

HOW TO DEFEND SCIENTIFIC ANTI-REALISM

1. MOTIVATING THE DISCUSSION

I think that many philosophers have, at one time or another, felt some sneaky mistake lurking within *the question* of realism about science: should we believe that our current best scientific theories are approximately true, or not? While this unease has traditionally arisen over the matter of how truth interfaces with accounts of meaning and reference in scientific language, more recent attention has turned to the implicit universal quantifier in the question itself: it demands a single line of reasoning for whether one should think that *all* (or most) of our best scientific theories are true. Both the no-miracle argument and pessimistic metainduction attempt to justify such a conclusion about all or most of our current best scientific theories.

It may be wrong to interpret the question of realism in this way. Magnus and Callender [7] have argued that common arguments about realism are predisposed to committing a tenacious and familiar fallacy, where the only patch to fix the fallacy would push the arguer into question-begging territory against their opponent. An examination of what commonly goes wrong is first necessary.¹

2. THE BASE RATE FALLACY

The fallacy committed by arguments like the no-miracles argument and the pessimistic metainduction is called the base rate fallacy, or base rate neglect. To

¹Magnus and Callender suggest that distinguishing *wholesale* from *retail* arguments may help in accounting for why arguments commit this fallacy [7, p.321]; I will address this distinction in relation to Stanford's *new induction* in what follows, and argue that whether an argument is retail or wholesale is not a deciding factor in whether it commits the fallacy.

illustrate this fallacy, consider the following scenario. Jim goes to his doctor to get tested for a disease D . The doctor uses a test which is 80% reliable in that if someone has the disease it will produce a positive result with 80% percent reliability and if someone does not have the disease it will produce a negative result with 80% percent reliability. The doctor administers the test to Jim and the result is positive. Is it pretty likely that Jim has the disease? One might reason as follows.

1. Jim tested positive.
2. If Jim had the disease and tested positive, that's to be expected.
3. If Jim didn't have the disease and tested positive, that'd be surprising.
4. So, Jim probably has the disease.

In other words:

1. $Pr(Pj) \gg 0$
2. $Pr(Pj|Dj) \gg 0$
3. $Pr(Pj|\neg Dj) \ll 1$
- ∴4. $Pr(Dj|Pj) \gg 0$

This argument is fallacious. In seeing why, imagine that there are no occurrences of the disease in the population of which Jim is a member. Now, knowing the rate of the disease in the relevant population (the base rate), what do we learn about Jim having the disease from a positive test result? Nothing, since by hypothesis Jim does not have the disease; the result was just a false positive. The only way to obtain a non-fallacious estimate of the probability that Jim has the disease requires knowing the rate of the disease in the relevant population (or ranges of rates at different levels of confidence). Let's say the disease is common, and that Jim is a member of a 100 person population, 5 of which have the disease.

All we know is that Jim is from this population and that he tested positive. He could be a member of the positive-testing group that *has* the disease; since the test is 80% reliable and five people have the disease, we would expect this group to only consist of four members. The only other option, given that Jim tests positive, is that he will be among those that tested positive but *do not* have the disease; since there are 95 of these, and a 20% chance of testing error, we would expect 19 false positives from this group. So, one ought to expect exactly 23 members of this population to test positive, of which only four have the disease. The chance that Jim is one of these four is given by $4 \div 23$; if Jim is in a population where 5% of the people have the disease, and an 80% reliable test says he is positive for the disease, then Jim ought to raise his confidence that he has the disease from a 5% chance to a 17% chance. To argue that a reliable test says one has a disease and therefore that one probably has the disease is to commit the base rate fallacy.

Now examine the no-miracles argument: we use our best theories to make incredibly accurate predictions, send rockets to the moon, and cure diseases; if our theories were false, this degree of success would be miraculous, and such success is what we would expect if our theories were (at least approximately) true, so our theories are probably true. Magnus and Callender suggest that this argument fits nicely into the exact same fallacious structure as the argument above, if one replaces “testing positive” with “being a successful theory” and “having the disease” with “being true” [7, p.323]:

1. One of our best scientific theories, *h*, is successful.
2. If *h* was true, success is what one would expect.
3. If *h* was false and was successful, that'd be surprising (miraculous).

4. So, h is probably true.

Translated:

1. $Pr(Sh) \gg 0$
2. $Pr(Sh|Th) \gg 0$
3. $Pr(Sh|\neg Th) \ll 1$
- ∴ 4. $Pr(Th|Sh) \gg 0$

Granting, for the sake of argument, that the sort of success in question is a 99% reliable indicator of truth, does the conclusion follow even on this assumption? Not without knowing the *base rate of theory truth* in the population (in the relevant set of theories, call it \mathcal{H}). For if this rate is *zero* then if a theory h is a member of this set then h would be false with certainty. In order to avoid committing this fallacy, the realist would have to plug in a base rate of truth in the population of current scientific theories. The realist cannot *assume* the base rate that they would like (a high one), because this would beg the question against the anti-realist. The realist could *argue* that the base rate of truth must be high, *but that was supposed to be the point of the no-miracles argument*; if such an additional argument was at the realists' disposal, for establishing the rate of truth among scientific theories, then the no-miracles argument would be redundant. So, the no-miracles argument *requires* a rate of theory truth if it is to establish that the rate of theory truth is high. This means that the argument fails, because it either commits the base rate fallacy or begs the question against the anti-realist. Thus, the no-miracles argument turns on pumping up our intuitions for a high probability of $Pr(Sh|Th \ \& \ h \in \mathcal{H})$ and then tries to cash it out as implying that it is probable that $Pr(Th|Sh \ \& \ h \in \mathcal{H})$ without specifying a rate

of truth in \mathcal{H} ; but, as Magnus and Callender point out, that's fallacious [7, p. 325-6].

The anti-realist, on the other hand, will fight against premise (3) in the above argument: that is, they will seek to drive up the probability of $Pr(\neg Th|Sh)$, that a theory is false given that it is successful. This is typically done by amassing cases of past successful theories now believed to be false for an induction onto present successful theories probably being false as well. However, these theories will be from a set of past theories \mathcal{H}_p . Suppose that there is a perfect rate of falsehood in our past theories. This still doesn't deliver the conclusion that $Pr(\neg Th|Sh)$ is high; really, it could mean one of two things: that $Pr(\neg Th|Sh)$ is low because successful theories generally are false or that $Pr(\neg Th|Sh)$ is low because *past* theories are generally false. The realist will argue that new theories are importantly different than old theories in ways relevant to believing that they are true, and that for any current mature theory h , h is a member of a very different set than \mathcal{H} ("What? h ? In that ol' batch!?"): it's a member of the set of shiny new theories, \mathcal{H}^{new} . Of course, the rate of truth in this batch is anyone's guess, since they are still the current ones. The anti-realist will have to argue that the rate of falsity now is the same as it was in the past, *but that was supposed to be established by the pessimistic meta-induction*. The realist will say that our theories now are relevantly different than past theories, and as a result, that the rate of falsity has been decreasing over time. So, the anti-realist must know the base-rate of truth among current theories in order to get the pessimistic induction to establish a rate of truth in current theories, so the argument fails, for it also either begs the question or commits the base rate fallacy.

3. A NEW ANTI-REALIST ARGUMENT

In his 2006 book *Exceeding Our Grasp*, Kyle Stanford points out similar weaknesses of the common arguments in the realism debate, including the anti-realists' argument from *underdetermination*. The problem of underdetermination is typically thought of as the problem of *empirically equivalent* theories (though this identity is an oversimplification): for any theory, one could specify another which makes all the same predictions and would be equally confirmed by the empirical evidence. But, as Stanford points out, the evidence for thinking that this must be some sort of logical truth about science does so “only by transforming the problem into one venerable philosophical chestnut or another” [10, p.16]. Algorithms which have been proposed by those such as Kukla [3] and van Fraassen [11] do less toward introducing serious rival scientific theories than they do to instantiate worries of global skepticism, like Descartes' evil demon [10, p.12]. Other examples tend to be mere instances of what is known as the “tacking problem” in confirmation theory: for any theory with observable consequences, one can construct another theory that is *the same* as the original but with the addition of a sentence that yields no observable consequences; the two theories will be equally confirmed by the observable evidence. Thus, the Michelson-Morley experiment would confirm the theory of special relativity equally with the theory of special relativity conjoined with the additional hypothesis “there are invisible widgets everywhere that have no observable effects on anything.”

Stanford notes that philosophers and scientists could agree that there have been convincing cases of different rival theories that are thought to be empirically equivalent and thus underdetermined by the evidence (or evidence in the foreseeable future).² But this is insufficient evidence for the sweeping claim that

²Laudan and Leplin [4] convincingly argue that we cannot examine any two rival hypotheses and conclude that they *must* be empirically equivalent, as this would require knowing (1) that

all or most theories face worries of underdetermination. The trouble is, if one cannot straightforwardly show that most theories have empirical equivalents without drifting into pre-existing problems in epistemology and confirmation theory, then it is not clear what the problem of underdetermination is, if it is a separate problem for science [10, p.16].

Stanford, however, has staked out a new anti-realist argument, which is something of a hybrid of (1) an underdetermination worry (which he calls the *problem of unconceived alternatives*) coupled with (2) an improved pessimistic induction over the history of science (which he calls the *new induction*). Stanford holds that upon inspection of the history of *scientists*, a problematic pattern emerges. Past scientists have often engaged in *eliminative inferences* in choosing one theory over others from a small set of plausible alternatives (a process of elimination, so to speak). Roughly, one can think of this sort of inference as an inference to the theory that best explains the data out of the set of theories that scientists “*have managed to come up with so far*” [10, p. 31, original emphasis]. From this notion of eliminative inference on a set of presently conceived alternative theories, Stanford correctly notes that this type of inference is reliable *only if* scientists are confident “in their ability to have exhausted the space of likely or plausible explanations in the first place” [10, p.31]. If it could be shown that scientists are

future observations will not distinguish between them and (2) what theories will eventually be adopted into the body of scientific background knowledge and used as auxiliaries, since new auxiliary theories mean new derivable consequences of a hypothesis which cannot be generated by current reflection. Kukla responds [3], brandishes *gruesome* algorithms showing that it must be the case that there can always be two rival hypotheses that are empirically equivalent – better yet, the van Fraassenian move of making a theory θ' out of all the appropriate empirical consequences of a theory θ plus the claim that θ is false. Summary: Laudan and Leplin complain that logico-semantic skullduggery is afoot and a consequent burden of proof battle ensues over whether such trickery shows that there are in fact empirically equivalent rival theories.

generally bad at conceiving of all the plausible alternative theories *and* that many of the best theories of the past and present were arrived at through eliminative inference, then we may have good reason to doubt that our current best theories of mature science are true.^{3,4}

In chapters 3, 4 and 5 of his book, Stanford provides strong historical evidence, from the works of Darwin, Galton and Weismann on heredity, that scientists have been more than willing to commit to the approximate truth of a theory by means of an eliminative inference, inferences that can *now* be seen to constitute failures to exhaust plausible alternative models of explanation.⁵

For the purposes of this paper, I will want to grant the following two points, the second of which I will gesture at a justification for. The first is that the history of scientists does in fact support a strong pattern of theorists (a) failing to conceive of relevant alternative hypotheses and then (b) choosing theories

³Importantly, the ‘best of the bad lot’ worry about abduction is irrelevant here for Stanford; having even a non-unique best lot is not sufficient, since the mistake he is worried about could apply in this case too. Suppose the ‘lot’ that scientists are eliminating theories from is tied for first with some other unconceived lot. One may be mistaken in *believing* a theory is *true* on the basis of eliminative inference, since they performed poorly at exhausting all the possible theories, leaving out an equally plausible best scientific explanation [10, p.47-8].

⁴Whether a theory that is embraced by scientists in the year 2000 can be expected to have been seen as a plausible or ‘scientifically serious’ alternative to scientists in 1800 is a puzzling question, a problem in which ‘philosophical chestnuts’ quickly surface yet again. I will not entertain this issue here. See Magnus [5] for a fuller discussion.

⁵Stanford could certainly give more examples of such illicit eliminative inferences. In chapter 6 he mentions other such instances, such as the claim by Maxwell in *A Treatise on Electricity and Magnetism* [8] that “*whenever energy is transmitted from one body to another in time, there must be a medium or substance in which the energy exists after it leaves one body and before it reaches the other* (1955[1873 Vol. II 493]).” [10, as quoted and emphasized by Stanford 2006, p.152].

using the subsequently unreliable eliminative method. The second point is that Stanford means to provide a wholesale ‘blanket’ anti-realist argument, a single line of reasoning for thinking that we should not believe that most of our current successful scientific theories are approximately true.

Consider the following passage at the end of Chapter 1, where Stanford states what he hopes to establish in his book:

For if the problem of unconceived alternatives is as pervasive as I suggest and has the implications that I claim, *the natural conclusion to draw will be that the fundamental theories of contemporary science should be regarded, like their historical predecessors, simply as powerful conceptual tools for action and guides to further inquiry rather than accurate descriptions of how things stand in otherwise inaccessible domains of nature* [10, p. 24-5, my emphasis].

Stanford makes clear in the above passage the conclusion he is interested in is a blanket conclusion — one that, when and if the argument is complete and successful, will apply to all of mature science. However, I think that he may not intend that a single line of reasoning actually *establish the truth* of his conclusion. What he may mean is that unconceived alternatives, if they are “pervasive” as he suggests, would *allow us to establish* the blanket conclusion *but in a piecemeal fashion*,⁶ by examining many types of scientists working different domains of inquiry. Granting that Stanford does intend his argument to function as a piecemeal blanket argument for anti-realism, I wish to make another distinction precise. Magnus and Callender use a *retail/wholesale* distinction, where retail arguments for scientific realism are “about specific kinds of things such as neutrinos, for example,” and where wholesale arguments sell a conclusion “about all

⁶Magnus and Callendar [7] credit Arthur Fine [2] with coining the term *piecemeal realism*, i.e. establishing scientific realism through individual cases.

or most of the entities posited in our best scientific theories” [7, p.321]. Stanford’s new induction is a wholesale argument, since theoretical entities are not what is at issue; the content of a theory is quite irrelevant as to whether the new induction has something to say about believing it. The relevant aspect of theories for the new induction is the *process through which theories are chosen*, a positive instance for the induction being a theory which was arrived at using the historically suspicious eliminative method. The relevant factor then is only a ‘property’ of the actual theories as a matter of speech; the property, a theorist eliminatively selecting the theory from an anemic set of plausible alternatives, is a property of *scientists*.

Stanford thinks this relevant property of scientists is widespread. Although he thinks that there are clear cases of theory selection which do not depend on eliminative inference, such as the hypothesis “that a chunk of pure sodium will burst into flame when placed in water,” or “that dinosaurs roamed the earth long ago,” he does suggest that there is a *very* wide range of current theories, ones which were (probably) selected for on the basis of eliminative inferences, that are the target of his new induction [10, p. 33]:

Consider the claims that nothing can travel faster than the speed of light, ...that spiders and human beings share a common ancestor in the distant past, [and] that the deformation of spacetime produced by massive bodies is responsible for their mutual gravitational attraction ...The reasons we can offer for believing them [are] limited to the fact that each of the fundamental hypotheses in question offers the most powerful and convincing systematic account we have for explaining, predicting and intervening with respect to a wide range of empirical phenomena ...*and we can neither offer nor even imagine any alternative hypothesis whose performance in these respects would be equally impressive.* ... [10, p. 34, my emphasis].

But theory selection by elimination under these cognitive conditions is precisely what Stanford argues is unreliable. Stanford's intent is (largely) clear: if the history shows that this kind of eliminative inference is a poor method of theory choice, then if a large range of our best current science were chosen on these grounds, then we should doubt the truth of many of our best scientific theories, *viz.* those selected for on an eliminative basis where scientists could not really have been sure that they had exhausted all other possible explanations.

But what is the set from which Stanford seeks to make an inductive inference? It might be a set of scientists, but I think it is actually a set of theories. Scientists could use eliminative inference very often, but when a theory is actually selected over others to join the ranks of the best theories, eliminative inference need not have been used. And if a theory is especially successful, this might be reason to doubt that it was chosen through unreliable means. It seems to me that the set from which Stanford wants to generalize is the set of theories which were, in their respective timeframes, thought to be among the best theories of mature scientific research. Call this set \mathcal{H}_* . One can see Stanford then as taking samples from this set, ones which he can check if they had an underdetermined rival at the time they were 'added' to \mathcal{H}_* by a scientist relying on the eliminative inference method. If enough of theories in \mathcal{H}_* fit this pattern, then Stanford can generalize: current theories in \mathcal{H}_* are probably underdetermined with respect to unconceived rivals as well.

4. GETTING WHAT THE NEW INDUCTION NEEDS

If what I have attributed to Stanford so far is plausible, then it is an interesting question whether Stanford's new induction, another wholesale argument for anti-realism, can escape the pitfall which claimed the pessimistic meta-induction. The question I now turn to answer is whether the new induction can deliver its

conclusion (that many current best scientific theories are underdetermined with respect to unconceived rivals) *without* requiring a base rate of current underdetermination that it cannot help itself to.

As the proponent of the pessimistic induction pushed for a high probability of a theory being false given that it is successful, i.e. $Pr(\neg Th|Sh)$, Stanford seeks to establish that there is a high probability of any particular theory of mature science (indexed to some time in science) being underdetermined (at that time) with respect to an *unconceived* alternative theory. He argues that an examination of \mathcal{H}_* shows that many theories in \mathcal{H}_* were established by the suspicious eliminative inference method (E). This method is thought to be generally unreliable because it has a history of leading to the selection of a theory that is underdetermined with respect to an unconceived rival (U). Stanford can be seen as trying to establish this: $Pr(Uh|Eh \ \& \ h \in \mathcal{H}_*) \gg 0$, that is, that there is a high probability that a theory is underdetermined with respect to an unconceived rival theory *given that* the theory is a member of \mathcal{H}_* and was arrived at by means of E .

Now, there is no immediate probabilistic conclusion about *all* theories in \mathcal{H}_* , but only those which were also arrived at through the eliminative method. If this is what Stanford is after, then he is left with another job if he indeed wants to reach a general conclusion about current best scientific theories: the rate of E among current best scientific theories. If one wants a probability estimate for whether any random current best scientific theory is underdetermined with respect to an unconceived rival, there are two options. Either argue that the current rate of E should not have changed much over time, i.e. that scientists still use eliminative inference and that their cognitive capacities for exhausting all the possible alternative explanations is roughly the same, or alternatively,

establish the rate of E among current theories by investigating the practices of current theorists.

Now it is time to take stock of the position of the new induction. Stanford, like the proponent of the pessimistic induction (PI), must amass cases of past successful theories which now have a desired anti-realist characteristic; for PI, it was that the old theories turned out to be false, and for the new induction, it is that they are now known to have been underdetermined (at the time scientists subscribed to them) with respect to an unconceived rival theory. Notice that the set of theories that the PI proponent needs to do an induction on is roughly the same set as new induction: theories which were (and are) successful and/or considered (previously or currently) to be among the best of mature science. The problem with PI was that *even if* past theories in \mathcal{H}_* had a perfect rate of falsehood, and thus there would be ‘enough’ instances for the anti-realist to safely generalize, this is insufficient for concluding that current theories in \mathcal{H}_* are false, since the instances could just be evidence that *past* theories are generally false. What was needed in addition was the claim that the rate of truth in current theories is going to be the same as past theories, but of course, that was what PI was supposed to establish. The realist will have plenty to say about why newer theories in \mathcal{H}_* are different in important respects, respects which the realist thinks are relevant to whether or not they are true.

But the new induction gains more traction at this point. The PI, which required that the rate of truth be uniform over past and present successful theories, left the anti-realist begging the question. However, the new induction would only require showing something like the following: that the rate of E in past theories in \mathcal{H}_* is not much greater than the rate of E in current theories in \mathcal{H}_* . Discovering the (approximate) rate of E in past theories in \mathcal{H}_* is utterly achievable by empirical

means (it would be hard work, certainly), but certainly would not push the anti-realist into out-of-bounds territory against the realist. The anti-realist will get to use historical research methods to investigate relatively current theory selections, methods involving the interpretation of written passages, inference to the best explanation, and psychological analysis perhaps. These methods beg no questions against a scientific realist. So, the methods for determining what the new induction requires to work, the rate of E among current theories in \mathcal{H}_* , will not beg questions against the realist, since establishing the rate of E among current scientific theories is simply investigative research.

Furthermore, the criteria for being a positive instance of E is relatively clear and is both retrospectively and currently applicable: a theory was or is being selected because (1) in the opinion of theorists, it performs better than all the other theories scientists can come up with in regards to explanation, prediction and intervention, and (2) theorists cannot at the time imagine what another kind of theory would be that could perform as well in the relevant respects. As mentioned earlier, there are at least two ways to establish a current rate of E : either argue that the current rate of E should not have changed much over time, or establish the rate of E among current theories by investigating current theorists. Neither of these rates must be established by begging any points against the realist. Where the PI leaves the anti-realist right where they started (wondering about the current rate of truth among successful theories) the new induction leads to a promising research program. Establishing a high probability of $Pr(Uh|Eh \ \& \ h \in \mathcal{H}_*)$ requires, *prima facie*, no question begging assumptions. Stanford has already begun driving up the rates of E and U in \mathcal{H}_* .

The realist is then left in a puzzling situation. They have so far not been given an *obvious* place to dig in their heels. If *current* rates of E are relatively high, and

there is a historical pattern of E -theories being underdetermined, then the realist would need to come up with reasons why current successful theories of mature science which are E -theories wouldn't also have *unconceived* rivals, *despite the fact* that they were selected using a method predisposed to leaving the selected theory underdetermined with respect to unconceived rivals.⁷ The prospects of this foothold look much less promising for the realist, and it certainly seems to allow for less leverage than they were given by the PI, where the realist could simply point to the many ways in which current theories are more successful than past theories.

The new induction leaves less options for the realist on another front as well. Typical methods of dispelling anti-realist arguments from underdetermination, as seen in Sklar [9], work by denying that current scientific theories are actually underdetermined with respect to *empirical equivalent* different rival theories. Here's what these moves are like: (1) if two theories have all the same empirical consequences, either one could be true *as a matter of convention* (thus the current scientific theory could be believed to be true *because* it was chosen), or (2) the meaning of a theory is exhausted by its empirical consequences (so the current scientific theory would have no rival empirical equivalent — the 'two' theories would differ only *linguistically*). These rebuttals miss the point against Stanford's new underdetermination argument, for there is no reason to think that an

⁷The question ends up being: is $Pr(Uh|Eh \ \& \ h \in \mathcal{H}_{past}) \approx Pr(Uh|Eh \ \& \ h \in \mathcal{H}_{now})$? The value of the term on the left of the identity can be determined empirically. The value lends itself to an inductive generalization from past theorists to current ones, which leaves the realist having to generate reasons why current theories wouldn't also be currently underdetermined with respect to an unconceived rival theory. Of course, the value of the right hand side of the equation isn't directly answerable empirically (yet — save by induction on past theorists), but that's the anti-realist's point.

unconceived rival to a current best scientific theory will be empirically *equivalent* to the current theory.

I have not suggested that the realist is without *any* options against the new induction. Magnus has suggested that any statistical argument about realism that generalizes on a set of theories (such as \mathcal{H}_*) would be flawed unless it met both of the following conditions: (1) the instances of the induction would need to be representative samples from \mathcal{H}_* , and (2) that “the cases involved form a homogenous reference class,” [6, from personal discussion; quote from P.D.’s blog]. The idea in (1) is that if the instances cited for the induction are ‘lemons’ or from too few sectors of scientific research, then the induction to non-lemons or members of unsampled subsets of is illicit. The idea in (2) is the following. If the conclusion of the new induction, together with an anti-realist research program described above, is supposed to deliver an actual probability of some current scientific theory having an unconceived rival, then if this probability is to be interpreted as a frequency probability rather than as an individual propensity interpretation (as suggested by the inductive argument), then there must be a specified reference class. If the reference classes involved in the cases are heterogeneous, that is, if the probability of a theory being U is not homogenous with respect to it being E , then the realist could always argue that for any given current theory the frequency probability could fail due to any number of other possibly relevant factors; that is, the realist could argue that the probability of some current theory being an E -theory might be relatively high, but that this alone no longer entails that the theory has a high probability of being a U -theory. And the anti-realist has no other means to get the current rate of U in a non-question begging way (apart from employing either demonic minions, logico-semantic trickery or instances of a confirmation theory paradox).

Although having U be homogenous with respect to E could guarantee an undisturbed projection of the induction, it's not as though the realist couldn't in principle find a new theory trait X which could significantly reduce the rate of a theory being U and E . Nothing I have said here suggests that a successful research program of finding more cases for the new induction would be the end of realism, to be sure. But it may require the realist's philosophical methodology to change significantly (and save tables in philosophy departments from further thumping).

What such an trait as X would look like is an interesting question, and one that realists now *must* work to flesh out. I am skeptical of a recent claim by Chakravartty [1] that the new induction poses no new problems over and above those raised by the pessimistic induction. Charkravartty argues that, in some sense, it is "silly" for a realist not to grant the soundness of the pessimistic induction, since "the real question of interest" is whether and how the realist can flesh out exactly what is the "principled continuity across scientific theories over time, [that] allows realists to latch on to certain aspects of theoretical descriptions as likely being approximately true" [1, p.5]. This is, in a way the question of what the X property is. He continues to say that the realist will simply grant the pessimistic induction but deny that it implies that we cannot have knowledge of the *approximate* truth of things through science. Chakravartty further says that any move to skirt the pessimistic induction would do the same for the new induction.

But, by the new induction, although we may know that many current theories are underdetermined with respect to unconceived rivals, *we don't know which*

parts of (or to what minimal or vast *extent*) our current theories will share similarities or great differences with their rivals, since the rivals are unconceived.⁸ Chakravartty is certainly right that “the fact that we have not conceived of theories that we may adopt in future does not *preclude* believing that *some* aspects of current theories pertaining to unobservables are approximately true” [1, p. 5, my emphasis]. But the new problem for the realist, after the new induction, is to specify how we could know what the true parts of an approximately true theory are, or, granting the consequences of the new induction, how we could know that there are any approximately true parts of our theories if we cannot say what they are. How strange this *X*-property must be, if it could allow us to know that parts of a theory are approximately true without knowing what those parts are.

In closing, it seems that Stanford’s new induction is a significant improvement over the pessimistic meta-induction, and ought to take its place as the central anti-realist wholesale blanket argument. Importantly, it avoids a pitfall which has claimed the previous flagship realist and anti-realist wholesale arguments. In doing so, it does not seem like a “novel red herring,” as Chakravartty suggests [1, p. 5], but a new pressure in the debate, and perhaps a turning point in what the relevant arguments will look like.

⁸Stanford says as much as warns the “champions” of underdetermination that “...unlike constructing empirical equivalents, [the new induction] does not allow us to say just *which* actual theories are underdetermined by the evidence, nor anything about what the (unconceived) competitors to present theories look like [10, p. 22].

REFERENCES

- [1] Chakravartty, Anjan. “What You Don’t Know Can’t Hurt You: Realism and the Unconceived”. *Philosophical Studies*, forthcoming.
- [2] Fine, Arthur. “Piecemeal Realism”. *Philosophical Studies*, 61: 79–96.
- [3] Kukla, Andre. “Does Every Theory have Empirically Equivalent Rivals?”. *Erkenntnis*, 44(1996): 137–166.
- [4] Laudan, Larry; Leplin, Jarrett. “Empirical Equivalence and Underdetermination”. *The Journal of Philosophy*, Vol. 88, No. 9 (1991): 449–472.
- [5] Magnus, P.D. “What’s New About The New Induction?”. *Synthese*, 148(2006): 295–301.
- [6] Magnus, P.D. “Working Retail”. *Footnotes on Epicycles ...the Philosophical Footfaraw of P.D. Magnus*, <http://laser.fontmonkey.com/foe/>: Monday, December 3, 2007.
- [7] Magnus, P.D.; Callendar, Craig. “Realist Ennui and the Base Rate Fallacy”. *Philosophy of Science*, 71(2004): 320–38.
- [8] Maxwell, James Clerk. *A Treatise on Electricity and Magnetism Vol. 2 [1873]*. Oxford University Press, 1955.
- [9] Sklar, Lawrence. *Space, Time, and Spacetime*. University of California Press, 1977.
- [10] Stanford, Kyle P. *Exceeding Our Grasp: Science, History, and the Problem of Unconceived Alternatives*. Oxford University Press, 2006.
- [11] van Fraassen, Bas. “Empiricism in the Philosophy of Science”. In B.A. Churchland, P.M.; Hooker, editor, *Images of Science*. University of Chicago Press, 1985.