**State of the field: Why novel prediction matters**

Heather Douglas and P.D. Magnus [1]

**ABSTRACT:**
There is considerable disagreement about the epistemic value of novel predictive success, i.e. when a scientist predicts an unexpected phenomenon, experiments are conducted, and the prediction proves to be accurate. We survey the field on this question, noting both fully articulated views such as weak and strong predictivism, and more nascent views, such as pluralist reasons for the instrumental value of prediction. By examining the various reasons offered for the value of prediction across a range of inferential contexts (including inferences from data to phenomena, from phenomena to theory, and from theory to framework), we can see that neither weak nor strong predictivism captures all of the reasons for valuing prediction available. A third path is presented, Pluralist Instrumental Predictivism; PIP for short.

Imagine three scientists. The first uses her theory to make a novel prediction, tests it, and the prediction is successful, garnering support for her theory. The second collects some evidence, and then alters her theory to accommodate the evidence. The third secretly gathers some evidence in her lab, accommodates her theory to the evidence, then publicly makes a prediction using her new theory that the evidence she already gathered will obtain (which it does again). The third scientist is obviously deceiving people about something; more than that, the deception seems to undercut the legitimacy of her results. As Gardner writes, it seems that lying about whether one had actually predicted or only accommodated evidence "would be a form of scientific dishonesty akin to fabricating experimental results" (Gardner, 1982, p. 11). But why should the third scientist's misbehaviour seem relevant to the evidential status of her results?  A natural answer is that she is

---

[1] This paper was entirely collaborative, and authors are listed in alphabetical order.

pretending to more epistemic warrant than she has, but this only makes sense if mere accommodation does not provide as much support for the theory as prediction. Thus, our reaction to the third scientist depends in part on whether we think that the first scientist has better support for her theory than the second.

The intuition that the first scientist does in fact have better support for her theory than the second or third is the intuition of *predictivism*, the idea that successful predictions generally provide better epistemic support than merely accommodating the evidence. Predictivists will say that the first scientist has the best support for her theory, and that the third scientist has done something horribly dishonest, not just ethically but epistemically. The opposing view, *accommodationism*, holds that there is no epistemic difference between the three scientists and the third scientist lied only about incidental things, not about anything epistemically significant. The question is not just which view is right — predictivism or accommodationism — but why and in what contexts.

Predictivism in any form requires consideration not just of the available body of relevant evidence but also how and when that evidence was gathered. Thus, one can think of predictivism as a form of scientific assessment that requires some historical understanding of how a scientific claim was developed and tested. This enters into the long-standing debate between (on the one hand) those who view the epistemic support for theories as resting solely with the logical relationship between evidence and theory and (on the other hand) those who view epistemic support as resting at least in part with how the theory was developed and tested. As Musgrave (1974) noted, Leibniz, Whewell, and Duhem all subscribed to the idea that predicting data was *ceteris paribus* better than accommodating it, while Mill and Keynes disagreed (Musgrave, 1974, pp. 1–2; see also Maher, 1988, p. 1).

Concern over novel predictions has been embedded in particular debates, but has also emerged as a philosophical topic in its own right. In the 1970s, the debate centered on the Lakatosian idea that only novel predictions made a program "progressive" (Lakatos, 1978). With such a strong requirement for novel predictions, much discussion concerned what counted as "novel." Philosophers from Zahar (1973) to Gardner (1982) argued against the requirement of *temporal novelty* in which prediction must be made before evidence gathered. In order to maintain 'prediction' as paramount while still accommodating philosophical intuitions and the history of science, the intuitive concept of novelty as temporal was replaced with heuristic or use-novelty. Evidence counts as a use-novel prediction of a theory if the scientist did not use

2

knowledge of the evidence in constructing the theory. Such expansion in the meaning of novelty came at a price. It is relatively easy to determine whether a prediction was temporally novel; deceptions of the kind in the our opening paragraph are rare. It is more difficult to assess use-novelty; see sec. 3.2, below.

By the 1980s, the debate shifted away from Lakatosian accounts of theory change and assessments of research programmes to whether formal theories of confirmation (especially Bayesian) could accommodate predictivist intuitions and to what role novel predictions could play in the debates about realism. The discussion on Bayesian methods and predictivist intuitions flourished to such an extent that by the mid-1990s, all possible positions on the role of predictivism within Bayesian confirmation theory had been defended.[2]

Predictivism, as both its own term and within the frame of the current debate, emerged in the late 1980s (Maher, 1988). Rather than focusing on debates about theory change, confirmation theories, or realism, the question of whether and to what extent successful novel prediction conferred stronger epistemic weight became its own topic of discussion. Since the late 1980s, the debate has centered largely on the following terrain: First, predictivists largely agree that accommodated evidence does provide some epistemic support.[3] So novel prediction is not the only way to gain epistemic support, as it was for the Lakatosians. What predictivists argue is that prediction provides more or better epistemic support than accommodation. Second, novel prediction is now taken roughly to mean heuristic or use-novelty; the question is whether scientists considered the already known evidence when constructing the theory: If they did, they performed an accommodation. If not, their theory successfully 'predicted' the evidence (even though the evidence was already available). Because the view that only novel predictions provide epistemic support is no

[2] Brush writes that "philosophers have defended all four possible positions: Bayesian analysis is (i) valid because it favors novel prediction, (ii) valid because it does not favor novel predictions, (iii) invalid because it favors novel predictions, and (iv) invalid because it does not favor novel predictions" (Brush, 1994, p. 134).

[3] The question of what kind of support and support of what can vary. Thus, Worrall argues that accommodated evidence only supports a particular version of a theory if scientists have already accepted the more general version of the theory and that accommodated evidence provides no support for the general version of the theory. (Scerri and Worrall, 2001, p. 425–6)

longer on the table, predictivists do not need to worry as much about the problems of assessing when a case is prediction or accommodation. If it is hard to tell whether a scientist made a use-novel prediction or accommodation, we can simply presume it was an accommodation and then assess the strength of epistemic support *qua* accommodation — that will still count for something.

Even with such a mild form of predictivism (in which accommodation counts too but just not as much, and where heuristic or use novelty is sufficient for 'novel prediction') it has still been a challenge to articulate why prediction is preferable to accommodation. Sometimes the argument is couched in naturalist terms: do scientists grant prediction special status or not? This literature is interesting, as it shows at least in part how difficult it can be to assess a particular historical example and whether it supports predictivist or accommodationist tendencies. Key historical examples have been evaluated and reevaluated, making it vexed as to whether or not actors treated novel prediction as special. Regarding diffraction, Fresnel's wave theory of light, and the bright spot, see *inter alia* Giere (1983) contra Worrall (1989). Regarding Einstein, relativity, and the perihelion of Mercury, see Zahar (1973) contra Earman and Glymour (1978). One of the most detailed debates on predictivism in the history of science has concerned Mendeleev and the predictions he made using the periodic law; see Brush (1996; 2007), Scerri and Worrall (2001), Worrall (2005), Scerri (2005), McIntyre (2001), Barnes (2005), and Schindler (2008). Despite the centrality of predictivism in the case of Mendeleev for Maher (1988) and Lipton (2004), it remains unclear whether and to what extent scientists were persuaded by predictions of new elements or revised atomic weights, rather than by accommodations of already well known chemical phenomena. Given the complexity of history, such uncertainty is not surprising. As Harker notes, "examples and counterexamples fuel, but don't settle, the debates" (Harker, 2008, p. 444). In many of the key historical examples, both stunning novel prediction and elegant accommodation of previously known data occur together. This makes trying to determine what actually impressed scientists at the time historically challenging.

This complexity points to a third possibility on the value (or not) of novel prediction. Many contemporary authors maintain that prediction is better than accommodation, but that this is not an irreducible virtue. Rather, prediction is good just because it reliably correlates with some other epistemic virtue. Such philosophers are willing to abandon prediction in cases where the other good is clearly in view. For

example, Harker (2008) argues that prediction just serves as a defeasible sign of explanatory power. Where it is possible to assess explanatory virtue directly, on his view, prediction is superfluous. This move distinguishes between strong predictivism (the view that there is some intrinsic value to prediction) and weak predictivism (the view that prediction is a surrogate for something else).[4]

Finally, there is a view nascent in the literature that is orthogonal to the strong/weak opposition. Suppose that prediction is not an intrinsic epistemic virtue, but rather that it serves as an instrumental indicator. Yet it can be a proxy for disparate sources of epistemic assurance. So it does not reduce to one particular virtue (as in weak predictivism), but neither does it have some kind of intrinsic epistemic virtue (as in strong predictivism). We dub this *Pluralist Instrumental Predictivism* (in sec. 5), and it is a position that can be articulated by coming to grips with the present state of the debate. As we proceed in this overview of the literature, we will focus on the recent philosophical debates on prediction, organizing as we go the reasons to think that novel prediction provides some instrumental assurance in a range of cases. We will then assess what this collection of reasons means for weak predictivism in its various guises.

To look ahead, our survey is structured along the following lines. We will first structure the levels of scientific inference for which novel prediction can be relevant. We consider inferences from data to phenomena, inferences from phenomena to theory, and inferences from theory to framework. We distinguish these levels in the next section, and then proceed to examine reasons for why novel prediction might be useful at each level. We conclude our survey by considering what we have gleaned from the literature for the views of Worrall and weak predictivists of various stripes.

## 1  Data, Phenomena, Theories, and Framework

Before surveying the range of reasons one might value novel prediction in scientific practice, we distinguish four levels across which scientists make inferences. Making these distinctions will serve as the organizing basis for explicating the value of novel prediction.

In this paper, there are four levels of concern: data, phenomena, theory, and framework. On the side of evidence, we draw from Bogen and Woodward (1988) to distinguish *data* from *phenomena*. On the side of theory, it is important to distinguish particular hypotheses

---

[4] See Harker (2008), p. 430.

subject to testing from broader, less-testable systems of background theories. We will call the former *theory* and the latter *framework*. These levels range from the rawest observation to the most sophisticated general commitments. To describe each in more detail:

The *data* are raw observations, such as the individual numbers or points on a scatter plot. They must be "relatively easy to identify, classify, measure, aggregate, and analyze in ways that are reliable and reproducible" (Bogen and Woodward, 1988, p. 320). Particular observations are idiosyncratic; they are disturbed randomly by noise and error.

The *phenomena* are the patterns in the world that are indicated by data, such as the true curve which underlies a scatter plot. The transformation of data into phenomena is a matter of statistics, data analysis, or data reduction. The phenomena "have stable, repeatable characteristics which will be detectable by means of a variety of different procedures, which may yield quite different kinds of data" (Bogen and Woodward, 1988, p. 317). Repeating an experiment and reproducing the result will not generate the same data, but it should provide new data that instantiate the same phenomena.

The *theory* typically predicts and explains phenomena. It is general, in that it will apply across more than just one kind of situation and often ties apparently different phenomena together. One can think of theories as sets of models, laws, or postulated entities and causal relations that provide explanations and predictions of the phenomena of concern.

The *framework* is the partly implicit background against which particular theories are formed. It includes auxiliary hypotheses and related commitments which are necessary in order for the theory to yield determinate predictions, such as commitments about when conditions or entities are the *same*. (As a tacit but unobjectionable example: It does not matter for a physics experiment which day of the week you do it on.) It also includes the motivation for particular theories along with general assumptions about the domain of phenomena and about how various pieces of experimental equipment do or do not function.

All four levels will be prominent in at least some cases. For example, consider rival explanations of growth perturbations and small size among people raised in adverse conditions. Schell and Magnus (2007) describe two rival theoretical approaches. The adaptationist approach tries to explain scrawniness as an adaptation to harsh conditions. The medical approach explains it instead as a disease or disfunction, harm done to people by their adverse conditions. These are

each general frameworks (in our sense) which underwrite specific theories. For example, the adaptationist approach might suggest the hypothesis that smaller people would be better off than bigger ones under conditions of scarcity because they need less food. This is a particular theory (in our sense) and it can be tested. Schell and Magnus argue that this distinction is important to understanding the rivalry between the two approaches.[5] The test of a particular hypothesis involves considering a large collection of specific measurements. The raw numbers are taken to indicate a general relation, perhaps between body mass, food intake, and physical performance. The raw measurements are the data, and the relation is the phenomenon. The relation is in turn used as a proxy for fitness, so the phenomenon can serve as evidence for or against the theory. If the theory succeeds, then this provides support for the framework; if it fails, then it is a demerit for the framework.

Nevertheless, it may not be possible to distinguish all four levels in every case. For example, there might be no distinction between phenomena and theory for some phenomenological or statistical hypotheses; cf. Glymour (2000). Regardless, there are typically more than two levels. In what follows, we will survey various reasons in the literature that novel prediction underwrites — with more epistemic assurance than accommodation — inferences among the adjacent levels of the hierarchy. Different, albeit conceptually related, reasons for the increased epistemic assurance arise in these different contexts.

## 2 Values for novel prediction

In this section, we glean from the literature reasons that novel prediction provides epistemic value over and above accommodation at each boundary between levels — moving from data to phenomena, from phenomena to theory, and from theory to framework.

Often, this value derives from the incomplete nature of our knowledge. Perhaps, as many have argued, the value of novel prediction would diminish were we to have complete epistemic clarity regarding the available evidence and the claims made about it. For example, White considers the limits of predictivism by positing problems that arise for agents that are "thoroughly familiar with the content of [the theory] and all the independent evidence supporting it"

---

[5] Following Lakatos, Schell and Magnus call the rival frameworks *research programmes*.

(White, 2003, pp. 657–8).[6] Even then, it would not evaporate entirely. Minimally, the process of making and testing a prediction produces additional relevant evidence. We concede, however, that the problem posed from a standpoint of perfect information makes predictivism seem thin. The problem, it seems to us, is with the standpoint and not with predictivism. For those interested in the value of novel prediction for actual scientific practice and inference, it matters that scientists are never in a situation of perfect knowledge and transparency. Taking seriously the actual epistemic context of science, where theories are rarely perfectly articulated and logical vision is limited, provides a clearer understanding for the value of prediction over accommodation. We can see this across a range of inferential contexts in science.

## 2.1 Data-phenomena relations: The danger of overfitting

Fitting a curve to data requires moving from a plot of discrete, particular observations to a formula which correctly defines the relation between quantities; that is, it is a matter of using the data to figure out the phenomenon. As a matter of mathematics, one can almost always write down a high-order polynomial which passes precisely through each and every data point.[7] Yet this is never what one actually does. It would be a mistake to make the curve match the data exactly, because we know that the phenomena do not do so. The actual data include error and noise. Indeed, this is one reason for the data-phenomena distinction.

Instead of writing down a baroque, perfect-fit polynomial, one draws a simpler curve that gets pretty close to the data points. There are various formal methods for doing so. Sober and coauthors apply one in particular, the Akaike Information Criterion (AIC), to defend the value of simplicity (Forster and Sober, 1994) and novel prediction (Hitchcock and Sober, 2004). Given some data generated with random error, one can do a better job predicting *future* data by drawing something simpler than the best-fit polynomial. As the amount of data increases, one can justifiably draw a higher-order better-fit curve. The

---

[6] See also Collins (1994), p. 222.

[7] We say 'almost always' because there may be data points with the same x value but different y values, and then y cannot be represented as a function of x. Throwing out troublesome points, we might instead say: one can always write down a polynomial that passes through almost all of the data points.

8

AIC provides a quantitative procedure for deciding how complicated the curve should be.

The connection between AIC and predicting future data makes a (qualified) case for predictivism. If a scientist overfits the data with a grotesque curve, then he will make poor predictions about future data. Contrapositively, if a scientist makes successful predictions, then we can infer that he has not overfit the data. Although AIC is easiest to illustrate with the example of fitting polynomial curves to plots of data, it readily generalizes to cases where some amount of data is treated with models that have some number of parameters. Too many parameters will be a sign of overfitting, and predictive success will be a sign that there are not too many parameters. AIC provides a precise way to determine when the amount of data is sufficient to justify the better fit provided by an additional parameter. Because AIC provides a way of determining this transparently, however, this is only a qualified defense of predictivism. Hitchcock and Sober argue that there is no further value to novel prediction once we know that AIC has been met.

Although AIC can be applied generally to any case of fitting n-parameter models to N points of data, such cases are only for inferences from data to phenomena. The examples which Hitchcock and Sober give underscore the fact that we are not talking about *theories* in any strong sense; e.g.:

> Marsha measures the width of her desk with a tape measure. She holds the end of the tape with '0' on it at one end of the desk, and observes that the other end of the desk coincides almost exactly with the line on the tape labeled '150'. On the basis of this observation, she hypothesizes that the desk is between 149 and 151 centimeters wide. (Hitchcock and Sober, 2004, p. 6, also p. 24)

They offer this example to show that accommodation is not always bad; it seems OK that Marsha forms her belief about the length of the desk after making the measurement. The data is one measurement made with a tape measure; the implicit one-parameter model is just that the desk has a width $x$. The conclusion made from this measurement is that the actual parameter value is close to 150; that is, the desk is approximately as wide as it was measured to be. The width of the desk is a phenomenon, because it is the real value which must be inferred from the data. Even more general facts, like the mean width of desks on campus, would still be phenomena in this sense. It is far from being a theory of anything. The general problem of curve fitting, although more complex than just measuring desks, is still primarily a matter of

figuring out the phenomenon on the basis of data. Hitchcock and Sober's argument shows at most that the advantages of novel prediction *at this level* can be replaced by the machinery of the AIC, if it is properly applicable.[8]

Churchland (1989, esp. pp. 179–181) makes a point regarding over-fitting which is similar to Sober *et al.*, although in an entirely different formal framework. For Churchland, the data are represented by the input nodes of a neural network. Network weights are set by a learning algorithm in response to a training set of data. Each input in the training set has a normative output. Churchland gives the example of a network trained to distinguish rocks from mines based on sonar data. If the network has too much internal structure, the weights can effectively encode each input from the training set; but a network which has just learned the idiosyncrasies of the training set will not give correct outputs when faced with novel data. A network with simpler internal structure cannot master the training set just by encoding facts about each distinct data point; to succeed, it must find a general pattern. Although Churchland calls the representation at the hidden layer a 'theory', in our terminology the general pattern is clearly the phenomenon.

For Churchland — as for Sober and Forster — a needlessly complex account is likely to *not* be predictively successful because it is likely to have overfit the data. As such, an account that *is* predictively successful is shown not to have overfit the data and so not to be needlessly complex. Successful prediction is an assurance that one has not overfit the data, and that future predictions will also be successful. Such assurance is obviously very valuable in science. The question is whether prediction is a mere surrogate that can be generally replaced by a formula like AIC or an analysis of network complexity. We think not. Instead, novel prediction provides a crucial supplement and support to the proper use of tools like the AIC. Here are some reasons why:

---

[8] The adaptationist approach to growth, discussed above, illustrates this. It is clear how AIC might be used to identify relations between body mass, food intake, and physical performance. Given data, it can be used to identify the phenomena. Yet performance at specific physical tasks is used as a proxy for *fitness*, a term which appears at the theory level and which cannot be measured directly. It is not clear how the AIC could function at this level, in relating theory to phenomena.

First, the AIC only applies if certain assumptions hold (see Forster and Sober, 1994). Although the number of data points and free parameters can readily be counted up, we can never be entirely sure that the formula applies. Whether or not the assumptions needed to apply AIC are present cannot itself be shown by using AIC. Different, external assurance is required. Novel prediction can provide some of that assurance, showing us that AIC and how we are applying it is on track.

Second, applying AIC requires dividing models into natural families. In the curve-fitting case, we can treat orders of polynomials as the families of models. Yet a high-order polynomial can be made into a model with one free parameter by writing in all but one of the coefficients. It would be perverse to do this arbitrarily, but in the general case of n-parameter models it may be possible to motivate specific values for some of the parameters. There is no formal rule for when this is or is not legitimate. Forster and Sober call this the *subfamily problem*. Identifying the families of legitimate models (i.e., solving the subfamily problem) requires constraints from what we have called the levels of theory and framework — levels where AIC does not readily apply. Appealing to AIC as grounds for thinking that accommodation is just as good as prediction, in a specific case, thus relies on background commitments which themselves cannot be based on AIC. Novel prediction can help provide crucial assurance that the way we are thinking about families of models is appropriate, thus again assisting with the application of AIC.

Third, the general idea of prediction as insurance against over-fitting does not require the AIC approach; recall that Churchland makes the same point in a very different formal idiom. This suggests that the precise judgement which AIC yields is not decisive. It is not the singular algorithm for determining when prediction is of no value over accommodation. Instead, it is one (nice and precise) way of assessing the assurance against over-fitting that novel prediction also generally supplies. Moreover, the fact that insurance against over-fitting is more general than AIC provides some license for appealing to novel predictive success where, for technical reasons, AIC does not apply.

One may object that these are quibbles. After all, we agree with Hitchcock and Sober that this advantage of prediction over accommodation, insurance against overfitting, diminishes as the amount of data increases. In extreme cases, where we take phenomena to be established matters of fact, a prediction that the pattern will continue carries negligible epistemic weight (but then so does any additional accommodation). Measuring the length of a desk and then

predicting that the desk will have the same length tomorrow is unimpressive, because the phenomenon is entirely familiar. Where we need epistemic assurance is not in the cases of the familiar, but rather in the cases of the unfamiliar. In cases where the phenomena are not an entirely settled matter, confirmation of a novel prediction will always retain some value. Minimally, it will add evidence. Additionally, it can bolster our assurance that the more formal AIC approach (or some other formal approach) is being applied properly. And in cases of inference from data to phenomena that are more qualitative, to which AIC simply cannot apply, novel prediction can still be a source of assurance that we have neither oversimplified nor overfit. Novel prediction thus typically provides more assurance than accommodation in inferences from data to phenomena.

## 2.2 Theory-phenomena relations: The power of relevance

What of inferences from phenomena to the explanations and theories postulated to account for them? Novel prediction offers several epistemic advantages over accommodation for inferences made here as well.

For the purposes of this survey, we take it that theories (including models and laws) relate to phenomena primarily through explanatory relations. Explanations explain phenomena by invoking theoretical machinery such as causal relations and laws, and thus show how these theoretical ideas link up with phenomena on the ground. Theories, and the explanations that invoke them to explain phenomena, help us to conceptually organize an otherwise chaotic world of phenomena. In the generation of theories that identify these explanatory relations, accommodation plays a key role. Already-known phenomena constrain scientists' creativity in generating new explanatory ideas, which are ultimately the source of new theories. But as with the relationship between data and phenomena, prediction (usually) goes one step further and (typically) provides additional epistemic merit.

The close working relationship between explanation and prediction is central to the generation of this epistemic merit. The explanations at the heart of a scientific theory not only organize already known phenomena, they are also the tools that should be used to predict new phenomena. As Douglas (2009) has argued, all scientific explanations have cognitive aspects that enable us to move from what is postulated by a theory to the next novel prediction. Examining deductive-nomological, causal, mechanistic, and unifying explanations, Douglas shows how each kind of explanation is the generator of predictions in a

well-functioning theory. Phenomena are predicted through the use of explanations, and the predictions then guide tests to provide evidence for the theory. The prediction thus plays a crucial part in testing the explanation and the theory of which it is a part.

Successful novel predictions tell us several things about the explanations (and theories) that produced them. First, most obviously, novel predictions add additional pieces of directly relevant confirmatory evidence for the explanations. New evidence, particularly evidence of new relevant phenomena, is a key epistemic bonus in science. New evidence or new evidential relations throw an evidential gauntlet down for competitor theories — they must accommodate the new evidence or risk epistemic demerit. Accommodation is usually possible, but the point here is that accommodation did not produce the new evidence nor did it bring new evidence into the purview of the theory, showing that it was relevant to an area previously unconnected to it.

Note that this advantage obtains even if the predicted fact was already known, provided it was not used in the construction of the theory; i.e., if the prediction was only *use-novel*. In such cases, the evidence was not connected to the particular theoretical domain before, not brought under any competitors' explanations. Seeing even previously known evidence as relevant to a new area is an important explanatory and epistemic accomplishment.[9]

Second, novel predictions provide some inductive ground for confidence in the ability of an explanation to help guide us to further new predictions. In turn, this can open up new evidential terrain for exploration. An explanation that has been used to produce a successful novel prediction has proven itself fruitful, giving us reason to think that it can help us think productively about the world. Such an explanation is a powerful tool, which matters greatly to scientists, who need such conceptual tools. Novel prediction here gives us an assurance about the *evidential* fecundity of the explanation and theory that produced it which accommodation cannot. Of course, this assurance is fallible. It might be that, next time, the explanation fails to produce successful novel predictions. But its past success is a modest inductive assurance.

Third, and finally, novel prediction gives us some assurance that we have not fiddled around too much with our core explanatory structures

---

[9] If the evidence is not new to the domain, it is unlikely it would be use-novel. Only in cases of formal parameter setting, much discussed by Worrall, can good cases be made that *known-to-be relevant* evidence is indeed use-novel. More on this below.

to make them fit the available phenomena. (Lipton, 2004, p. 170) Just as projections of phenomena can be overfit to data, theories can be overfit to known phenomena. Unlike the curve-fitting context, however, there is no formal measure like AIC to help at the level of explanations and theories.[10] Lacking such a formal measure, novel prediction can provide a crucial guard against overfitting. Accommodation provides no such assurance. Novel prediction is thus a potent tool for constraining our explanatory imagination, because it forces us to pit our favored explanations and theories against the world in a way that accommodation does not. Novel predictions are a crucial way of telling the extent to which our theories may be tuned too much to the available evidence and not projectable beyond it; they provide a piece of instrumentally valuable assurance.

The link between explanation and prediction which gives us confidence in explanatory theories that successfully predict also makes it clearer why predictive flukes are worrisome. Predictive success severed from explanatory or theoretical contexts makes it very unclear what we are supposed to be assured about. Novel predictions provide assurances for the explanations and theories we use to generate those predictions. Divorced from this context, it is hard to know what to think. For example, consider coin-tossing cases. Imagine a series of coin tosses and a prognosticator who predicts the outcomes successfully, such that we do not know how the prognosticator is managing it. The predictive success appears free-floating and, plausibly, a fluke. Barnes illuminates such cases by expanding the example to include a whole community of coin-flipologists. If that community is large enough, our intuitive trust in a successful predictor — who now just looks like a lottery winner — drops dramatically (Barnes, 1996, p. 74). But specific coin toss results are data rather than phenomena, so consider another example. Bode's Law was a mathematical formula that accommodated and then predicted the size of the planetary orbits, clearly phenomena rather than data. It was a modest predictive success before 1846, predicting the distance to

---

[10] It is impossible to group together classes of theories, which would be required to solve the relevant subfamily problem; e.g. Newtonian mechanics might be one or many theories, depending on how we organize constraints and force laws. Additionally, it is often not clear how to count the number of free parameters; the Newtonian description of a specific system has a determinate number of parameters, but Newtonian systems generally do not.

Uranus and the location of the asteroid belt.[11] Without an explanatory theory, it was easy to think of the rule's two predictive successes as flukes — particularly once Neptune was discovered, far from where Bode's Law said it should be. If some mechanism had produced the regularity and if our theory of that mechanism also successfully predicted other phenomena, it would not have been dismissed so quickly. With both Bode's Law and the coin tossing predictions, we have no theoretical or explanatory claims to which we can tie the predictive success.

Contrast this with the case of Fresnel and the ability of his wave mechanics for light to 'predict' diffraction fringes.[12] As Worrall (1989) notes, there was only one parameter to fix, the wavelength of light, to make the theory fit the phenomena. Once Fresnel had a measure of the wavelength of light he was using, he could use his mathematical theory of diffraction to predict diffraction fringes from any configuration, and very precisely. So, Fresnel's theory can be considered to produce the use-novel prediction of the diffraction measurements as well as the novel prediction of the spot. All provided stunning support for his theory, and given the precision of the diffraction measurements and the numerous predictions made and measured there, it is not surprising the French Academy's prize commission of 1819 would focus on them. The explanatory structure of wave interference (which causes diffraction patterns), made precise in Fresnel's theory, produced novel and successful predictions (some use-novel, others temporally novel) and garnered profound support.

Consider also the case of Mendeleev. The theoretically thin predictions which he made on the basis of his periodic table at least had the periodic law behind them — some explanatory apparatus, plus a plethora of diverse accommodated evidence. His predictions came directly out of thinking through the implications of the periodic law, including the revised atomic weights and additional elements (some of which were erroneous predictions, as Scerri and Worrall (2001, p. 439) point out). Initially, the noble gases were thought to be a threat to

---

[11] For more details, see McIntyre (2001), p. 315-316.

[12] We use cautionary quotes because it is a prediction in the use-novel sense only; that is, he did not use the specific diffraction data to derive his theory. Because this is a case of precise parameter fixing, whether Fresnel had the specific diffraction fringe measurements in hand before developing the mathematics was irrelevant. Although a case of use-novel prediction, it illustrates the value of novel prediction generally.

Mendeleev's theory. Yet it was quickly suggested that one needed to simply add another column to the periodic table to accommodate argon. Successful predictions of neon and other noble gases swiftly followed (Scerri and Worrall, 2001, p. 442–447). This case shows the tight relationships between accommodation, explanation, and prediction when the theory is working well. Minimal accommodation, when needed, should lead quickly to new predictions, which then bolster our confidence in the explanations provided. If the addition of a column to the periodic table had not yielded successful new predictions, it would have looked like an *ad hoc* strategy merely to save the table. Because of the successful predictions, it looks like a legitimate move.[13] In scientific practice, accommodations need to prove themselves with additional predictions.

Novel predictions of phenomena, predicted using the core explanations of a theory, provide the following epistemic goods: (1) new relevant phenomena are discovered and/or known phenomena are shown to be relevant to a theory, (2) the evidential fruitfulness of the explanations and their reliability in helping us navigate the world is demonstrated, and (3) we are assured that we have not been excessively creative in our development of any needed accommodations, explanations, and theories, i.e. that such conceptual work is not 'overfitted'. Accommodation has nothing to offer for (1) or (2), and it is worry over accommodative practices that is the reason we need prediction to provide assurance of kind (3). Novel predictions provide us with particular, if imperfect, epistemic assurance over accommodations when examining theories and the phenomena they try to explain.

## 2.3 Theory-framework relations: Seeing the big picture

Finally, consider inferences from theory to framework.

Eric Barnes (2008) advocates a version of the realist *no miracles* argument which he dubs the 'miraculous endorsement' argument.[14]

---

[13] This illustrates a point we emphasize in our discussion of weak predictivism below: Novel prediction bolsters our confidence more than direct assessments of simplicity, which are contentious. In the Mendeleev case, was adding another column simple? Simplicity of theories is a cipher, whereas the successful novel predictions can at least readily be seen as such.

[14] Parts of this section follow the exegesis of Barnes given by Magnus (2011).

Rather than arguing for the existence of specific phenomena or the truth of specific theories, the miraculous endorsement argument is concerned with the success of scientific practice in general. It roughly takes this form: Scientists make predictions which are verified. If these scientists were merely guessing or accommodating prior evidence, then this success would be an inexplicable miracle. Alternatively, this success could be explained by scientists' general reliability. So scientists are (probably, mostly) reliable. The reliability of their methods only makes sense if scientists' background theories are true. So those background theories are (probably, approximately) true.

As Barnes acknowledges, this argument echoes ones made by other philosophers. Richard Boyd, for example, argues that the general reliability of the frameworks employed by scientists is the best explanation of their success in developing particular theories.[15] Peter Lipton (1996) also draws a connection between reliable methods and true background beliefs.[16]

It is important to note that the miraculous endorsement argument is not about a particular, predictively-successful theory. Rather, it is about the ability of scientists to develop such theories. It explains their ability in terms of the general reliability of their methods, methods which are reliable only because they correspond to some true beliefs. The argument looks not at specific, isolated theories but instead at the background of theories which underwrite scientific methods — that is, not at what we have called *theories* but instead what we have called *frameworks*.

One can think of the framework as scientists' *background beliefs*, as Barnes does when explicating his argument; he says "novel success is not direct evidence of theory truth — it is rather evidence for the credibility of the endorser. More specifically, it is evidence for the truth or empirical adequacy of the endorser's background beliefs" (Barnes, 2008, p. 140). Barnes' claim, both in this passage and throughout his book, is that novel predictive success is a reason to conclude that a scientist is a reliable expert in her domain. This expertise will partly be a matter of background belief, in the sense of believing true propositions, but it will also be a matter of 'background belief' in a broader sense which includes approaches to theory development and

---

[15] See e.g. Boyd (1982); importantly, this argument is different than the one often called the 'Boyd-Putnam argument' for realism.

[16] Lipton's argument is that judging the merits of specific theories requires reliable and hence approximately true background theories.

tacit commitment to reliable methods. What we have called the level of theory consists of explicit claims and models, but the level of framework consists of a less easily circumscribed raft of commitments.

Following Barnes: if a scientist is able to generate predictively-successful theories, then this gives us reason to trust or adopt the framework which she employs in doing so. Notice here that the relevant evidence is not merely facts about data and theories — when data was observed, how the theory was initially formulated, and so on. As Barnes emphasizes, what matters is scientists' endorsement of theories which ultimately enjoy predictive success. That is what reflects well upon scientists' background commitments. This is what Barnes calls *endorsement novelty*.

Although Barnes, Boyd, and Lipton all give these arguments in the context of thinking about scientific realism, the assurance provided by novel prediction can be acknowledged while remaining neutral on the question of realism. Barnes recognizes that winning the day for realism requires solving further problems. In the passage we quote above, he allows for the possibility that predictive success might just be explained by the *empirical adequacy* of background theories. Instead of concluding that methods and commitments are truth conducive, one might infer instead that they are instrumentally powerful at producing reliable theories.

The anti-realist can grant to Barnes that scientists should and do endorse theories that are predictively successful. Moreover, scientists' being able to do so reflects well upon the background framework the scientists employ in developing these theories. Of course, the anti-realist will insist that this only gives us reason to believe that the framework is empirically adequate — not that the framework's claims about the unobservable world are true.[17] Both realists and anti-realist empiricists can recognize a framework which consistently delivers predictive success as reliable, and both will prefer reliable frameworks over unreliable ones. The realist just understands *reliability* as the tendency to generate true theories, whereas the anti-realist understands it as the tendency to generate empirically successful theories.

The arguments given by Barnes, Boyd, and Lipton concur in saying that scientists' development of predictively successful theories reflects well upon scientists' framework commitments. For them, this favorable light justifies belief; for the anti-realist, it justifies acceptance instead. In either case, novel predictive success does important epistemic and evidential work. The fact that a framework generates predictively

---

[17] This is what Barnes calls *the anti-realist challenge*.

successful theories gives us *prima facie* reason to adopt it or (if we have already adopted it) to stick with it.

This assurance is parallel to the one that we get when a theory generates novel and accurate predictions of unanticipated phenomena. That gives us a reason to adopt or continue employing that theory. When a framework or research programme generates theories which can enjoy that kind of success, then we have a reason to adopt or continue employing that framework. In the example from Schell and Magnus (2007) discussed above, the rival adaptationist and medical approaches are not things that can be directly tested. Rather, they inform the construction of specific theories. If practitioners employing one framework can generate predictively successful theories but practitioners employing the other cannot, then we have a reason to prefer the former framework to the latter. Schell and Magnus diagnose this situation for the medical and adaptationist approaches respectively.

Schell and Magnus draw an explicit connection to Lakatos (1978). Indeed, there are parallels between predictivism for theory-framework relations and what Lakatos called *the methodology of scientific research programmes*. To use his terminology: A research program is progressive if it leads to theories which underwrite novel predictions and if at least many of those predictions are accurate; a research program is degenerating otherwise. Whereas Lakatos offered this as a general account, suggesting that progress and degeneration were the *only* considerations for or against a research program, we offer the more modest claim that they do matter. Progress in Lakatos' sense is both methodologically and epistemically to a framework's credit.

An additional element introduced by Barnes' argument is that — unlike data, phenomena, or theory — a framework might be largely a matter of implicit commitments. Where that is so, we cannot write down the framework that receives credit for the scientist's success. Instead, the success accrues to their tacit expertise; that is, to the scientist as expert. Realists and antirealists disagree about whether the expertise is a matter of truth or empirical reliability, but it is a legitimate scientific merit in either case. And it is novel prediction that gives us such assurance, not accommodation.

## 3  Weak Predictivisms

As noted in the introduction, defenses of novel prediction may seem somewhat unfashionable. Many philosophers have argued that novel predictive success is nothing special, that it is better understood merely as a guide to some other confirmatory virtue and that the exploration of

that other virtue is where the real philosophical action is. Other philosophers, most notably John Worrall, substantially reinterpret what counts as prediction. On Worrall's account, the intuitive situation in which a scientist predicts something previously unknown becomes just a special case — and not even the most important case. We now turn to these deflationary and ersatz predictivisms.

## 3.1 Would-be reduction and deflation

Some contemporary philosophers admit that predictive success might *seem* to have some value, but they suggest that it is just indicative of something more important and that the importance of predictive success reduces to this other more important thing.

Harker (2008) and Collins (1994) are exemplary of such views. Both insist that prediction is just a superficial proxy for considerations of explanatory unification and that the theoretical virtue, explanatory or unificatory power, is what really matters. Collins writes that "the ability of a theory to make correct new predictions is not what is ultimately important; rather, what is important is the theory's general structural/relational features" (Collins, 1994, p. 222). Given the overview of the instrumental value of prediction given above, there are several problems with their approach to reducing the value of prediction.

First, although explanatory power and unification are important theoretical virtues, the accounts that Collins and Harker provide fail to show that novel prediction's value reduces to such virtues. While novel prediction may sometimes or even often accompany virtues such as unification, explanatory power, and simplicity, it seems possible that a theory could be predictively successful without exhibiting them. Indeed, in the early stages of theory development, novel predictive success may be our main indicator that there is something worth pursuing further. Unification can come later. The assurance novel prediction provides does not simply dissipate in a context where the other virtues are absent or contested.

Second, the strategy of reducing novel prediction to a theoretical virtue trades a marker that is relatively clear with a marker that is problematically more vague. As Kuhn (1977) emphasized, throughout the history of science there is a substantial degree of subjectivity in the assessment of epistemic virtues like scope, simplicity, and explanatory power. Individual judgements differ, and community standards change over time. Thus, whether a theory has more or less explanatory or unifactory power than its competitors can be a point of disagreement

among scientists. Novel prediction is less contestable. The presence of successful novel prediction is more clearly assessable than explanatory power or unification even if there is disagreement about how strong a test the novel prediction provides. Take the example of quantum mechanics: It made successful predictions from the beginning, but the current assessment that it is explanatory arose partly because judgments about what is explanatory have matured in response to quantum mechanics.

Third, critics of predictivism often presume an epistemic situation in which a theory is fully articulated and the theory's structural virtues are transparent. Indeed, Collins insists that "the fundamental issue underlying the accommodation/prediction debate is whether or not the information that a theory predicted instead of accommodated a set of data should increase our confidence in its truth *given that we already know the relevant structural/relational features of the theory*" (Collins, 1994, p. 215, our emphasis). As noted in the beginning of sec. 2, such perfect transparency of the relevant features of a theory is hard to come by in the midst of scientific practice. It is a philosophical idealization, so recommendations contingent on it provide no guidance for actual practice. In the flow of practice, because the structural features of a theory are never all transparent, novel prediction can provide useful assurance.

Fourth, treating explanatory resources as features of a theory *tout court* ignores the important connection between explanation and prediction in the practice of science. Rather than being merely abstract features of a theory, explanatory resources are a crucial predictive resource, as noted in sec. 2.2 above. (And vice versa: Novel predictions, successful or not, provide grist for developing explanations.)

Fifth, the theoretical virtues are primarily confined to one level of inference, the relationship between the theory and the phenomena. As we elaborated above, scientific inference also occurs in relating data to phenomena and in relating theory to frameworks. Novel prediction provides epistemic assurance at these levels of inference as well, as described in sec. 2.1 and 2.3.

One might accept some of our objections and attempt to reframe Harker's argument. If we cannot assess theoretical virtues directly, then we must use proxies. But why, Harker asks, should be use novel predictive success rather than other available proxies? He writes, "Supporters of weak predictivist theses have done nothing to establish that the advantage of predictive success over accommodative success is any more reliable at tracking the epistemic virtues of a theory than

considerations pertaining to the reputation of the author, advocacy of the expert and so on" (Harker, 2008, p. 438). The argument is that predictive success, although some indicator of confirmatory virtues, might not be the best indicator available. Harker suggests that considerations of expertise might be better.

Of course, this objection might well be decisive in some particular cases. We have only argued that predictive success can often do and typically does epistemic work. Yet Harker's worry supposes that predictive success and expertise can, in general, be treated as separate indicators. Barnes' argument (Barnes, 2008, discussed in sec. 2.3 above) shows that this separation is untenable.[18] A theory's predictive success provides assurance for the framework which led to the theory, and some of this credit extends to the scientists who employed the framework, developed the theory, and endorsed the prediction. Because the framework involves tacit as well as explicit elements, this recognition of expertise cannot be eliminated in favor of purely formal confirmation of general theories. Novel predictive success interacts with expertise, and thus it would not be possible to jettison novel prediction while retaining considerations of expertise.[19]

The kind of reduction proposed by Harker and Collins is not the only one in the field. Mayo (1996) attempts a different reduction of prediction's virtues. She argues that novel prediction is only valuable because it is frequently associated with severe tests; that is, tests where the hypothesis under consideration would not have passed the test had the hypothesis been false. Novel prediction, on her view, is thus just an indicator for what really matters: viz., severe tests. Here too, the reduction fails for scientific practice. In order to assess whether a test was in fact severe, we need to account or control for all the sources of error. While this kind of assurance can be gained in controlled experimental contexts, it is much less clear how we are to achieve this in more complex contexts. Indeed, the applicability of Mayo's approach to all of scientific inference is a central concern for her critics; see recent discussions by Achinstein, Musgrave, and Worrall (in Mayo and Spanos, 2010). Thus, the value of novel prediction across the full range of scientific contexts resists Mayo's reduction.

---

[18] At minimum, Harker's claim (Harker, 2008, cited above) that supporters of predictivism have nothing to say about the relation between novel predictive success and expertise is now false.

[19] In any case, the revised argument avoids at most the first three of the five worries about deflation that we raise earlier in this section.

## 3.2 Worrall's Predictivism

John Worrall has been one of the most original commentators on the value of novel prediction. Worrall's view focuses on the assessment of theories and issues of parameter fixing. Within this context, he is concerned with the value of *heuristic* or *use-novel* predictions: the 'predicted' claim might already be common knowledge, as long it does not figure in the actual construction of the theory. For Worrall, this is a logical rather than historical matter. In some sense, for Worrall, "novelty is not the issue at all" (Worrall 2002, p. 194). Importantly, timing is irrelevant. For example, Fresnel's wave theory of light had one free parameter (wave length) which could be fixed using experimental data (measuring the wave length with which the scientist is working). With this parameter fixed, Fresnel's theory could be used to derive other experimental results, such as straight-edge diffraction. Even though straight-edge diffraction was a known phenomena to Fresnel when he constructed his theory, the derivation of the diffraction pattern using the wave length and Fresnel's theory counts as a use-novel prediction. For Worrall, it is this kind of use-novelty which provides genuine empirical support to a theory, and nothing else does.

Worrall maintains that accommodating data within a theory is appropriate *given* that we have independent reasons to accept the theory, and he holds that use-novel predictions provide those independent reasons. Thus, it is fine to use data to fix a specific parameter in a set of equations which are otherwise well motivated. Yet this motivation is conditional. Worrall explains, "*Given* that a general framework… is already accepted, then the data give… conclusive support" (Worrall, 2002, p. 203). (Here he means 'framework' to be something that is fully explicit, and so it is closer to what we here call a *theory* than what we call a framework.) In order to motivate accepting a theory, however, there must be some *other* kind of evidential support, and he suggests that there is a special value for predictions which were not exploited in constructing the theory itself. The special value of these use-novel predictions explains how evidence can justify accepting such a theory in the first place.[20]

In one sense, Worrall's predictivism is of the strongest sort, in that only use-novel evidence counts as genuine support for a theory. Yet, the structural aspects of Worrall's view limit its applicability. Worrall's view cannot be the whole story. First, it is too stark. Where it

---

[20] This is a long-running theme in his work. For a clear, recent statement, see Worrall (2002).

acknowledges a special value for prediction, it cannot allow accommodated evidence to count for anything. Second, applying it requires being able to clearly assess whether a prediction is use-novel. This requires a clear mathematical structure to a theory and a clear assessment of its parameters, clarity which is often lacking. So, despite its merits, Worrall's account cannot be the whole story. Recognizing the value of novel prediction outside of the confines of Worrall's view requires that we once again consider the temporal aspects of novel prediction, as we did in sec. 2. We add that temporal, endorsement novelty assures us of use-novelty, and so it provides the support that Worrall articulates in cases where we cannot independently assess use-novelty. Strict use-novelty is hard to assess when parameters and the mathematical structure of the theory are not transparent, so true temporal novelty can provide the epistemic support Worrall seeks for less precise theories.

There is a broader lesson here. Attempts to reduce the issue of predictivism to the logical structure of theories and evidence or particular epistemic virtues, often by demanding perfect transparency of evidential theory relationships, will miss precisely what is so valuable about predictions: When these things are obscure, when we do not know that we have accounted for all sources of error, then novel prediction provides the a useful epistemic assurance that we are on the right track. It is often only possible to characterize scientific success in precise terms long after it is a settled matter, so genuine novel prediction is an important resource in the flow of practice.

## 4 Pluralist instrumental predictivism

The survey provided here shows that prediction provides epistemic goods and additional assurance above and beyond accommodation at a number of points in scientific inference. For inferences from data to phenomena, from phenomena to theories, and from theories to framework, novel prediction provides distinct epistemic advantages. These epistemic assurances include: insurance against overfitting, evidential relevance to the explanatory structure of a theory, and reliable production of successful theories. In addition, novel predictions provide other epistemic goods, such as uncovering new relevant evidence and new evidence/theory relations.

These assurances are all defeasible, but the point of epistemic assurance is not to be infallible. Rather, it is to indicate that our knowledge production is on the right track and likely to be reliable.

Novel prediction is neither a perfect warrant nor universally valuable, but rather it is commonly, locally valuable.

As we have seen, some authors defend predictivism but consider predictive success to have just one function. We have cast our net more widely and have found more ways that predictions provides epistemic assurance. We make no claims to completeness, though. We have discussed the advantages of prediction under three headings, but we do not think of novel predictive success as having a perfect trinity of virtues. Predictions may be connected to the explanatory structure of a theory in several ways, for example, and so the advantages at the theory-phenomena boundary might themselves be numerous.

The value of novel prediction is especially clear for inferences in actual, everyday science, because the complete structure of a theory under development is unavailable and because the features of a framework in use are not fully articulated and may remain tacit. Even though novel predictions are not a perfect assurance, and are not even optimally-assuring in all contexts, they often provide epistemic assurance that cannot be easily or feasibly bought in other ways. More philosophically-pure epistemic checks are often expensive or would require ideal circumstances. In contrast, for real cases, novel predictions provide a readily available and applicable epistemic check. For the state of incomplete knowledge in which epistemic actors actually find themselves, the assurance of novel prediction has clear advantages over attempting to assess the qualities of accommodation.

Thus, novel prediction typically does give us reason to credit both the scientist who made the prediction and the representations used in making it. This reason is not because of an intrinsic feature of novel predictive success, but neither is it reducible to any one other dimension of confirmation. Novel success is typically an indication of numerous confirming factors. Attempts to make these indications superfluous will usually fail, both because there are so many disparate indications involved and because scientists are often unable to infallibly make direct assessment of them. Novel predictive success has an *instrumental* rather than an intrinsic value, but its typical instrumental value is sufficiently multifarious that no other tool can readily replace it. We dub this position *Pluralist Instrumental Predictivism*; PIP for short.[21] We predict that developing the various

---

[21] We use 'pluralist' only to mean that there are many instrumental values for prediction, not to invoke any broader philosophical pluralisms. We use 'instrumental' only in opposition to 'intrinsic', not to suggest anti-realism about the prediction or theory.

advantages of prediction provides across a range of contexts will be a fruitful avenue for understanding scientific inference in practice.

In sum, for scientific inference at various levels — from data to phenomena, from phenomena to theory, from theory to framework — novel predictions generally provide better epistemic assurance than accommodation. The epistemic assurance novel prediction provides, imperfect as it is, is of such general applicability in such a wide range of contexts and is so difficult to replace with anything else that novel prediction should be considered generally instrumentally valuable to scientific practice. Even if novel prediction is a surrogate in some cases for other strong sources of assurance (such as AIC in inferences from data to phenomena or severe tests in experimental contexts), it may be more readily available in particular cases and thus can still provide valuable epistemic assurance. The correct way to view the value of novel predictions is one of pluralist instrumentalism, acknowledging it as an important epistemic resource in scientific practice.

## References

Barnes, E. C. (1996). Social predictivism, *Erkenntnis* **45**, 69–89.

Barnes, E. C. (2005). On Mendeleev's predictions: Comment on Scerri and Worrall, *Studies in History and Philosophy of Science* **36**, 801–812.

Barnes, E. C. (2008). *The Paradox of Predictivism*, Cambridge University Press.

Bogen, J. and Woodward, J. (1988). Saving the phenomena, *Philosophy of Science* **97**(3), 303–352.

Boyd, R. N. (1982). Scientific realism and naturalistic epistemology, *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association (1980)* **2**, 613–662.

Brush, S. (1994). Dynamics of theory change: The role of predictions, *PSA 1994* **2**, 133–145.

Brush, S. (1996). The reception of Mendeleev's periodic law in America and Britain, *Isis* **87**(4), 595–628.

Brush, S. (2007). Predictivism and the periodic table, *Studies in History and Philosophy of Science* **38**, 256–259.

Churchland, P. M. (1989). *A Neurocomputational Perspective: The Nature of Mind and the Structure of Science*, The MIT Press, Cambridge, Massachusetts.

Collins, R. (1994). Against the epistemic value of prediction over accommodation, *Noûs* **28**(2), 210–224.

Douglas, H. E. (2009). Reintroducing prediction to explanation, *Philosophy of Science* **76**(4), 444–463.

Earman, J. and Glymour, C. (1978). Einstein and Hilbert: Two months in the history of general relativity, *Archive for History of Exact Sciences* **19**(3), 291–308.

Forster, M. and Sober, E. (1994). How to tell when simpler, more unified, or less *Ad Hoc* theories will provide more accurate predictions, *The British Journal for the Philosophy of Science* **45**, 1–35.

Gardner, M. (1982). Predicting novel facts, *The British Journal for the Philosophy of Science* **33**, 1–15.

Giere, R. (1983). Testing theoretical hypotheses, *in* J. Earman (ed.), *Minnesota Studies in Philosophy of Science*, Vol. X, University of Minnesota Press, Minneapolis, pp. 269–298.

Glymour, B. (2000). Data and phenomena: A distinction reconsidered, *Erkenntnis* **52**, 29–37.

Harker, D. (2008). On the predilections for predictions, *The British Journal for the Philosophy of Science* **59**(3), 429–453.

Hitchcock, C. and Sober, E. (2004). Prediction versus accommodation and the risk of overfitting, *The British Journal for the Philosophy of Science* **55**(1), 1–34.

Kuhn, T. (1977). *The Essential Tension*, The University of Chicago Press.

Lakatos, I. (1978). *Falsification and the Methodology of Scientific Research Programmes: Philosophical Papers*, Vol. 1, Cambridge University Press.

Lipton, P. (1996). Is the best good enough? , *in* D. Papineau (ed.), *The Philosophy of Science*, Oxford Readings in Philosophy, Oxford University Press, pp. 93–106.

Lipton, P. (2004) *Inference to the best explanation*, second edn, Routledge, London.

Magnus, P.D. (2011). Miracles, trust, and ennui in Barnes' *Predictivism*, *Logos&Episteme* **2**(1), 103–115. http://logos-and-episteme.proiectsbc.ro/? q=node/73

Maher, P. (1988). Prediction, accommodation, and the logic of discovery, *in* A. Fine and J. Leplin (eds), *PSA 1988*, Vol. I, Philosophy of Science Association, East Lansing, MI, pp. 273–285.

Mayo, D. G. (1996). *Error and the Growth of Experimental Knowledge*, The University of Chicago Press.

Mayo, D. G. and Spanos, A. (eds): 2010, *Error and Inference*, Cambridge University Press.

McIntyre, L. (2001). Accommodation, prediction, and confirmation, *Perspectives on Science* **9**(3), 308–323.

Musgrave, A. (1974). Logical versus historical theories of confirmation, *The British Journal for the Philosophy of Science* **25**(1), 1–23.

Scerri, E. R. (2005). Response to Barnes's critique of Scerri and Worrall, *Studies in History and Philosophy of Science* **36**, 813–816.

Scerri, E. R. and Worrall, J. (2001). Prediction and the periodic table, *Studies in the History and Philosophy of Science* **32**(3), 407–452.

Schell, L. and Magnus, P.D. (2007). Is there an elephant in the room? Addressing rival approaches to the interpretation of growth perturbations and small size, *American Journal of Human Biology* **19**(5), 606–614.

Schindler, S. (2008). Use-novel predictions and Mendeleev's periodic table: response to Scerri and Worrall, *Studies in History and Philosophy of Science* **39**, 265–269.

White, R. (2003). The epistemic advantage of prediction over accommodation, *Mind* **112**(448), 653–683.

Worrall, J. (1989). Fresnel, Poisson, and the white spot: The role of successful predictions in the acceptance of scientific theories, *The Uses of Experiment*, Cambridge University Press, pp. 135–157.

Worrall, J. (2002). New evidence for old, *in* P. Gärdenfors, J. Woleński and K. Kijania-Placek (eds), *In the scope of logic, methodology and philosophy of science*, Vol. I, Kluwer, pp. 191–209.

Worrall, J. (2005). Prediction and the 'periodic law': A rejoinder to Barnes, *Studies in History and Philosophy of Science* **36**, 817–826.

Zahar, E. (1973). Why did Einstein's programme supersede Lorentz's (ii), *The British Journal for the Philosophy of Science* **24**, 223–262.