MIRACULOUS CONSILIENCE? CONSTRAINTS ON FORMULATIONS OF THE NO MIRACLES ARGUMENT

W. John Koolage
Department of History and Philosophy, Eastern Michigan University
2013

Abstract

The intuition that the success of science is a mark of sciences power to track the nature of our world, to permit us confidence that the descriptions of our best sciences are true (or approximately true), is not new. Puntam (1975) popularized an argument based on this intuition; this argument is known as the No Miracles Argument. Providing a convincing formulation of this argument has proven difficult, and recently Magnus and Callender (2004) have suggested that any probabilistic formulation (given form by Bayes theorem) is fallacious. Earlier formulations suffer from problems just as deep; best explanation formulations are question-begging, for example. In this article, I propose two constraints for those attempting to formulate the argument: first, we ought to narrow the scope of the argument, focusing on only a special set of successes, and, second, the epistemic principle that connects the intuition to the realist conclusion ought to be one that is defeasible (rather than probabilistic or categorical). These constraints serve to avoid problems already set out in the literature, and make the No Miracles Argument a sort of local inference – one that is licensed by the case at hand, rather than a general rule. That is, you may only *sometimes* infer to the truth or approximate truth from success.

Keywords: No Miracles Argument, Scientific Realism, Selective Realism, Agreement of Independent Measurements, Base Rate Fallacy

Introduction

If there are any miracles, certain successes of science are not among them. Since Hilary Putnam's work in 1975, arguments that rest on our intuition that the success of science cannot be a miracle have been popular tools for scientific realists. Scientific realism, at its core, hinges on the idea that our best sciences are truth tracking. That is, science, in its best theories, provides us with truthful, nearly truthful, or approximately truthful descriptions of the world. A simplistic way to put Putnam's "no miracles argument" for realism is this: either the best explanation of scientific success is that science tracks the truth or the explanation is that the success is a cosmic coincidence. Of course, either postulating a cosmic coincidence is unwarranted or it is no explanation at all. Thus, we should take it that our best science tracks the truth. If one is struck that this argument is no more than a false dilemma, do not fret. My primary goal in this paper is to refine this argument so that it isn't so obviously troubling (read, transparently question begging).

An interesting new attack on this argument involves the claim that the argument rests on an altogether different fallacy: the base rate fallacy. Magnus and Callender (2004) argue that the no miracles argument rests on this probabilistic fallacy, which occurs when reasoners fail to take into account pertinent statistical information that bears on their assessment of the probability of some proposition. They are correct in their assessment of the no miracles argument as a probabilistic argument, but I argue that all this demonstrates that the argument is not a probabilistic one. This will be my first refinement of the argument.

Further, Magnus and Callender suggest that wholesale arguments for realism, like the no miracles argument are bound for the circular file.[4] Wholesale arguments are attempts to vindicate realism that are not sensitive to particulars of the case; these arguments are supposed to make us realists regarding *all* of (our best) science. They suggest that there may be good realist arguments for the existence of electrons, but not, say, top quarks; they call these arguments, retail arguments. A refinement of the no miracles argument, a selective version, puts it among the retail arguments. This will be my second refinement to the argument.

Finally, the scientific realism that is supposed to be vindicated by the no miracles argument tends to be quite overstated. Not only does the no miracles argument require some refinement, but our basic notion of scientific realism needs some refinement, if not an overhaul. As an extension of the refined no miracles argument I develop here, I will make some simple suggestions regarding a refinement of scientific realism. These suggestions will be in line with a greater project, the development of Consilient Realism.

[4] Gerald Doppelt's *Reconstructing Scientific Realism*, in Philosophy of Science (Vol. 74) provides a novel take on the No Miracles argument. He constructs a wholesale argument by expanding the explanandum of the Inference to the Best Explanation used in the no miracles argument in order to defend scientific realism for our best current theories. While I find this idea oddly compelling, my strategy differs in two substantial ways: (1) rather than expanding the explanandum, I suggest we diminish it, and (2) I suggest we draw back from the standard Inference to the Best Explanation epistemology.

The Base Rate Fallacy and the No Miracles Argument

A statistical version of the No Miracles Argument is implausible according to Magnus and Callender, since it rests on a common statistical fallacy – the Base Rate Fallacy. Consider the following example from Tversky and Kahneman (1982):

> A cab was involved in a hit and run accident at night. Two cab companies, the Green and the Blue, operate in the city. 85% of the cabs in the city are Green and 15% are Blue.
>
> A witness identified the cab as Blue. The court tested the reliability of the witness under the same circumstances that existed on the night of the accident and concluded that the witness correctly identified each one of the two colors 80% of the time and failed 20% of the time.
>
> What is the probability that the cab involved in the accident was Blue rather than Green?

As it turns out, most people fail to assign the correct probability (and they miss by a lot). The diagnosis: most individuals fail to take into account the base rates – the distribution of cabs in the city. The correct probability assignment here is 41%. If you didn't guess right, don't sweat it; there are several reasons that may have happened.

An obvious problem with this experimental set up is found in interpreting the problem. What is the probability the cab involved in the accident was Blue rather than Green? Well, do they mean Blue rather than Green given that the witness said it was Blue? Presumably, yes. But that is far from obvious. Furthermore, Jonathan Koehler argues, persuasively, that base rate neglect (committing the base rate fallacy) is not as pervasive as we have been lead to believe. (1996) The problem is not base rate neglect, but employment of base rates in a non-Bayesian manner. That is, we do make use of base rates, but we do not always follow the Bayesian standard. Further, he argues that there may be other viable (rational) standards for deploying base rates. If he is right, then subjects may not be irrational when they fail to state that the probability is 41%. The point: while I agree with Magnus and Callender that the No Miracles Argument is in trouble when it comes to base rates, I do not agree with them that the intuition that drives the No Miracles Argument is in danger simply because people are notoriously bad at certain probabilistic judgments.

How do we get the "right" answer in the cab problem? Consider the conditional probability in question: Pr (The cab in the accident was Blue | The witness identified the cab as blue). To determine the value of this conditional probability, we will employ Bayes' Theorem:

$$Pr(A|B) = Pr(B|A)Pr(A)/Pr(B)$$

And, the Theorem of Total Probability, which says the following:

Pr(B) = Pr(B|A)Pr(A) + Pr(B|~A)Pr(~A)
So:
Pr(A|B) = Pr(B|A)Pr(A)/[Pr(B|A)Pr(A) + Pr(B|~A)Pr(~A)]

The likelihood, Pr(B|A), is given by the reliability of the witness as described in the story. The witness reports the cab is Blue, when it is Blue, 80% of the time. We also know Pr(A), called a prior probability: 15% of cabs in the city are Blue. Pr(~A) is also well defined in this case; since all the remaining cabs that are not Blue are Green, Pr(~A) is 85%. Finally, the probability the witness says the cab is Blue, when it is Green, is also well defined; it is 20%, as given by the reliability of the witness. So:
Pr(A|B) = (.8)(.15)/[(.8)(.15)+(.2)(.85)] = (.12)/[(.12)+(.17)] = (.12)/(.29) = .41

Here, though subjects often fail to do so, we can determine the "correct" probability that the cab was Blue, given the witness' report, since we have good access to the base rates and we can employ them in Bayes' Theorem.

This is not the case in the No Miracles Argument. Magnus and Callender (2004) suggest that a formal version of the No Miracles Argument looks like this:

> For any theory x, let Sx stand for the expression 'x is successful' and let Tx stand for the expression 'x is true.' Let ~A be the negation of A and let Pr(A|B) be the probability of A conditional on B. We may now gloss the argument in this way for some current theory: [1] The theory h is very likely successful. [2] If h were true, it would be very likely to be successful. [3] If h were false, it would not be likely to be successful. [4] Therefore, there is a high probability that h is true. Formalizing this version of the argument yields:
> Pr (Sh) >> 0 (1)
> Pr (Sh|Th) >> 0 (2)
> Pr (Sh|~Th) << 1 (3)
> Therefore,
> Pr (Th|Sh) >> 0 (4)

Reintroducing our Bayesian formalism and replacing A with Th and B with Sh, we get the following:

Pr(Sh|Th) = Pr(Th|Sh)Pr(Sh)/[Pr(Th|Sh)Pr(Sh) + Pr(Th|~Sh)Pr(~Sh)]

Recall that Pr(A) and Pr(~A) were set by the distribution of cabs in the population; in that case, it was easy to see how this calculation could be performed. In the No Miracles version of this calculation, it is not see easy to see how we get a value for Pr(Sh), and in turn Pr(~Sh). That is, to avoid a base rate fallacy, we must correctly determine this probability. What is the probability some theory h is true? Well, that will depend on the distribution of successful theories in some population of theories.

Before we continue with our exploration of the No Miracles argument as a base rate fallacy, let me take a short detour. It strikes me that the problem isn't a base rate problem, but actually a much older problem – the problem of priors. It is often caviled that the assignment of a probability that some theory is true, say to something like Newton's theory, makes little sense. Considering the probability that it is true that some theory is successful induces a similar vertigo. Thus, base rate fallacy or problem of priors, the No Miracles argument formulated in the way Magnus and Callender suggest is in grave danger. With respect to realism and anti-realism disputes, Jon Dorling (1992) argues that the problem of priors can be resolved (and the base rate problem dissolved) if we take a subjective Bayesian approach to the problem. That is, if the priors, Pr (A) and Pr (~A), are assigned in line with a subjects own degree of belief, then there is no need for an objective estimate of the Pr (Sh) – you merely draw it from your standing psychological state. While this resolves the problem, it does so in a somewhat unsatisfying manner.[5]

Returning, then, to the question of how to get the base rates (determine the distribution of successful theories in some population of theories), Magnus and Callender make several suggestions and find them wanting. Obviously, the h in the No Miracles argument is drawn from a population of successful mature theories. At first, this additional assumption seems to be a good way to resolve the dispute in favor of Pr(Sh) >> 0. Magnus and Callendar make much hay of this, but it seems to me the right thing to say about this fact is this: there is an equivocation over "success" that matter here. Either (1) success means the same in 'h is successful' and 'h is drawn from a pool of successful theories', and we have Pr (Sh) = 1, or (2) success comes in degrees and h's being drawn from a "successful" pool of theories does nothing to tell us just how high Pr(Sh) should be set, and we have still have a standing problem of priors (or base rate determination). In case (1), the posterior Pr (Th|Sh) = 1. This means that every theory is true, given that it is successful; this cannot be the case. There are obvious counter examples to be found in the history of science. In case (2), the problem for the no miracles argument is simply not resolved; we still have no objective way of assigning P(Sh) >> 0.

[5] Sober (1990) and Hitchcock and Sober (2004) provide an interesting discussion of a likelihood version of the No Miracles Argument. The later paper provides an additional worry for the No Miracles Argument: the kind of success we expect of our models involves selection in line with avoiding overfit to the data. Thus, the best explanation as to the predictive success of a particular currently employed model is that it avoids overfit, not that it is true.

As it stands, I am aware of no way to resolve this problem for the No Miracles Argument, short of conceding to Subjective Bayesianism, so I suggest the following: Modus Tollens – the No Miracles Argument is not a probabilistic argument after all.

Abduction and the No Miracles Argument

Another epistemology that has been traditionally employed in the realist literature is Inference to the Best Explanation. Consider André Kukla's formulation (1998):

[1] The enterprise of science is (enormously) more successful than can be accounted for by chance.

[2] The only (or best) explanation for this success is the truth (or approximate truth) of scientific theories.

Therefore,

[3] We should be scientific realists.

First, let me take issue with the conclusion. It is odd to state the conclusion of the abduction as a behavioral (or psychological) imperative. For now, let us simply revise the conclusion as follows: the enterprise of science is truth tracking. We will consider a further revision of this later, but for now this should allay some concerns.

Each of the two premises of the argument are in need of revision, so let us consider each of these in turn.

Success

I suggest, in a manner contrary to Doppelt (2007) and in line with Psillos (1999), that we refine the explanandum. The explanandum for our Inference to the Best Explanation is represented in [1] by the idea that the enterprise of science is (enormously) successful. As Psillos notes, predictive success is just too easy to get. This suggests the following strategy: we should restrict the explanandum, requiring that truth tracking be the best explanation for only some predictive successes, but not all the predictive successes of our sciences. In this way, we will also move from a wholesale argument for realism to a retail one – we are realists only about some of the descriptions of our best current theories. We are going to be *selective* realists; an accurate account of the epistemology of science only warrants belief that the descriptions of science are true in select cases.

Consider the beam balance example, introduced to the literature by Malcolm Forster (2000): Place an object a a distance x_1 from the fulcrum. Then move an object b left and right on the beam on the other side of fulcrum from a until the beam balances. Measure the distance y_1 that b is from the fulcrum when the beam is balanced. Repeat this for different x_1 values (distances of a from the fulcrum) and record all the x_1 and y_1 pairs.

This is the data from Experiment 1. In Experiment 2, we do the same thing for objects *b* and *c*, and in Experiment 3, the same for objects *a* and *c*.

The series of experiments might be described by a single model, call it DUM (for DisUnified Model), which is comprised of three equations:

DUM: (1) $y_1 = \beta_1 x_1$ (2) $y_2 = \beta_2 x_2$, and (3) $y_3 = \beta_3 x_3$.

For each experiment there are two variables, x and y, which are measureable distances, and one adjustable parameter, β, which is estimated from the data. Once the β for a particular experiment is estimated, it provides a fixed mathematical expression of the relationship between the two objects and the two distances – its mathematical structure, the predictive curve, represents some fact about the experiment in question.

For the sake of illustration, let us say you do the first experiment and get the following results:

x (distance)	y (distance)
1	2
2	4
3	6

From these results, we can estimate the parameter β for the first experiment – of course, the estimate will be 2. Thus, the mathematical equation for the first experiment is $y_1 = 2x_1$. We can continue in a similar manner for the other two experiments.

If, however, we are allowed make use of the resources of a theory, like Newtonian Mechanics, we can generated a model that has additional predictive successes. Mechanics when combined with a number of auxiliary assumptions, including idealizing ones, allows for the deduction of the following model, call it UNI (for Unified Model):

UNI: (1) $y_1 = (m_a/m_b)x_1$, (2) $y_2 = (m_b/m_c)x_2$, and (3) $y_3 = (m_a/m_c)x_3$.

Notice that when we employ the concept of mass, given by the Newtonian theory, we get a model that makes more predictions than its earlier counterpart! UNI provides a relation among its sub-models which DUM does not, namely that $m_a/m_c = (m_a/m_b)(m_b/m_c)$; let's call predictions that can be made on the basis of this constraint emergent predictions – they emerge from the model generated by the introduction of some concepts from a theory. For example, we can predict β_3 in two different ways, (1) by doing the third experiment and (2) by using the previous two experiments and the constraint above to determine β_3. If these two measurements agree, we have an agreement of independent measurements.

Now, we are in a position to refine the first premise of the No Miracles Argument:

[1'] The agreement of independent measurements is a success of science that cannot be accounted for by chance.

This captures the intuition that there is something impressive about UNI that is not as impressive as the success of DUM. Further, anti-realists and realists alike can agree that this is a predictive success of science that is certainly worthy of note in developing an account of theory selection in science.

Explanation

The second premise of the No Miracles Argument hinges on the idea that the best explanation of the success of science is that its descriptions are true (or approximately true). To get from this premise (and our earlier one) to the conclusion, philosophers have suggested the following inference pattern: Inference to the Best Explanation. In essence, we can infer from the "bestness" of an explanation to the truth of that explanation. Let us now turn to the second premise and this inference pattern, which requires the truth of the second premise.

Much to do has been made over the details and inductive strength of Inference to the Best Explanation. A trenchant criticism of Inference to the Best Explanation is that it is question begging. Explanation is merely a pragmatic virtue, not an epistemic one, on the anti-realist account of scientific inference; thus, we cannot infer from the best explanation to the truth of that explanation, or we have already won (or conceded) the debate regarding realism (or anti-realism).[6] That is, the strategy in deploying a No Miracles Argument is to apply an inference commonly used in science (Inference to the Best Explanation) to science itself; since Inference to the Best Explanation is objectionable as an inference pattern properly deployed in science, it is again objectionable at this "meta" level.

In light of this concern, I suggest that we pare down the No Miracles Argument. Rather than employ an inference pattern from the practice of science (like Inference to the Best Explanation), I suggest a more modest inference. Let us employ the following principle: unless there is reason to think that some regularity is the result of cosmic coincidence, take it to have a cause in the world. This principle may account for certain events in the history of science, like the dispute over hidden variables in Quantum Mechanics; until there was good reason (Bell's Theorem and related experiments) to think there were no hidden variables, the research program to vindicate hidden variables seemed perfectly rational. This principle

[6] Kukla (1998), Laudan (1981), and Fine (1984) all make this point.

doesn't suffer from the same problem (at least not to the same extent) as Inference to the Best Explanation, as anti-realists can also accept that there is a cause of the regularities in nature – they merely assert that we are epistemically removed from the nature or proper description these causes.

In light of this weaker principle, we can rewrite [2] as follows: [2'] Unless there is reason to think that the agreement of independent measurements is the result of cosmic coincidence, take it to have a cause in the world. Now, we can get a more modest conclusion: [3'] The agreement of independent measurements has a cause in the world. This has a realist flavor, but it isn't so obviously inflated as to be transparently question begging. Take the No Miracles Argument this way: where there are agreements of independent measurements, and there is no established reason to take this agreement as a cosmic coincidence, science has made a discovery about the world.[7] Of course, this isn't much of a victory for realism just yet, since there are many notable anti-realist positions that will agree with this conclusion.

Realism

So far, our modified no miracles argument is a retail argument for realism. It doesn't suggest that we should be realists writ large, merely that there are places where science tracks the truth (makes discoveries about the world). One might begin to worry that this is not much like traditional realism that holds that the descriptions of the world provided by our best theories are true or approximately true. This idea that there is a discovery made where there is an agreement of independent measurements doesn't sound much like that brand of realism.

Notice that theory plays a role in the generation of UNI. In fact, it is difficult to see how we would generate UNI at all without the resources (the concept of mass in particular) of Newton's theory. Now, many cavalier realists are inclined to take these predictive successes as demonstrating that Newton's theory provides truthful descriptions of the world. We have already seen that this inference is in danger, since a probabilistic account of this inference fails and an abductive account begs the question. Realism needs to be more conservative in at least two ways. (1) Since it is the model that provides for the predictive success of UNI over DUM and not the theory (writ large), a realist should be careful not to endorse the whole of the theory

---

[7] My hunch is that this move is also objectionable on similar grounds to Inference to the Best Explanation. However, this move is valuable in putting form to our intuitions about cases like UNI versus DUM. I argue elsewhere that taking the agreement of independent measurements to be a discovery about the world has additional epistemic virtues – it is employed in the identification of artifacts and it allows for the detection of false predictions in earlier theories. Work on this version of the argument is forthcoming.

in light of this success. And, (2) as of now, our No Miracles Argument only gets you to the fact that the intersection of Newtonian Mechanics and the agreement of independent measurements (the success of UNI) has provided us with a discovery; this says nothing about which scientific descriptions are true. However, it is difficult to see how we could escape endorsing some of Newton's theory as true, given its role in this increased predictive success of UNI.

Recognizing that I already pulled explanation from the No Miracles Argument, at some point explanation has to play a role in scientific realism. Good realists take science to be meeting two goals: (1) increasing predictive success (like the development of models such as UNI), and (2) providing an explanation of these predictive successes. Traditional, cavalier, realists take these to be simply tied together. We have seen that these goals are tied together; the concept of mass, from Newton's theory, plays a salient role in the generation of UNI and the discovery UNI allows. However, meeting the second goal is not as simple as reading the true descriptions (given by a theory) off of the predictive successes of the theory, as the traditional realists seem to think. Here I suggest a different strategy: we should use the theory to read into select predictive successes (agreements of independent measurements). Here, we take the theory to provide true descriptions of the world only where we need it to read in an explanation of the discovery endorsed by (the non-miracle of) the agreement of independent measurements.[8]

This should allay the concern that our No Miracles Argument fails to get any sort of robust realism. Of course, the No Miracles Argument only gets us to the conclusion that we have made a discovery where there is an agreement of independent measurements. Once this discovery is made, the realist then asks, what is it that has been discovered? The anti-realist, on the other hand, is happy with the increase in predictive success, and leaves it at that. A realist, on my view can answer her question by looking to the theory in question to provide an explanation of the predictive success (and in turn the discovery).

In restricting our realism to descriptions directed at the discovery made by models of the theory, we have generated a selective realism. This

[8] In ongoing work, I argue that it isn't even the current theory (the one we used to make the discovery) that is used to read in an explanation of the discovery itself, rather it is a future theory (whether a development of the current one or a revolutionary one) that is used to provide the explanation of the discovery. Consider the way in which Relativity explains the successes (and failures) of Mechanics as an example of what I have in mind. Taking this diachronic view of the second goal of realism (providing explanation of the predictive success of a theory) has additional epistemic virtues – for example, it allows us to identify what we thought were predictive successes that were, in fact, failures.

realism also focuses on the connections among theories, models, and certain kinds of predictive successes (the agreement of independent measurements). These are the two most fundamental building blocks of a conservative realism that I call Consilient Realism.

Conclusion

Wholesale arguments for scientific realism, arguments indented to vindicate the claim that all of our best current theories track the truth, tend to face serious difficulties. One such argument is the No Miracles Argument. Attempted analyses of this argument, as a probabilistic argument and as an Inference to the Best Explanation, run into immediate road blocks. It seems transparent that the reason for this is just that these inference patterns are difficult to vindicate even within the proper pursuit of science; for example, Bayesian inference patterns make good sense in some cases, like ones where the base rates can be objectively determined, but seem to make little sense in other cases. Turning these inference patterns to the whole of science (or even just the successes of science) becomes quite difficult. As a result, philosophers of science have begun to push for retail arguments for realism; that is, some inference patterns in science to get us to truth some of the time, and others do not.

In this paper, we explored a retail version of the No Miracles Argument. This argument, to avoid problems (like question begging!) had to be pared down. We weakened the argument so that it only provides us with a proclivity to say that certain predictive successes (that rely on models derived from theories) are truth related (they underwrite the claim that we have made a discovery). This doesn't seem all that out of line with most anti-realist arguments. The realist then takes science to be directed at a second goal, explanation; it is here that she parts ways with the anti-realists. Recognizing that theory plays a role in discovery (and predictive success), the realist seeks explanations of these discoveries. These explanations are often read in to the discovery from the theory used to generate the agreement of independent measurements, but that need not be the case.

The Consilient Realist, when the view is fully developed, will hang her hat on the notion that realism isn't so naïve as to think that all of our best sciences, in all their glory, track the truth.[9] This has been an embarrassment to realism for long enough. She will be pleased to note that realism is

[9] A recent literature has arisen regarding problems with selective realisms (which is the kind of realism being endorsed in this essay). Work by several philosophers, including, most recently, Elsamahi (2004), Chakravartty (2007), and Harker (2010, 2012) have considered this problem. The success of the sort of No Miracles Argument I advance in this paper will depend on how that dispute turns out, and whether a selective realism based on the ideas generated by Forster's beam balance, can survive the main line critiques.

compatible with anti-realism on one level, but that taking certain predictive successes as discoveries (rather than mere increases in predictive accuracy) plays a role other epistemically promising inferences (read, promissory note here).[10]

References:

Blackburn, S. (2002), "Realism: Deconstructing the Debate" *Ratio* XV 2: 113-133

Chakravartty (2007), *A Metaphysics for Scientific Realism: Knowing the Unobservable*. Cambridge University Press

Doppelt, G. (2007), "Reconstructing Scientific Realism" *Philosophy of Science*, 74: 96-118

Elsamahi, M. (2004) "A Critique of Localized Realism" *Philosophy of Science*, 72: 1350-1369

Fine, A. (1984), "The Natural Ontological Attitude" J. Leplin (ed.), *Scientific Realism.* Berkeley: University of California Press, 83-107.

Forster, M. (2000), "Hard Problems in the Philosophy of Science: Idealisation and Commensurability," in R. Nola and H. Sankey (eds) *After Popper, Kuhn, and Feyerabend*. Kluwer, pp. 231-250

Harker, D. (2010) "Two Arguments for Scientific Realism Unified" *Studies in History and Philosophy of Science*, 41: 192-202

Harker, D. (2012) "How to Split a Theory: Defending Selective Realism and Convergence Without Proximity" *The British Journal for the Philosophy of Science*, 0 (2012): 1-28

Hitchcock, C. & Sober, E. (2004), "Prediction Versus Accommodation and the Risk of Overfitting". *British Journal for the Philosophy of Science.* 55: 1-34

Koehler, J.J. (1996), "The base rate fallacy reconsidered: Descriptive, normative, and methodological challenges". *Behavioral and Brain Sciences* 19 (1): 1-53

Kukla, A. (1998), *Studies in Scientific Realism.* New York: Oxford University Press.

Laudan, L. (1981), "A Confutation of Convergent Realism". *Philosophy of Science.* 48: 19-49

Magnus, P.D. & Callender, C. (2004), *Philosophy of Science.* 71: 320-338

Psillos, S. (1999), *Scientific Realism: How Science Tracks Truth*. London: Routledge.

[10] My own current work focuses on questions surrounding correct inferences and coincidences. I am hoping the sort of reasoning suggested by Forster's Beam Balance example will play a key role in a good inference about non-coincidence. If it does, this would be a victory for the idea that the agreement of independent measurements (noted here as a "predictive success that is a discovery") plays a role in other successful inferences.

Putnam, H. (1975), *Mathematics, matter, and method: Philosophical papers* (vol. 1). Cambridge: Cambridge University Press.
Sober, E. (1990), “Contrastive Empiricism” in C. Wade Savage (ed.), *Scientific Theories.* Minesota Studies in the Philosophy of Science, vol. 14. Minneapolis: University of Minnesota Press, 392-410.
Tversky, A. & Kahneman, D. (1982), Evidential Impact of Base Rates. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under Uncertainty: Heuristics and Biases.* Cambridge: Cambridge University Press.