of predictivism are problematic precisely because of their reliance on contingent considerations. Thus, although contingent considerations are objective in character and can therefore be permitted to play a role in an objective theory of confirmation, they are not the kind of considerations we want such a theory to possess. Ridding ourselves from contingent considerations leaves only inferential and semantic ones to fall back on. Our first desideratum then is to demand that an objective theory of confirmation articulate appropriate inferential and semantic (but not contingent) conditions such that the confirmational judgments it issues remain invariant under anything other than the evidence and the hypothesis (plus any auxiliaries) in question. This demand was earlier encoded in Worrall's dictum – which, it is worth reminding, he unintentionally violates – 'same evidence, same hypothesis (plus any of the same auxiliaries), same confirmation'.

A second desideratum is suggested by one of the failures of Worrall's version of predictivism. You may recall that Worrall attempts to save his view from Mayo's counterexample by arguing that it does not amount to a case of genuine confirmation. In doing so, Worrall was in effect trying to dismiss the counterexample by conveniently reducing the set of cases that we may legitimately call cases of genuine confirmation. His failure to provide warrant for this reduction should be a cautionary tale for all those thinking of similar reduction moves. A theory of confirmation needs to be able to tell us whether or not *any* given piece of evidence supports, opposes or is neutral with respect to *any* given hypothesis. This is a form of comprehensiveness. Our second desideratum then is that an objective theory of confirmation needs to be comprehensive.

In Mayo's counterexample we can also find hints of a third desideratum. The basic set-up in that counterexample is a true hypothesis constructed out of a complete set of data. But what constitutes a complete set of data? Such a set *can* be thought of as involving all the deductive consequences of that hypothesis. Suppose E , the set used to construct T' , contains all the consequences of T' . Under this set-up, Worrall's account dictates that E cannot provide any real, i.e. unconditional, support to T' . Otherwise put, no consequence of T' lends support to it. This violates a highly intuitive principle that is related, though it is weaker, to Hempel's so-called 'entailment condition'. I call my principle the 'Consequence Principle' or (CP) for short.

(CP): All (non-trivial) deductive consequences of a hypothesis that are true confirm it and all false ones disconfirm it.¹⁷

Putting aside its intuitive appeal, I would like to put forth more robust reasons why we must pledge our allegiance to this principle. Nobody would deny that any one (non-trivial) consequence of a hypothesis, if it turns out false, is enough to refute that hypothesis as it is currently formulated. But if one consequence can refute a hypothesis, then surely hypotheses are never fully confirmed, i.e. are not true without exception, prior to checking that consequence. Hence that consequence possesses positive or negative confirmational weight depending on whether it is true or false respectively. Since

¹⁷ Trivial consequences include logical truths and irrelevant disjuncts. A good way to cash out the notion of 'non-trivial' consequences' is via the notion of relevant (as opposed to irrelevant) consequences. See the next section for an exposition of this notion.

this holds for any consequence, all consequences of a hypothesis possess positive or negative confirmational weight. Equivalently, all true consequences of a hypothesis confirm it and all false consequences of a hypothesis disconfirm it. That is, (CP) holds.

Our third desideratum is adherence to (CP). An attractive by-product of this adherence is that (CP) conflicts with strong predictivism. The outlandishness of strong predictivism becomes evident when one considers that the view violates not only (CP) but also a much weaker, and indeed very weak, principle that I call the 'Partial Consequence Principle' or (PCP) for short.

(PCP): At least some (non-trivial) deductive consequences of a hypothesis that are true confirm it.

Suppose that evidential set E contains all the (non-trivial) consequences of a hypothesis H . Suppose, furthermore, that E does not meet the distinctive condition(s) that make(s) a version of strong predictivism the version it is. Then H cannot get any support from its consequences, thereby leading to a violation of (PCP). As an illustration, consider the strong temporal novelty view. Its distinctive condition holds that the only source of support for a hypothesis is from evidence that becomes known after the hypothesis (plus any auxiliaries) was formulated or modified. Since by supposition E violates this condition, its contents cannot be used to support H . Moreover, since by supposition E contains all the (non-trivial) consequences of H this entails that no such consequence of H lends support to it. Hence, the strong temporal novelty view runs afoul of (PCP). Structurally identical arguments can be launched against other versions of strong predictivism.¹⁸ Unless the violation of a strong predictivism's distinctive condition(s) is logically impossible, and, frankly, I do not see how this could be the case, there is no escape for the strong predictivist.¹⁹

10. Post hoc and Other Monsters

No discussion of desiderata for an objective theory of confirmation would be complete without saying something about the much-derided post hoc monsters. Recollect that to meet the completion challenge requires, among other things, dealing with the niggling issue of whether or not, and if so to what extent, post hoc monsters can be confirmed. One approach to attain that goal has been to either demonise or penalise post hoc-ness itself. This approach has been taken up by the predictivist movement. Strong versions of predictivism demonise post hoc-ness by claiming that no post hocly constructed or modified hypothesis earns support from accommodated evidence. Weak versions of predictivism penalise it by claiming that no such hypothesis earns as much support from accommodated evidence as one that is not so constructed. The various counterexamples to predictivist views we saw earlier stand as a vivid reminder that not every post hocly constructed hypothesis is a monster. Indeed, some of them may even be true hypotheses constructed from true evidence. In short, the approach of demonising or penalising post hoc-ness is not subtle enough. As Allan Franklin jokingly put it in a title of a talk he delivered at the London School of Economics: "Ad hoc is not a four letter word" (see Worrall 2006, p. 42).

Another approach is to target not post hoc-ness itself but, more narrowly, post hoc monsters. Once again two options are available. The first seeks to demonise post hoc monsters. Goodman (1983) may be an advocate of this option. Although he doesn't directly talk about post hoc monsters he holds that hypotheses like the ones we considered earlier are not confirmed by their consequences. The second seeks to penalise post hoc monsters. I will opt for a variant of the second option. The

¹⁸ Both referees raised the concern that examples like this are atypical in science since most cases involve abductive inferences. One of them, referee 2, also offered the right sort of reaction to this concern: Any account of confirmation that purports to be universal should be able to cope with all cases, including those where the hypothesis can be deductively inferred from the evidence.

¹⁹ The failure of various forms of predictivism means that realists need a radical rethink of the notion of empirical success in the so-called 'no miracles argument', since at present that notion relies heavily on the ability of theories to make novel predictions.

reason I reject the first is that in my view even monstrous hypotheses deserve to be confirmed by accommodated evidence. Here's why. Take hypothesis $H_{AB}: A \wedge B$. To determine the truth of H_{AB} one needs to determine both the truth of A and of B . Thus some support for H_{AB} arises from the truth of proposition A . This holds even if A is accommodated evidence and propositions A and B are post hoc stitched together in a monstrous way. And since we haven't specified the content of A or B , the point is obviously general. Regardless of the route through which certain propositions get to be included in a hypothesis, that hypothesis will be confirmed by the truth of those propositions. Thus, no matter how counterintuitive it may sound, and subject to an important qualification I am about to make, creationism (plus auxiliaries), Ptolemaic astronomy (plus auxiliaries), Velikovsky's theory (plus auxiliaries), etc., earn confirmation from the true propositions they were designed to entail.

The above argument dictates a disavowal of the completion challenge. This should not be news. We have been on a direct collision course with the challenge for some time. Among other things, the challenge conflicts with the first desideratum as that desideratum demands that an objective theory of confirmation employs only inferential and semantic considerations. Where does this leave us? I do not mean to imply that there are no genuine concerns behind the completion challenge or behind the confirmation of post hoc monsters. Rather, I would like to propose a different approach to both of these issues.

Let us start with post hoc monsters. Monstrous hypotheses, like the great fictional monster in Mary Shelley's *Frankenstein*, are assembled out of a motley assortment of parts. Each individual part may be fine on its own but when all the parts are put together they give rise to a monster. Let us call a hypothesis 'monstrous' if and only if some of its content parts are disjointed. The notion of disjointedness I have in mind is as follows: Any two content parts expressed as propositions A, B are disjointed if and only if $P(\alpha/\beta) = P(\alpha)$ for all propositions α, β where α is a relevant deductive consequence of A and β is a relevant deductive consequence of B . The first thing to note about the concept of disjointedness is that it is articulated in terms of the concept of *probabilistic independence*. We say that two propositions α, β are probabilistically independent just in case $P(\alpha/\beta) = P(\alpha)$. The concept of probabilistic independence is apt here because it allows us to express the idea that two propositions are confirmationally unrelated. After all, the probability of the one is not affected if we assume the truth (or falsity) of the other.²⁰ The second thing to note about the concept of disjointedness is that to establish the confirmational unrelated-ness between two propositions A, B it is not enough to merely focus on the propositions themselves. We must also take into account their *deductive consequences*. The reason for this is that two propositions may be probabilistically independent even though some of their deductive consequences are probabilistically dependent.²¹ To rule out such cases we must demand that probabilistic independence holds all the way down, that is, between all – save for an exception to be discussed below – the deductive consequences of two propositions. This demand is an apt way to express the idea that no part of the content of the one proposition confirmationally affects any part of the content of the other proposition. The third and final thing to note is that unless we restrict our evaluation to *relevant* deductive consequences of propositions the concept of disjointedness would be unsatisfiable, i.e. no two propositions would ever qualify as being disjointed. The idea of a relevant deductive consequence is fully developed in Schurz (1991): "the conclusion of a given deduction is irrelevant iff the conclusion contains a component [i.e. a formula] which may be replaced by any other formula, salva validitate of the

²⁰ The probabilities in question are objective, i.e. they are determined by the structure of the world. The quality of our assessments about disjointedness is thus dependent on the quality of our information about the structure of the world.

²¹ Here's an example. Take $A: A_1 \wedge A_2$ and $B: A_1 \wedge B_1$. Suppose that A_1, A_2 are independent but also that A_1, B_1 are independent. We may assign values, e.g. $P(A_1) = 0.5, P(A_2) = 0.5, P(B_1) = 0.5$ and $P(B/A) = 0.25$, such that $P(A/B) = P(A)$ even though there is a proposition α that follows from A and a proposition β that follows from B , in both cases this proposition being A_1 , such that $P(\alpha/\beta) \neq P(\alpha)$. This is because $P(A_1/A_1) = 1$ and hence $P(A_1/A_1) \neq P(A_1)$.

deduction” (pp. 400-1). Here’s why we need it. Whatever the content of propositions A , B we can always validly derive consequences that are common to both. For example, using the classical rule of disjunction introduction we can derive the proposition $A \vee B$. The existence of such trivial common consequences guarantees that there is a pair of propositions α , β for which $P(\alpha/\beta) \neq P(\alpha)$ provided $0 < P(\alpha) < 1$. Obviously such consequences are irrelevant to the evaluation of the non-disjointedness between A and B . The restriction to relevant consequences forbids this kind of situation by ruling out irrelevant formulas, i.e. formulas such as $A \vee B$.²²

Under this conception, monstrous hypotheses need not be post hocly constructed but also non-monstrous hypotheses may be post hocly constructed. Notice also that hypotheses may possess both disjointed and non-disjointed content parts. To be exact, since disjointedness and non-disjointedness are relations that hold between various content parts of hypotheses, the claim is that hypotheses may possess content parts, some of which are disjointed and others non-disjointed from other content parts. It is thus more informative to speak about monstrous and non-monstrous relations between content parts of a given hypothesis rather than monstrous and non-monstrous hypotheses.

My suggestion on how to approach (post hoc and other) monsters involves the idea that disjointedness forms a barrier against the spread of confirmation.²³ Take hypothesis H_{AB} again. Suppose that A and B are disjointed. If that is indeed the case, then finding out about the truth of one (or any of its relevant deductive consequences) will leave the truth of the other (or any of its relevant deductive consequences) unaffected. This means that there is a confirmation barrier between A and B . Proposition A (or not- A) confirms (or disconfirms) only that part of the content of H_{AB} that corresponds to itself, namely A . Ditto for proposition B . Thus even though (post hoc and other) monsters get confirmed under my view, the confirmation they receive for a content part that is disjointed from other content parts doesn’t spread to those other parts. This is unlike what happens in cases of non-monstrous relations between content parts of a given hypothesis where the confirmation of a content part that is not disjointed from other content parts spreads to those parts.²⁴

We are now ready to formulate our fourth desideratum: Any objective theory of confirmation must ensure that support earned for disjointed parts does not spread beyond those parts. Taking stock from all that has hitherto been said, we can propose a revised version of the completion challenge that applies only to objective theories of confirmation: Objective confirmation theorists must tell us what is needed in addition to the four desiderata listed here to obtain a complete account of confirmation.

Acknowledgements: I have benefitted from discussions with, and would therefore like to thank, José Díez, Ludwig Fahrback, Clark Glymour, Carl Hoefer, Stathis Psillos, Sam Schindler, Paul Thorn and John Worrall. I am also thankful to two anonymous referees. Referee 2, in particular, provided exceptionally incisive comments. I gratefully acknowledge the German Research Foundation (Deutsche Forschungsgemeinschaft) for funding my research under project B4 of Collaborative Research Centre 991: The Structure of Representations in Language, Cognition, and Science. Part of this paper has been written while working on the project ‘Aspects and Prospects of Realism in the Philosophy of Science and Mathematics’ (APRePoSMa) during a visiting fellowship at the University of Athens. The project and my visits are co-financed by the European Union (European Social Fund - ESF) and Greek national funds through the Operational Program ‘Education and Lifelong Learning’ of the National Strategic Reference Framework (NSRF) - Research Funding Program: THALIS -UOA.

²² For more examples of irrelevant consequences see Schurz (1991).

²³ This idea is similar to Goodman’s own that in cases of ad hoc-ness “establishment of one component endows the whole statement with no credibility that is transmitted to other component statements” (1983, pp. 68-69). Alas, Goodman and I don’t see eye-to-eye as he effectively rejects (CP).

²⁴ Note that it may spread to some sub-parts but not others.

References:

- Barnes, E. C. (2008). *The paradox of predictivism*. Cambridge: Cambridge University Press.
- Duhem, P. ([1914] 1991). *The aim and structure of physical theory*. Princeton (NJ): Princeton University Press.
- Earman, J. and C. Glymour (1978). Einstein and Hilbert: Two months in the history of general relativity', *Archive for History of Exact Sciences*, vol. 19: 291 – 308.
- Glymour, C. (2008). The paradox of predictivism (book review). *Notre Dame Philosophical Reviews*, <http://ndpr.nd.edu/news/23561-the-paradox-of-predictivism/>
- Goodman, N. (1983). *Fact, fiction and forecast* (4th edition). Cambridge, MA: Harvard University Press.
- Harker, D. (2008). On the predilections for predictions. *British Journal for the Philosophy of Science*, vol. 59: 429-453.
- Hitchcock, C. and E. Sober (2004). Prediction versus accommodation and the risk of overfitting. *British Journal for the Philosophy of Science*, vol. 55: 1-34.
- Lakatos, I. (1968). Changes in the problem of inductive logic. In I. Lakatos (Ed.), *The problem of inductive logic* (pp. 315-417). Michigan: North Holland Pub. Co.
- Lange, M. (2001). The apparent superiority of prediction to accommodation as a side effect: a reply to Maher. *British Journal for the Philosophy of Science*, vol. 52: 575-588.
- Laudan, L. and J. Leplin (1991). Empirical equivalence and underdetermination. *Journal of Philosophy*, vol. 88: 449–72.
- Leplin, J. (1997). *A novel defense of scientific realism*. Oxford: Oxford University Press.
- Maher, P. (1988). Prediction, accommodation, and the logic of discovery. In A. Fine and J. Leplin (Eds.), *PSA 1988*, vol. 1 (pp. 273-285), East Lansing, MI: Philosophy of Science Association.
- Mayo, D. (1996). *Error and growth of experimental knowledge*. Chicago: University of Chicago Press.
- Musgrave, A. (1974). Logical versus historical theories of confirmation. *British Journal for the Philosophy of Science*, vol. 25: 1-23.
- Votsis, I. (2011). Data meet theory: Up close and inferentially personal, *Synthese*, vol. 182: 89–100.
- Whewell, W. (1847). *Philosophy of the inductive sciences, founded upon their history*. London: John W. Parker.
- Williamson, J. (2010). *In defence of objective Bayesianism*, Oxford: Oxford University Press.
- Worrall, J. (1985). Scientific discovery and theory confirmation, In J. C. Pitt (Ed.), *Change and progress in modern science* (pp. 301–22). Dordrecht: D. Reidel.
- Worrall, J. (2002). New evidence for old. In J. Wolenski, & K. Kijania-Placek (Eds.), *In the scope of logic, methodology and philosophy of science* (pp. 191–212). Dordrecht: Kluwer.
- Worrall, J. (2005). Prediction and the 'periodic law': a rejoinder to Barnes, *Studies in History and Philosophy of Science*, vol. 36: 817–826.
- Worrall, J. (2006). Theory-confirmation and history. In C. Cheyne, & J. Worrall (Eds.), *Rationality and Reality: Conversations with Alan Musgrave* (pp. 31–62), Dordrecht: Kluwer.
- Zahar, E. (1973). Why did Einstein's programme supersede Lorentz's? (Part I and II). *British Journal for the Philosophy of Science*, vol. 24: 95-123 and 223-62.