

MALCOLM R. FORSTER

**HARD PROBLEMS IN THE PHILOSOPHY OF SCIENCE:
IDEALIZATION AND COMMENSURABILITY**

In R. Nola and H. Sankey (eds.) (2000) *After Popper, Kuhn & Feyerabend: Issues in Theories of Scientific Method, Australasian Studies in History and Philosophy of Science*, Kluwer.

ABSTRACT: In the 1960s, Kuhn maintained that there is no standard higher of rationality than the assent of the relevant community. Realists have seek to evaluate the rationality of science relative to a highest standard possible—namely the truth, or approximate truth, of our best theories. Given that the realist view of rationality is controversial, it seems that a more secure reply to Kuhn should be based on a less controversial objective of science—namely, the goal of predictive accuracy. Not only does this yield a more secure reply to Kuhn, but it also provides the foundation on which any realist arguments should be based. In order to make this case, it is necessary to introduce a three-way distinction between theories, models, and predictive hypotheses, and then ask some hard questions about how the methods of science can actually achieve their goals. As one example of the success of such a program, I explain how the truth of models can sometimes lower their predictive accuracy. As a second example, I describe how one can define progress across paradigms in terms of predictive accuracy. These are examples of hard problems in the philosophy of science, which fall outside the scope of social psychology.

NOTE ON AUTHOR: Malcolm R. Forster was partially educated in Applied Mathematics and Philosophy at the University of Otago, New Zealand, and is currently Professor of Philosophy at the University of Wisconsin-Madison, U. S. A. Interests range from the foundations of quantum mechanics, causal modeling, statistics, neural networks, and phylogenetic inference, to Newton and the Whewell-Mill debate. For more information, visit the author's homepage at <http://philosophy.wisc.edu/forster>

1 INTRODUCTION

Many philosophers underestimate the general disillusionment in the philosophical outlook on science caused, in part, by Kuhn's *Structure of Scientific Revolutions*. The response to Hume's problem of induction has always kept the issue of scientific truth at the forefront of philosophical research, and philosophers have expended great energy in defending a broad spectrum of replies to Hume's skepticism, ranging from the view that theories are merely instruments for the control and prediction of

nature, to realist views of science (which hold that science aims at the truth about the world, and is rational in the pursuit of this goal). Kuhn (1970, p.171) now insists that this approach to studying science is unhelpful, and many outsiders have followed his lead:

Does it really help to imagine that there is some one full, objective, true account of nature and that the proper measure of scientific achievement is the extent to which it brings us closer to that ultimate goal? If we can learn to substitute

evolution-from-what-we-know for evolution-toward-what-we-wish-to-know, a number of vexing problems may vanish in the process. Somewhere in this maze, for example, must lie the problem of induction.

For Kuhn (1970), and his followers, the rationality of science has nothing to do with truth. Rather: ‘As in political revolutions, so in paradigm choice—there is no standard higher than the assent of the relevant community’ (Kuhn, 1970, p.94).

To be a card-carrying philosopher of science it is almost obligatory to reject Kuhn’s point of view. It is natural that any intellectual community attends mostly to the internal issue that divides the community, rather than defending their shared beliefs. Kuhn was right about intellectual communities in this regard. But should they behave in this way? Perhaps, philosophers of science should unite against the common enemy. If so, then the strongest and most convincing rebuttal of Kuhn’s position must be based on the weakest and most secure premises, which are the ones on which most philosophers of science agree. Almost all philosophers of science agree that scientific theories are (successful and rational) instruments for prediction. The strategy of the present essay is to do as much as possible with as little as possible.

Many philosophers of science believe that the many replies to Kuhn are already completely adequate for this purpose. I do not share that conviction. To support my viewpoint, I present a problem to philosophers of science—the problem of idealization (section 5)—which appears to support Kuhn’s view that rationality with respect to truth is a bankrupt notion. The problem is not merely that idealizations are used everywhere in science. The problem is that such falsehoods can actually *increase* the predictive accuracy of the resulting equa-

tions. There is necessarily a need, in some cases at least, to trade off truth in one respect for truth in another respect. It is this trade-off that threatens the rationality of truth as a univocal goal of science.

The traditional ‘solution’ to the problem of idealization is well represented in Musgrave’s writings, especially his 1981. There, Musgrave argues that idealizations are rational from a realist point of view if either they are lead to no losses in predictive accuracy (‘negligibility’ assumptions), or they are only presented for heuristic reasons (‘heuristic’ assumptions) to make it easier to learn the full theory. The solution assumes that if the idealization were removed, then the resulting equations would bring us closer to the truth, or at least no further from the truth. That is, the solution assumes that it is possible to simultaneously optimize ever aspect of truth at the some time, so that truth is a univocal goal of science. It is exactly this assumption that has been shown to be false in recent research on idealizations in science (Forster and Sober, 1994; Forster, 1999; Forster, 2000; Forster, forthcoming).

To formulate and to solve the problem of idealization, one needs to make a clear distinction between three levels of theorizing (section 2)—*theories*, like Newton’s theory of motion, at the most general level, *models* applied to concrete systems in the middle, and *predictive hypotheses* at the lowest level, which result from fitting models to data. The essential point of this tripartite distinction is that predictive accuracy is a property of predictive hypotheses at the very bottom of the hierarchy, and there is traded-off against the truth at the next level up—the level of models.

The same distinction is also useful for the explication of Kuhn’s views about science (section 3). In particular, normal science concerns the development of the middle layer of theory—at the level of models. Revolutionary science involves a change of

theory at the top. If theory change is rationally motivated by the success or failure of normal science, and normal science consists in the development of models, and models are not evaluated according to their truth, then there is a *prima facie* problem here for realists. Kuhn may not have explicitly pointed to the problem of idealization, but it supports his view of science, at least on the surface.

Therefore, philosophers of science who want to defend a standard of rationality higher than the assent of the relevant community must address the problem of idealization. To do this, they need a finer-grained definition of the goal of science than truth *simpliciter* (section 4). Traditionally, philosophers of science have made a distinction between epistemic and pragmatic goals of science. Epistemic goals include all goals that depend on what is true of the world. This includes not only the truth of theories, but also the predictive accuracy of predictive hypotheses. Truth and predictive accuracy operate at different levels of theorizing, so that they depend on each other in complicated ways. The three levels of theorizing are essential to the correct formulation of the problem. If the problem is not understood correctly, it cannot be solved correctly.

Once these distinctions are in place, a space of possible philosophical positions is opened up, and the core instrumentalist view of science is strengthened in the process. It is not only survives the problem of idealization, but it explains the use of idealization in a way that makes *essential* reference to epistemic values. It's not that idealization is epistemically harmless, as Musgrave believes. It has positive epistemic value, which cannot be explained except by reference to predictive accuracy.

Predictive accuracy, like truth of theories, is something that hypotheses do not wear on their sleeves. But unlike the truth of theories, it can be directly tested by see-

ing whether the predictions come out to be true, or approximately true. This requires that a hypothesis constructed from one set of data is tested against a different set of data. This is quite different from testing hypotheses against the combined set of data. The difference might be described as *diachronic* testing as opposed to *synchronic* testing. The suggestion is that models and theories should be evaluated according to their survival of diachronic tests.

Once the problem of idealization resolved, one needs to determine whether the truth of competing theories can be rationally evaluated. The problem that Kuhn presents in this regard is the problem of incommensurability, which, in part, denies the comparability of the theoretical content of rival theories. I have no argument against incommensurability in this sense. Rather, I present it as a non-problem. If theories are to be rationally compared according to their truth, or verisimilitude, then the judgment should supervene on the degree of predictive accuracy that can be obtained within each theory. In section 6, I explain what this means, and describe the difficulties that crop up in making such judgments. Kuhn's incommensurability is not on that list.

In one sense, the solutions presented here are small achievements relative to the wide diversity of methodological issues in science. Modest though they may be, they go beyond the assent of the relevant scientific community in an essential way. They go beyond the assent of any scientific community because our understanding of predictive accuracy is relatively recent—scientists have been unaware of positive *epistemic* benefits of idealization. It is not therefore a part of any story about the psychological goals of scientists, or of community of scientists, or their beliefs. Nevertheless, the payoff is real, and its explanation argues for the rationality of science in the objective

sense recommended by many philosophers of science, and rejected by Kuhn.

To take the argument further—towards establishing the rationality of science with respect to the full realist goal of truth—is an unsolved problem. However, to respond to the common Kuhnian enemy, the first step is the essential one. The modest problem-solving exercise described in this essay is sufficient to establish that the rationality of science should not be conflated with the rationality of scientists.

2 THEORIES, MODELS AND PREDICTIVE HYPOTHESES

Perhaps the easiest way of introducing the distinction between theories, models and predictive hypotheses is to consider how observational predictions are derived from theories. For this purpose, I will suppose that predictions are *logically deduced* from theories. Such an assumption will not be true in statistical theories, which have only probabilistic consequences. However, the idealized picture is sufficient for the task at hand.

Suppose that E is an observational statement about the position of a planet relative to the fixed stars at some particular time, or the frequency of light emitted by burning a certain substance, or the rate at which a species will colonize a new volcanic island. Let T stand for the fundamental theoretical principles involved in making such a prediction, like Newton's laws of motion, the laws of quantum mechanics, or the principles of population ecology. Everyone agrees that it is impossible to logically deduce E from T because there are missing premises. I will divide these additional assumptions into two kinds. First, there are the *background empirical data*—statements of past observation that are used in the theory to fix initial conditions and estimate parameter values. Let me refer to this background data by the letter D ('D' for data). However, there are other

assumptions needed, which are not directly determined by past experience. I will refer to these as *auxiliary assumptions*, denoted by the letter A . On this analysis, a prediction E is deduced from a theory T via the logical entailment $T \& A \& D \Rightarrow E$.

Auxiliary assumptions are typically more theoretical than those included in the background data. They most commonly include simplifying assumptions about the *absence* of interfering factors, like the absence of confounding causal factors in causal modeling, the absence of other forces like air resistance in Newtonian mechanics, the purity of a chemical substance in chemistry, or the absence of genetic mutations in population ecology. Auxiliary assumptions also include the assumptions made by applied mathematicians when they omit high order terms of a Taylor expansion, or when they drop terms on the basis of an order of magnitude analysis. They are often *known* to be false, in which case we refer to them as *idealizations*.

It is important to distinguish between theories and models. Unfortunately, the term 'model' has several unrelated uses in the philosophy of science. Here are three senses in which the term will *not* be used in this essay. (1) A 'model' as in a model airplane. Such models do appear in science, such as in the 'model of DNA' Watson and Crick used to 'model' the helical structure of the DNA molecule. But it is not the sense of 'model' used here. (2) 'Model' in the sense used by mathematicians in model theory (e.g., Sneed, 1971; Stegmüller, 1979). This has a rather technical meaning, which corresponds roughly to what logicians call an *interpretation of a language* (an assignment of objects to names, a set of objects to properties, a set of object pairs to relations, and so on). It is not the sense of 'model' used here. (3) I have heard people speak of Darwin's model of evolution, where they are referring to the core postulates of the theory.

‘Model’ in this instance refers to what we are calling a ‘theory’, and is not the sense of term used here.

I am more concerned with the way in which scientists speak of models. To capture their usage, it is better to say that a *model* is a theoretical statement (often in the form of an equation) that is specific enough to be applied to a concrete system. Theories do not have this specificity. For example, Newton’s *theory* of gravitation says that every body in the solar system attracts every other body in the solar system in a certain way without making any implications about the number or nature of such bodies. Nor does it say whether the system should be treated as isolated, or whether electromagnetic forces play a role. This is the function of auxiliary assumptions. That is, a model M is obtained from the theory T with the aid of a set of auxiliary assumptions A . In symbols, $(T \& A) \Rightarrow M$. Note that the entailment does not work the other way. Models do not ‘contain’ the theory from which it is derived—in fact, it is not essential that they be derived from theories at all. I have explained the meaning of models by their relationship to theory only because that is when we need to be careful about the distinction.

Consider a famous case in the history of planetary astronomy. In the 60 years before the discovery of Neptune in 1846, there was a series of Newtonian models of planetary motion which assumed that Uranus is the outermost planet in the solar system. When the discrepancies between the predictions of this model and the observed motions of Uranus remained after the interactions of the known planets were taken into account, Le Verrier and Adams adopted a model that assumed the existence of an 8th planet. Let me label this new model as M' . M' postulated the existence of an 8th planet, but made no precise assumptions about its position. But when it was combined with the data, D ,

about the past positions of Uranus and the other planets, the new model did predict its position, whereupon Neptune was discovered when telescopes were pointed towards the predicted position of the planet. That is, $M' \& D \Rightarrow E$, where E is a statement about the position of the eighth planet.

A model is unable to make precise predictions because it postulates a number of free parameters, like mass values, or initial conditions, whose values are not given by the theory, or the auxiliary assumptions. Let a *predictive hypothesis* be a *version* of the model together with a precise numerical assignment of values to all adjustable parameters. It is a *predictive hypothesis* because once all parameters have precise numerical values, the hypothesis is able to make precise numerical predictions. There are *many* such versions of the model, so the model is really a *family* of predictive hypotheses. Logically speaking, M is an open-ended disjunction that says that one of its members is true. (Scientists often refer to predictive hypotheses as ‘models’ as well—‘fitted models’ might be the appropriate translation in most instances).

There are two kinds of models—statistical and non-statistical. Philosophers of science mostly think about non-statistical models, which make precise predictions. Most the predictive hypotheses in such an M are logically inconsistent with the background data D . Naturally, we only want the unrefuted members of M to play a role in prediction. Ideally, only one member of M is consistent with D , in which case $M \& D$ is singles out a unique predictive hypothesis. If no members of M are consistent with D , then M is falsified by D . If many members of M are consistent with D , then the predictions will be imprecise.

In the case of models that make only probabilistic assertions, D may be logically consistent with every member of M (*e. g.*, if we assume Gaussian error distributions),

although some members will always fit the data D better than others. In that case, a unique member of M is picked out by choosing the best fitting member of M , where 'best' is defined by some statistical measure of fit, as in the method of maximum likelihood or the method of least squares. Therefore, in either case, *the role of background data is to single out a unique predictive hypotheses from a model*. If we label this predictive hypothesis by H , then $(T \& A \& D) \Rightarrow H$, or equivalently, $(M \& D) \Rightarrow H$.

A theory may be thought of as a family of models. Different models are derived from a theory using different idealizations, different simplifying assumptions, and different auxiliary hypotheses. Many different models can be derived from a single theory. For instance, if we assume that there are six planets, which are small point masses, then we get one Newtonian model of the solar system. But if we assume that there are 7 planets, or if we model the Earth as bulging at equator, then we get a different Newtonian models of the solar system.

Not all theories are as precisely formulated as Newton's or Einstein's theories of motion. For example, connectionist modeling (Rumelhart *et al*, 1986) of animal or human behavior is based on the idea that behavior is caused by information processed by neural networks. There is a collection of basic models, or what Kuhn would call the exemplars of connectionist science, which serve to guide the construction of new models. But there is no well articulated procedure for constructing models from something like Newton's three laws of motion. That is why science at the level of models and predictive hypotheses is perhaps the most important. Models appear in every science, while theories do not. That is why it is essential to include models as a species of scientific hypothesis—to speak only in terms of theory and auxiliary assumptions is to either exclude such sciences from the

discussion, or to conflate models and theories.

3 KUHN'S PICTURE OF SCIENCE

Some of the most notable examples of science are those spawned by the great books in science, such as Copernicus's *De Revolutionibus*, Newton's *Principia*, and Darwin's *Origin*. These books were not the end, but the beginning, of highly productive periods of science. For Kuhn (1970), these periods of science are examples of *normal science*. Normal scientific research is conducted under a *paradigm*, or *disciplinary matrix*. He lists four elements of a disciplinary matrix; *symbolic generalizations*, *metaphysical presumptions*, *values*, and *exemplars*. I will assume that a model is a kind of symbolic generalization, and that the goals of research come under the heading of 'values'.

In contrast, *revolutionary science* is the process by which one paradigm is replaced by another. In the history of science, Kuhn sees periods of normal science punctuated by revolutions followed by new periods of normal science. While this broad picture is consistent with the traditional philosophies of science, Kuhn's explanation of how and why the changes come about is quite different.

In the terminology of the previous section, normal science is about the construction of new *models* or the improvement of old ones, whereas paradigm change or revolutionary science is about *theory* change. Kuhn's account of these two different kinds of change is sketched below.

Anomalies are the driving force behind normal science. For Kuhn (1970, p.52) an *anomaly* is a violation of 'the paradigm-induced expectations that govern normal science' In terms of the previous section, an anomaly is one of two things: (a) a *discrepancy* between the best worked-out model of a theory at the time and the known phenomena or (b) a *discrepancy* between two mod-

els, each of which is accepted as a good representation of different parts of the phenomena. The ‘puzzles’ of normal science are puzzles about how to change the ‘paradigm-induced’ expectations so that an anomaly is removed. Scientists solve these puzzles by constructing new models that removes the anomaly without creating too many new ones. Model construction is the engine of normal science, while anomalies provide the fuel. Kuhn does not use the term ‘model’ in this sense, but I believe that it does fit well with how contemporary scientists describe the activities of normal science.

Kuhn’s account of model construction departs from the deductive account described in the previous section, which is more familiar to philosophers. He tends to downplay the role of the formal derivation of models from a background theory, and, instead, suggests that models are constructed by analogy from *exemplars* of the science taught in textbooks. By an ‘exemplar’ Kuhn (1970, p.187) refers to ‘the concrete problem-solutions that students encounter from the start of their scientific education, whether in laboratories, in examinations, or at the ends of chapters in science texts.’ ‘All physicists, for example, begin by learning the same exemplars: problems such as the inclined plane, the conical pendulum, and Keplerian orbits; instruments such as the vernier, the calorimeter, and the Wheatstone bridge.’ Exemplars provide the scientist with a kind of *tacit* knowledge that cannot be articulated explicitly, but is nevertheless an essential part of the paradigm. The advantage of ‘substituting paradigms for rules’ is that it ‘should make the diversity of scientific fields and specialties easier to understand.’ (Kuhn, 1970, p.48) Kuhn therefore rejects the deductivist view that models are the logical consequences of theory and auxiliary assumptions. If we are concerned about the psychology of scientists, then

Kuhn is right that scientists do not always follow a rigorous pattern of deduction. However, philosophers of science are not concerned with the psychology of models but with their evaluation. The deductive picture may be useful for this purpose, although I will be neutral on this point during the remainder of this essay.

Kuhn recognizes that the removal of anomalies by using ‘fudge factors’, or *ad hoc* gerrymandered changes in auxiliary assumptions, is not an acceptable puzzle solving strategy in science. At the same time, he is denying the existence of rules for the construction of models, so he is not able to say that *ad hoc* models are disallowed because they violate the *rules* for model construction. So, how are they disallowed? One such example that he considers is Ptolemaic astronomy:

Given a particular discrepancy, astronomers were invariably able to eliminate it by making some particular adjustment in Ptolemy’s system of compounded circles. But... astronomy’s complexity was increasing far more rapidly than its accuracy and... a discrepancy corrected in one place was likely to show up in another. (Kuhn, 1970, p.68)

This passage is ambiguous. On the one hand, it could point to a practical difficulty in fitting a particular Ptolemaic model (defined by a specific number of epicycles assigned to each celestial body). The adjustment of radii and periods of motion to remove one discrepancy might fail to provide good fit with other known data. This is a problem concerning synchronic fit with data. On the other hand, he may be referring to fact that after a complex model successfully fit all known data, it was likely to fail in its prediction of new data, and therefore the corrected discrepancy would show up in another place. This concerns a diachronic

notion of fit. The second diachronic concept of fit is the epistemologically important notion, for it is a well known character of complex models that accommodation is easy and prediction is hard. This is therefore the more charitable reading of Kuhn.

This brings us to Kuhn's description of revolutionary science, in which the concept of a 'crisis' plays a key role. A *crisis* in normal science occurs when puzzle-solving breaks down; either because no solutions are found, or because the discrepancy corrected in one place shows up in another. Although crisis is necessary to end a period of normal science; it is not *sufficient*. A second requirement is that there is a competing paradigm that shows greater promise in puzzle-solving potential. 'The decision to reject one paradigm is always simultaneously the decision to accept another, and the judgment leading to that decision involves the comparison of both paradigms with nature *and* with each other.' (Kuhn, 1970, p.77)

At the time of publication, Kuhn introduced the new and controversial idea that scientists do not see anomalies, or even crises, as *testing the paradigm* itself.

Though they may begin to lose faith and then to consider alternatives, they do not renounce the paradigm that has led them into crisis. They do not, that is, treat anomalies as counterinstances, though in the vocabulary of philosophy of science that is what they are. (Kuhn, 1970, p.77)

For Kuhn, '...science students accept theories on the authority of teacher and text, not because of evidence' (Kuhn, 1970, p. 80). All the standard confirmation theories of the time, assumed that scientists are constantly evaluating predictive hypotheses, models, *and theories* against the latest empirical evidence. However, writers like Popper recognized that the mediation of auxiliary assumptions often protected theories from

direct falsification. The issue was whether there were ever occasions when the auxiliary assumptions were sufficiently well tested independently of the theory so that the arrow of modus tollens could be directed at the theory some of the time. If Kuhn is right to claim that no such process actually takes place in normal science, then it is a genuine embarrassment for the Popperian point of view. This is still a controversial issue. However, at best it undermines one particular account of how theories are evaluated. It does not preclude the possibility that theories can be objectively evaluated in a different way.

However, at the level of models, Kuhn (1970, p.80) concedes that 'Normal science does and must continually strive to bring theory and fact into closer agreement, and that activity can easily be seen as testing or as a search for confirmation or falsification.' Scientists may try out a number of solutions to a puzzle, 'rejecting those that fail to yield the desired result' (Kuhn, 1970, p.144). Scientists do, therefore, test their *models*. This is an important difference between theories and models on Kuhn's account. *Models are constantly evaluated in normal science, whereas the theory is only evaluated in times of crisis, and only against a competing theory.*

The lesson is clear: If Kuhn is right, then there is a huge difference between the way scientists evaluate theories and the way they evaluate models. Philosophers of science have paid too little attention to normal science. If one has no clear concept of 'model', then one has no clear conception of normal science. To consider only the conjunction of theory and auxiliary assumptions, *T&A*, will not do, because a 'model' in the proper sense does not imply such a conjunction. Otherwise it would be impossible to derive *true* models from a false theory, or false auxiliary assumptions. The proper concept therefore allows for the

separation of the questions: (A) Is the theory true, and (B) Are the models true? If normal science aims at true models, then it may not matter that the theory is false. It may make sense that the truth of theory is not questioned in normal science, because its falsehood does not preclude the success of normal science.

Nevertheless, there is a need to refine the question. As I shall argue in the following sections, the use of idealizations makes it difficult (though not impossible) to defend the view that normal science aims at true models. It is better to argue that normal science aims at predictively accurate models, and then to ponder how this can lead to truth at a higher level of theorizing. To defend the objective rationality of science, I believe that it is important to decompose the problem in this way.

4 HEMPEL'S CRITICISM OF POPPER AND KUHN

Hempel (1979, pp. 50-51) makes the point that a procedure 'can be called rational or irrational only relative to the goals the procedure is to attain.' He notes that 'Popper, Lakatos, Kuhn, Feyerabend, and others have made diverse pronouncements concerning the rationality or irrationality of science...without...giving a reasonably explicit characterization of their conception of rationality which they have in mind and which they seek to illuminate or to disparage in their methodological investigations.' The point is not that there is one unique sense of scientific rationality, for there are many goals of science (some of which are arguably more essential to science, *qua* science, than others). The point is merely that clarity demands that the goals are made explicit, and that rationality with respect to different goals should be discussed separately, one at a time.

Suppose that we can agree that *one* goal of planetary astronomy, from Ptolemy to

Einstein, was to search for the *true* trajectories of the planets in the future and the past. That seems clear enough. But is it entirely clear? I think that the goal implied by this statement is the *predictive accuracy* (Forster and Sober, 1994) of a predictive hypothesis, rather than its truth. The trouble with 'truth' as a goal is that there is no automatic criterion of partial success. The only obvious criterion for achieving truth is black and white—you either achieve it, in which case you are 100% successful, or you do not, in which case you are entirely unsuccessful. This is not the way we understand the 'search for true trajectories'. Some false trajectories are better than others, and predictive accuracy defines what counts as better. Some false hypotheses are predictively more accurate than other false hypotheses, even though none is truer than any other. This feature of predictive accuracy is good.

Other features of predictive accuracy appear to be bad.¹ Predictive accuracy is defined by first fitting a model to one data set, D_1 , and then considering the accuracy of its predictions in another data set, D_2 . The predictive accuracy is the expected fit with respect to D_2 , or equivalently, the fit with the true hypothesis within the domain of data defined by D_2 . For example, suppose that D_1 is the set of observations of Haley's comet available to a group of scientists at the present time. That set will include observations at an assortment of times over a fixed period of times. Suppose we find the Newtonian hypothesis that best fits this data. There are at least two kinds of predictive questions we may ask: (1) If there other observations of Haley's comet during the past that we have not seen, how well does our hypothesis predict these data. (2) How well will our hypothesis predict future positions of Haley's comet. There are two distinct kinds of predictive accuracies at issue here—the first involves the interpola-

tion of our observations to the past, while the second involves their extrapolation to the future.

So, predictive accuracy is subject to Hempel's warning—one should be precise about the notion of predictive accuracy appealed to. Moreover, the distinction amongst different kinds of accuracy allows philosophers of science to raise an interesting variety of methodological questions. For example, suppose that there is no method that will do the best job at optimizing the accuracy of interpolation and extrapolation at the same time. Then there is no unique answer to questions about objective rationality. The objective answers are conditional in nature: If you are interested in interpolation, then use method 1; and if you are interested in extrapolation, then use method 2.

In Forster (2000) I describe one computer simulation that shows that such trade-offs do exist (see also Busemeyer and Wang, 2000). One of the two methods compared involved synchronic fit with data together with a complexity factor. The standard methods of model selection, including the method of maximum likelihood, AIC, and BIC, are of this kind. They performed reasonably well at interpolation, but performed poorly at extrapolation *even in the limit of infinitely large data sets* (sorry, convergence theorems (Earman, 1992) do not help here because you can't converge on the truth if the true predictive hypothesis is not in any of your models). The second method involved a diachronic method of fit, whereby a model fitted to one data set was tested against new data. Unsurprisingly, perhaps, the past predictive success of models provided a better indicator of future predictive success than the standard methods of model evaluation.

Not only does the standard Bayesian philosophy of science (Earman, 1992) not *answer* these harder methodological ques-

tions, it does not allow for the formulation of the questions so long as it defines the rationality of science in terms of the truth or the probability of truth of hypotheses. Hempel's criticism of Popper, Lakatos, Kuhn, and Feyerabend was that the goals of science are not clearly specified. The problem with Bayesianism, in its standard form, is that it specifies a single goal, and is unable to consider other epistemic achievements (this criticism does not apply to decision-theoretic Bayesianism, but this is not the standard form of Bayesianism).

The point of this section has been to argue that the multi-faceted nature of predictive accuracy is actually one of its biggest advantages. For it provides a fine-grained analysis of the epistemic effectiveness of the many methods actually used in real science.

5 THE PROBLEM OF IDEALIZATION

What follows is an example of the explanatory work done by looking at predictive accuracy (only the kind of predictive accuracy associated with interpolation is considered here). The question is: why should idealizations be used in science even when more 'realistic' models are available? Why should complicated models *not always* supercede simpler ones? Also, why should Newtonian science flourish today even though Newton's *theory* is false? In brief, the explanation is that false theories and false models may sometimes help, rather than hinder, the search for truth at the level of predictive hypotheses. To understand when, and why, this should be the case, we need to examine the relationship between predictive hypotheses and models.

Recall that a model M must make use of background data D in order to make predictions. Suppose that there is a unique predictive hypothesis in M , namely H , that best fits that data D . Then it is this best fitting hypothesis, H , that is used to make predictions, and it is therefore the predictive accuracy of

H that defines the predictive accuracy of the model M at that particular time. The predictive accuracies of the other members of M are irrelevant.

In particular, it is irrelevant that there are some predictive hypotheses in the model that are more predictively accurate than H , and this is the key point. If we denote the most predictively accurate member of M by H^* , then H may not be close to H^* , in which case the predictive accuracy of the model is below its *potential* predictive accuracy.

Potential predictive accuracy is irrelevant if it is not actualized. Think carefully about this last statement. It implies the possibility that a true model (in which H^* is true) may achieve less accurate predictions than a false model. Let CIRCLE and ELLIPSE be competing models of a planet's trajectory, and suppose that the true trajectory is actually an ellipse. Then ELLIPSE is true, and CIRCLE is false. However, the data may be such that best fitting circle may be closer to the true trajectory than the best fitting ellipse. If this happens, it is because the best fitting ellipse is not close to the best circle. There are at least three reasons why this may happen.

One reason is that there are errors of observation in the background data, D . Consider the fact that it is always possible to fit an n -degree polynomial through n data points *exactly*. Thus, the H obtained from M will achieve perfect fit, but is unlikely to have high predictive accuracy. This is similar to the case of Ptolemaic or Copernican astronomy, where a model with a sufficient number of epicycles can fit any finite set of observational data to an arbitrary degree (as proven by Fourier's theorem). In such cases, it is necessary to sacrifice the *potential* predictive accuracy of the model in order to maximize the *actual* predictive accuracy of the model. Note that this makes sense of Kuhn's observation (section 3) that the complexity of Ptolemaic astronomy was

increasing far more rapidly than its accuracy in the sense that a discrepancy corrected in one place was likely to show up in another.

The second reason why H may not be close to H^* has nothing to do with observational error. Suppose that M is not true, and that there are no errors of observation. Then M will never fit the data perfectly. H , by definition, is the member of M that fits D the best, and different D will lead to different H , for no other reason than that they are sampled differently. It is impossible for all of these H 's to be equal to H^* . Therefore, H may be quite different from H^* .

The third reason is the data D may be unrepresentative of the domain over which the predictive accuracy is defined. In that case, even an infinite number of data may fail to pick out an H that is close to H^* . This is the case of 'extrapolation error' discussed in the previous section.

The magnitude of this effect is relatively small when two models have close to the same degree of complexity, such as CIRCLE and ELLIPSE. However, the effect is significant when comparing models of widely different degrees of complexity. This is philosophically important because it means that a theory should not be blamed for the poor predictive performance when the idealizations are removed. Or to put it another way, theories must be compared by the predictive success of their idealizations. Exactly how this can be done has yet to be worked out in detail.

Philosophers of science might take the ubiquitous use of idealizations in science to mean one of two things: (a) Science should avoid idealizations, because the goal of science is to obtain true theories and models. (b) Science should continue using idealizations, in which case truth is not the goal of scientific theorizing. For example, Cartwright (1983) opts for (b) in a book called *How the Laws of Physics Lie*. Our discussion shows that there is a third possibility:

(c) Science should continue using idealizations, because they are necessary in order to optimize the predictive accuracy of our models. The explanation is that science seeks *predictive hypotheses* that are as close to the truth as possible.

This shift of focus from truth *simpliciter* to predictive accuracy has subtle but important consequences for any view of testing or confirmation in science. First, there is a problem for any form of Bayesianism that assumes that theories, models, and predictive hypotheses should be evaluated by their probability of truth. If scientists ought to maximize the *probability of truth* of their models, then why should scientists be so indifferent about the falsity of their models? It is no good appealing to posterior probabilities to get around this objection, for the background data, *D*, will most often confirm that the world is really ‘messier’ and more complicated than the model assumes. There appears to be a conflict between what scientists actually believe about the probability of their models being true, and the decisions that Bayesians recommend in response to those beliefs. Are we to say that scientists are irrational, or should we resolve the conflict by supposing that their goal is something other than truth?

Philosophers of science frequently talk about hypothesis testing and selection in terms of ‘confirmation’, ‘justification’, ‘proof’, ‘warrant’, ‘credence’, ‘support’, ‘verification’, and ‘corroboration.’ All of these terms suggest that hypotheses are evaluated with respect to their truth or falsity. It is time that these terms were replaced by words that do not build in that assumption from the start. For that reason, I prefer to talk about hypothesis *testing*, *evaluation*, *appraisal*, or *assessment*. The main thesis of this section is that the problem of idealization *cannot* be solved by any philosophical theory of confirmation, given

the way that the term ‘confirmation’ is usually understood.

6 KUHNIAN COMMENSURABILITY

Kuhn (1970) and Lakatos (1970) say a lot about the comparison of programs and paradigms *as a whole*, but say little about the pairwise comparison of *models* in different programs or paradigms. Why is this? Einstein’s solution to the precession of the perihelion of Mercury was surely evaluated against the attempted Newtonian solutions. Planck’s model of black body radiation (which introduced the quantum hypothesis for the first time) was surely evaluated against the best classical solutions of the day. These are examples of inter-theory model comparison.

Perhaps the obstacle is Kuhn’s famous incommensurability thesis (IT), which says, roughly, that there is a failure of translatability between paradigms—that is, the puzzle solutions of one paradigm cannot be translated and understood in terms of another paradigm. Since I plan to argue that IT is not an obstacle to model comparison, it is appropriate for me to examine IT in more detail.

Kuhn traces his idea back to Butterfield (1962, pp.1-7), who claims that ‘of all the forms of mental activity’ in scientific revolutions, ‘the most difficult to induce...is the art of handling the same data as before,



Figure 1: The duck-rabbit visual gestalt

placing them in a new system of relations with one another by giving them a different framework.’ Kuhn (1970, p. 85) then carries on to remark that ‘Others who have noted this aspect of scientific advance have emphasized its similarity to a change in visual gestalt: the marks on paper that were first seen as a bird are now seen as an antelope, or vice versa.’

I am among those who found the gestalt analogy extremely vague, until I mapped it onto the Butterfield quote. Let me use the duck-rabbit visual gestalt as an example (Figure 1). The marks on the paper represent the data, and they are *intrinsically* the same whether or not the drawing is seen as a duck or as a rabbit. There is no incommensurability at that level. However, the different modes of perception imply a difference in the *significance* or the *salience* of features. For example, the kink at the back of the duck’s head is ‘noise’ when it is seen as a duck, while it is an essential feature when it is seen as the mouth of the rabbit. The same is true about the relationships amongst features. The fact that one ear of the rabbit lies above the other is a matter of accident when it is seen as a rabbit, whereas is the essential that one part of the duck’s bill is above the other when it is seen as a duck.

We see exactly these kinds of changes across scientific revolutions. Copernicus saw great significance in the fact that the retrograde motions of superior planets occurred when and only when those planets were in opposition to the sun. Ptolemaic astronomers did not, even though they had no trouble agreeing that it was a fact. Darwin saw great significance in the structural similarities (homologies) across species, whereas non-evolutionists did not. It is exactly these kinds of differences that Kuhn (1970, pp.118-9) finds in the Aristotelian and Galilean views of pendulum motion:

To the Aristotelians, who believed that a heavy body is moved by its own nature from a higher position to a state of natural rest at a lower one, the swinging body was simply falling with difficulty. Constrained by the chain, it could achieve rest at its low point only after tortuous motion and a considerable time. Galileo, on the other hand, looking at the swinging body, saw a pendulum, a body that almost succeeded in repeating the same motion over and over again ad infinitum. And having seen that much, Galileo observed other properties of the pendulum as well and constructed many of the most significant and original parts of this new dynamics around them. From the properties of the pendulum, for example, Galileo derived his only full and sound arguments for the independence of weight and rate of fall, as well as for the relationship between vertical height and terminal velocity of motions down inclined planes. All of these natural phenomena he saw differently from the way they had been seen before.

The gestalt analogy, and the scientific examples, are consistent with Butterfield’s idea that the new mode of perception handles *the same data as before*. In other words, there may be a failure of translation in some cases, but there is *no incommensurability of the background data, D*. If this is right, then IT claims *only* that the solution to a puzzle in one paradigm cannot be translated into a different paradigm.

Nevertheless, there is an unanswered objection here. If we base inter-theory comparison on predictive accuracy, then aren’t we assuming the existence of a theory-neutral language of observation? And aren’t there strong arguments against the existence

of a theory-neutral observation language, some of which Kuhn himself provided? I am inclined to concede that there is no such thing as an observation language that is neutral with respect to *all* theories, but to deny that this is required for inter-theory model comparisons (Sober, 1990). All we need is a formulation of the problem in terms that are *neutral with respect to the competing theories*. Aristotelians, and Galileans alike, had no trouble understanding what was meant by the ‘number of full swings of the stone in a given period of time’ or whether that number changed when the size of the swings decreased. Einsteinians and Newtonians had no disagreement about how the magnitude of the precession of Mercury’s perihelion should be measured, or about its observed value. And Planck did not introduce a new way of plotting observed light intensities against wavelengths when he introduced his new quantum model of black body radiation. In each case, the prediction *problems* were *easily translated* from one paradigm to the next.

Nor was there any reinterpretation of what counted as successful prediction or ‘good fit’ with the data. All of these examples confirm that the accuracy or inaccuracy of predictions is measured in the common currency of fit, usually defined by standard statistical methods such as the method of least squares.

Kuhn uses IT to argue that scientific knowledge is not *cumulative* despite the fact that the laws or models of earlier theories appear to be derivable as special cases of the later theory. Kepler’s laws appear to be special cases in Newton’s theory, and Newton’s equations appear to be *special* cases of Einstein’s equations. However, for Kuhn, this appearance is illusory because ‘the physical referents of these Einsteinian concepts are by no means identical with those the Newtonian concepts that bear the same name.’ ‘Newtonian mass is conserved;

Einsteinian is convertible with energy. Only at low relative velocities may the two be measured in the same way, and even then they must not be conceived to be the same.’ (Kuhn, 1970, p.102) However, this only serves to reinforce the *previous* interpretation of IT; namely that no Newtonian *model* can be translated into Einsteinian mechanics because they invoke a different set of relations and place them in a new framework.

Moreover, Kuhn (1970, p.102) is clear that the derivations do serve some purpose:

Our argument has, of course, explained why Newton’s Laws ever seemed to work. In doing so it has justified, say, an automobile driver in acting as though he lived in a Newtonian universe. An argument of the same type is used to justify teaching earth-centered astronomy to surveyors.

That is to say, the derivation of ‘limiting’ cases does serve to explain why the older models were so successful in their *predictions*.

Kuhn’s claim that the translatability of models across paradigms is impossible *in principle* is still controversial (Musgrave, 1979). For the purposes of this article, I will treat this issue as unresolved. I argue only that the acceptance of IT does not rule out the possibility of comparing Newtonian and Einsteinian models *with respect to the goal of predictive accuracy*.

So, exactly how is progress with respect to the truth defined across revolutions? My suggestion is that at a given time, the achievement of one program is greater than another if and only if its best worked-out model is predictively more accurate than the best worked-out model of its competitor. While the definition is vague if the domain of prediction is not explicitly specified, there is usually no problem in resolving this ambiguity in real cases. Planck’s formula was accurate over the full range of wavelengths,

whereas its predecessors were only accurate for either the low end of the spectrum or the high end of the spectrum, but not both. Everyone agreed that accuracy over the full spectrum of wavelengths was a goal of the research.

Lakatos (1970) suggested that competing research programs should be compared according to their rate of progressiveness at the time. So, for example, if one is progressive while another is degenerating, then the first receives a better evaluation than the second. I believe that Lakatos's idea is *plainly* wrong. Such an evaluation should compare the *achievements* of one program with the achievements of another. The rate of *improvement* within each program is not relevant to their current state of achievement, though it may be relevant to the question about how the current comparison should be projected into the future. My point is that those two questions should be clearly separated and Lakatos does not.

The best worked-out models of a young research program may not compete well with those of a more established competitor. Its best models have yet to be worked out, so an estimation of the unproven *potential* of the new program is largely an article of faith, similar to religious faith or blind political allegiance. Kuhn (1970, pp.157-58) describes this issue in exactly these terms, and he is right, not because of the incommensurability of competing paradigms, and not because models are incommensurable, but because nobody can predict the future course of science.

This talk of science based on faith brought forth various complaints about Kuhn and the irrationality of science, and from Lakatos (e.g., 1977, p.7) in particular. Of course, Lakatos's charge is unfair because a decision based on uncertainty is not necessarily irrational. However, I believe that Lakatos and Kuhn were talking past each other in any case, just as philosophers

and sociologists of science do today. The rationality of individual scientists, or even of scientific communities as a whole, is a different issue from the one that concerned Lakatos. Lakatos, like many other philosophers of science, was more concerned with whether science made sense as a knowledge-seeking enterprise. In other words (irrespective of what scientists *believe* they are doing) does science achieve knowledge in any sense, and what evidence exists for such a view? With respect to this question, the definition of what it means for one model to be more predictively accurate than another is relevant.

7 THE OBJECTIVITY OF SCIENCE

Kuhn's challenge to the philosophy of science was to defend the rationality of science with respect to the goal of truth. Philosophers have responded to this challenge, and Lakatos's methodology of scientific research programs is one such example. I have tried to argue that there are two obstacles in the way of evaluating this research. One problem is a failure to make a clear distinction between theories, models, and predictive hypotheses (section 2). A second problem is that the goal of scientific research is not always explicit (section 4).

As we have seen, Kuhn (1970) was mainly interested in the social psychology of science, while philosophers of science look to science as an objective source of knowledge. Rationality in this objective sense is not about what scientists believe. The question is not settled by taking a survey of scientists asking 'what is the goal of science?' or 'what are the standards of scientific community?' Nor is it concerned with what scientists think that science *ought* to be. It is about the *achievements* or the potential achievements of science, and the causes responsible for those achievements.

Let me expand upon the notion of causal responsibility. Consider any putative goal

of science, whether it be the truth of theories, the predictive accuracy of models, or the economic prosperity of the United States. Call the goal 'X'. Now consider two, or more, ways or methods of doing science. Call them *A* and *B*. It is now an *objective* question whether *A* is more effective than *B* in achieving *X*. True, it is a vague question until more is said about what 'effective' means. Secondly, the answer may not be univocal; that is, *A* may be more effective in some circumstances, but not in others. Let me refer to such questions as *goal-oriented* questions. These questions have nothing to do with what scientists believe.

Goal-oriented theses are answers to goal-oriented questions: So, 'A is more effective than *B* in achieving *X*' is a typical goal-oriented thesis by my definition. Such theses are *weakly normative* in the sense that they imply 'ought' statements *when coupled with goal statements*. For example, if one could establish that *A* is more effective than *B* in achieving *X*, and *X* is the goal of science, then it would follow that one ought to adopt *A* as the methodology of science.

There is a huge difference between social psychological questions and goal-oriented questions, although the distinction is not always clear. For example, compare the following normative arguments:

- (1) Scientist *Y* believes that method *A* is better than method *B* at achieving *X*.
Scientist *Y* wants *X*. Therefore, scientist *Y* *ought to* use method *A*.
- (2) Method *A* is better than method *B* at achieving *X*.
Scientist *Y* wants *X*. Therefore, scientist *Y* *ought to* use method *A*.

There is an important difference between these arguments. The first provides a *subjective* justification for using method *A*, while the second provides a more *objective* justification for the same action. The goal-oriented claim has nothing to do with the

beliefs of scientists, and supports a more objective rationality claim. The hard problems in the philosophy of science concern the objectivity of science as a goal-oriented process.

For example, in the problem of verisimilitude (Popper, 1963), philosophers of science seek to (a) define the goal of science in terms closeness to the truth, and (b) argue that science has made progress with respect to this goal (for a survey, see Niiniluoto, 1998). It is this kind of objectivity that is often lost in Bayesian decision theory, which currently dominates the philosophy of science in North America. The Bayesian theory is that a decision is rational only if the decision-maker succeeds in maximizing *expected utility*. The issue of whether the maximization of expected utility is causally effective in maximizing utility is the objective side of the problem, and it receives next to no attention. By couching the question of rationality entirely in normative psychological terms Bayesians lose sight of the hard problems in the philosophy of science (*e.g.*, Maher, 1993, especially section 9.4 on verisimilitude).

It is therefore important to me that an ambiguity in my formulation of the problem of idealization is well understood. The question was: Why should scientists use models that they know to be false? There are two different ways of answering this question: one in the style of argument (1) and the other in the style of argument (2). Or more exactly, there are two interpretations of the question. I have attempted to answer the question in the style of argument (2) by arguing for an goal-oriented thesis; namely, that idealized models are effective means to the goal of predictive accuracy. In brief, my claim was that idealized models may often promote the accuracy of predictions because of the way that scientists make predictions from models. Models are first fitted to background data, and this intro-

duces errors that may be far smaller for simpler models, even when the simplicity is obtained at the obvious expense of truth at the level of models. This answer refers to how scientists do science and what is achieved by what they do, and not to what they *believe* they are doing.

This solution to the problem of idealization is still a good one even if it turns out to be psychologically false. For example, scientists might use idealized models because they (truly) believe that they are mathematically more tractable, take less time to apply, and are far less prone to careless computational mistakes. In fact, I would hazard a guess this is right in many instances. However, there is no conflict between this explanation and mine because they address different aspects of science.

In a similar vein, I have tried to retrieve some of the objectivity of science, which Kuhn threw away in the name of incommensurability (section 6). I argued that while the normal science *solutions* to a Kuhnian puzzle may not be translatable across paradigms, there is frequently no real problem in translating the puzzle itself. Moreover, if the difference between good and bad solutions is defined only in terms of the common currency of predictive accuracy, then Kuhn's incommensurability thesis is no obstacle. This sense of progress is weaker than that sought in the verisimilitude program (Niiniluoto, 1998). But since there is no universally accepted definition of verisimilitude, I believe that some objectivity is better than none at all. And nothing I have said rules out the possibility of finding more.

Beginning students in the philosophy of science often enter our subject with a naïve faith in the objectivity of science. Kuhn's *Structure of Scientific Revolutions* challenges their faith, and many of us use it as a classroom text for that reason. However, the truth is always somewhere in between the

two extremes. Kuhn himself tries to restore faith in science by appealing to the standards of the scientific community and the less fickle nature of collective decision-making. However, for Lakatos and many other philosophers of science like myself, the objectivity of science is not rescued by the intersubjective agreement of scientists within a community (and might well be antithetical to it). Rather, the objectivity of science concerns the properties of science as a knowledge-seeking process. Is there progress in science? Is there any sense in which science provides knowledge of the real world using methods that are reliable to some degree in achieving those goals? These are the hard problems in the philosophy of science, and they will never be answered if the philosophy of science is left in the hands of social psychologists.

Acknowledgements

I am grateful to Robert Nola for inviting me to participate in *After Popper, Kuhn and Feyerabend* symposium at the 1997 Australasian Association of Philosophy meetings in Auckland, N. Z. Thanks also go to Ellery Eells, Elliott Sober and an anonymous referee for comments on an earlier draft of this paper.

REFERENCES

- Busemeyer, J. R. and Yi-Min Wang: 2000, 'Model comparisons and model selections based on generalization test methodology', *Journal of Mathematical Psychology*.
- Butterfield, H.: 1962, *The Origins of Modern Science*, The Macmillan Company, New York..
- Cartwright, N.: 1983, *How the Laws of Physics Lie*, Oxford University Press.
- Cohen, R. S., Feyerabend, P. K.. and Wartofsky, M. W. (eds): 1976, *Essays in Memory of Imre Lakatos*, Dordrecht, Holland, D. Reidel.
- DeVito, S.: 1997, 'A Gruesome Problem for the Curve Fitting Solution,' *British Journal for the Philosophy of Science* **48**: 391- 396.
- Earman, J.: 1992, *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*, The MIT Press, Cambridge.
- Forster, M. R.: 1999, 'Model Selection in Science: The Problem of Language Variance', *British Journal for the Philosophy of Science* **50**: 83-102.
- Forster, M. R.: 2000, 'Key Concepts in Model Selection: Performance and Generalizability', *Journal of Mathematical Psychology*.
- Forster, M. R.: forthcoming, 'The New Science of Simplicity' in H. A. Keuzenkamp, M. McAleer, and A. Zellner (eds.), *Simplicity, Inference and Econometric Modelling*, Cambridge University Press.
- Forster, M. R. and Elliott Sober: 1994, 'How to Tell when Simpler, More Unified, or Less *Ad Hoc* Theories will Provide More Accurate Predictions', *British Journal for the Philosophy of Science* **45**, 1 - 35.
- Hempel, C. G.: 1979, 'Scientific Rationality: Analytic vs Pragmatic Perspectives,' in T. H. Geraets (ed), *Rationality Today*, The University of Ottawa Press, Ottawa.
- Kuhn, T.: 1970, *The Structure of Scientific Revolutions* (Second Edition), University of Chicago Press.
- Lakatos, I.: 1970, 'Falsificationism and the Methodology of Scientific Research Programs' in I. Lakatos and A. Musgrave(eds), *Criticism and the Growth of Knowledge*, Cambridge University Press.
- Lakatos, I.: 1977, *Philosophical Papers*, vol. 1, Cambridge University Press.
- Maher, P.: 1993, *Betting on Theories*, Cambridge University Press.
- Miller, D.: 1975, 'The Accuracy of Predictions', *Synthese* **30**: 159-191.
- Musgrave, A.: 1976, 'Method or Madness,' in Cohen, R. S., Feyerabend, P. K.. and Wartofsky, M. W. (eds) *Essays in Memory of Imre Lakatos*, Dordrecht, Holland, D. Reidel.
- Musgrave, A.: 1979, 'How to Avoid Incommensurability' in I. Niiniluoto and R. Tuomela (eds) *The Logic and Epistemology of Scientific Change*, North-Holland Publishing Co., Amsterdam, 336-346.
- Musgrave, A.: 1981, 'Unreal Assumptions in Economic Theory: The F-Twist Untwisted', *KYKLOS* **34**: 377-389.

- Musgrave, A.: 1995, 'Realism and Idealisation: Metaphysical Objections to Scientific Realism', in J. Misiak (ed) *The Problem of Rationality in Science and its Philosophy*, Kluwer, 143-166.
- Niiniluoto, I.: 1998, 'Verisimilitude: The Third Period', *British Journal for the Philosophy of Science* **49**: 1-29.
- Popper, K.: 1959, *The Logic of Scientific Discovery*, Hutchinson, London.
- Popper, Karl (1963): *Conjectures and Refutations*. London: Routledge and Kegan Paul.
- Resnick, R.: 1972, *Basic Concepts in Relativity and Early Quantum Theory*, John Wiley & Sons, New York.
- Rumelhart, D. E., McClelland, J. *et al*: 1986, *Parallel Distributed Processing, Volumes 1 and 2*. MIT Press, Cambridge, Mass.
- Sneed, J. D.: 1971, *The Logical Structure of Mathematical Physics*, D. Reidel, Dordrecht, the Netherlands.
- Sober, E.: 1990, 'Contrastive Empiricism', in W. Savage (ed.), *Scientific Theories*, Minnesota Studies in the Philosophy of Science: vol. 14, Minneapolis: University of Minnesota Press, 392-412.
- Stegmüller, W.: 1979, *The Structuralist View of Theories*, Springer-Verlag, Berlin.

NOTE

¹ The alleged language variance of predictive accuracy (Miller, 1975; DeVito, 1997) is not on this list. For an explanation of why this is so, see Forster (1999).