

Scientific Realism and Empiricism

- Administrative:
 - All papers due December 18th (at the latest).
 - I will be available all this week and all next week ...
- Scientific Realism
 - What scientific realism is *not*.
 - * Metaphysical Realism *vs* Scientific Realism
 - Realism and Empiricism as claims about the aims of science
 - Two Arguments for Realism
 - * Realism as “the best explanation” of empirical success
 - Problems for Realist reconstructions of science
 - * Is Truth an *achievable* goal?
 - * Does truth provide a metric of *partial* success?
- Book recommendation: *Science and Relativism* by Larry Laudan.

Metaphysical Realism (not to be confused with *scientific* realism)

- According to *metaphysical realism*, the world is as it is independently of how humans (or anyone else, for that matter) take it to be. The objects the world contains, together with their properties and the relations they enter into, fix the world’s nature and these objects exist independently of our ability to discover or know them.
- Metaphysical *non*-realism asserts that there is no mind-independent fact of the matter concerning “the way the world is.” There are various forms of metaphysical non-realism or anti-realism (idealism, phenomenalism, conventionalism, constructivism, pragmatism *etc.*).
- This is *not* the kind of realism that is at issue when people talk about scientific realism. Almost everyone in the modern scientific realism debate is a *metaphysical* realist. So, this is not really what’s at issue for modern philosophers of science. So what is scientific realism, then?
- Van Fraassen gives a very nice definition ...

Scientific Realism as a claim about the aims/epistemology of science

- Scientific Realism is a two-part thesis:
 1. Science aims to give us, in its theories, (perhaps, *inter alia*) a literally true story of what the world is like.
 2. Acceptance of a scientific theory involves the belief that it is true.
- It is important to note that (1) and (2) are *not* to be understood as claims about *actual* scientists. Everyone knows that actual scientists may (and do) have all sorts of aims (prestige, money, fame, *etc.*) which have nothing to do with *science* as a rational cognitive enterprise.
- One way to think about (1) is as a hypothesis about the proper *rational reconstruction* of science as a cognitive activity.
- And, one way to think about (2) is as an idealization about “ideally rational” scientists who are in an “ideal” epistemological situation.
- These should become clearer as we go along ...

Scientific Realism and the Aims of Science

- It is helpful to contrast scientific realism with some alternatives.
- Van Fraassen’s *Constructive Empiricism*. According to VFCE, the aim of science is to give us *empirically adequate* theories. For VF empirical adequacy is still defined in terms of *truth*. But, it is not literal (or complete) truth — only truth of the *empirical component* of a theory. He has trouble saying exactly what the “empirical component” is.
- *Instrumentalism*. According to instrumentalism, the aim of science is to give us theories which are *useful* in various ways (*e.g.*, useful for making predictions of certain kinds or for building bridges, *etc.*).
- *Naïve Bayesianism*. According to naïve Bayesianism, science aims to give us theories which are *most probable*, given our total evidence.
- *Verisimilitude approaches*. According to verisimilitude approaches, the aim of science is to give us theories that are *as close to* the truth as possible (given the evidence and cognitive resources we have, *etc.*).

Scientific Realism and the “Acceptance” of Theories

- Van Fraassen’s *Constructive Empiricism*. According to VFCE, acceptance of a scientific theory involves belief that the theory is empirically adequate (*i.e.*, that its “empirical component” is true).
- *Instrumentalism*. According to instrumentalism, acceptance of a scientific theory involves commitment to the theory’s being the most useful (in some sense), as compared with the available alternatives.
- *Naïve Bayesianism*. According to naïve Bayesianism, acceptance of a scientific theory involves judging the theory to be most probable, given the total evidence (relative to the available alternatives).
- *Verisimilitude approaches*. According to verisimilitude approaches, acceptance of a scientific theory involves commitment to the theory’s being closest to the truth (relative to the available alternatives).
- NOTE: belief is an appropriate attitude toward theories which aim to be true. But, what about theories which aim at other things?

Arguments in Favor of Scientific Realism I

- *Cosmic Coincidence*. Smart (1963) and Putnam (1975) argue that, if scientific theories were false, then their empirical success (or other useful properties) would be a “cosmic coincidence”.
- Boyd asks why methods crafted by us and reflecting our interests and limitations lead to instrumentally successful science. Contrasting realism with empiricism and constructivism, he finds that realism is the *best explanation* of the success of our scientific theories.
- That is because, he argues, if we begin with truths the methods we have crafted for science produce even more of the same.
- Since it is only realism that demands the truth of our scientific theories, then realism wins as giving the best explanation for the instrumental success of science. Hence, like a scientific hypothesis, realism is most likely to be true and we should believe in it.
- Is this compelling? All of our theories have proved to be false . . .

Arguments in Favor of Scientific Realism II

- *Leplin’s “Novelty” Argument*. Leplin (1997) argues that realism must be preferred over empiricism as an explanation of the success of science.
- He argues that — according to empiricism — all that matters is conformance to (or goodness of fit with) our total evidence.
- And, he points out that a realist can explain why *novel* predictions (which *go beyond* our existing total evidence) are valuable. But, how can an *empiricist* justify the value of such *novel* predictions?
- This raises a very important question (one that VF struggles with). What are *empiricists* allowed to count as part of the *empirical evidence*? And, precisely how do empiricists understand *prediction*?
- If all empiricists are allowed to count as evidence are *extant data sets* or *observations*, then there really does seem to be a problem here about *novel* predictions. That is predictions about *hypothetical* data . . .
- This is most easily illustrated in terms of curve-fitting examples . . .

What is “Empiricism”? I

- Curve-fitting is a two-stage inference process.
 1. First, we choose a *model* M (*i.e.*, a family of curves). For instance, we may have two competing models: LIN and PAR. Where, LIN is the family of *linear* hypotheses (LIN: $y = ax + b$), and PAR is the family of *parabolic* hypotheses (PAR: $y = ax^2 + bx + c$).
 2. Then, we select a *hypothesis* H (*i.e.*, a single curve) from M , using our *extant* total evidence E . Usually, the hypothesis H is chosen by *maximizing the likelihood* $\Pr(E | H)$, over all hypotheses in M .
- If a dataset E contains $n - 1$ data points, then there will always exist an n -degree polynomial model M_n containing a (maximally likely!) hypothesis H_n such that $\Pr(E | H_n) = 1$ (*i.e.*, H_n *perfectly fits* E).
- However, we rarely choose such complicated models in step (1) of our curve-fitting inferences. Why? After all, n -degree polynomial families always contain hypotheses that *perfectly fit* the data. Isn’t that good?

What is “Empiricism”? II

- This is essentially Leplin’s point. If all empiricists can appeal to is *fit* with *extant* data, then it would seem that empiricists should always favor highly complex empirical models like M_n . But, of course, they do not. How is this possible? This is a crucial question ...
- Statisticians know why we shouldn’t always choose empirical models like M_n (even solely for the purpose of making empirical predictions). The reason M_n is usually not such a good idea is that M_n will tend to *overfit* the data — H_n will not *generalize* to *novel* predictions ...
- There is always some *noise* or *error* in our finite sets of observations. And, highly complex models will yield ML hypotheses which fit the noise or the idiosyncrasies of a particular, finite data set. They will not tend to fit the *trend* or the *signal underlying* the extant data.
- What we need here is to appeal to *more than just the existing dataset*. What statisticians do here is talk about *hypothetical* datasets ...

What is “Empiricism”? III

- Since we’re all *metaphysical* realists, we can all talk about the *true underlying distribution* t which generated our extant data set E .
- If we were to generate *novel* data sets E' using t , then the hypothesis H_n chosen from M_n on the basis of E would tend (on average) to do *poorly* on E' compared to hypotheses selected from *simpler* models M .
- Complex models yield hypotheses which fit the *actual* data *better*, but these hypotheses will tend to fit *novel* data *worse* (than simpler ones).
- Leplin’s challenge is to explain why performing well on *novel* predictions is a virtue. He is right that a naïve empiricist (who can only appeal to *actual* data) will have problems here.
- VF requires that the theory be *true* in its “empirical component.” It’s unclear how this would be applied to *statistical* models like M_n . Moreover, M_n will *always* be *true* with respect to the *actual* data!
- Of course, if $M = t$, then this would explain M ’s success!

What is “Empiricism”? IV

- So, realism certainly does have an advantage in this case, if empiricism is not allowed to appeal to *hypothetical* (*non-actual*) datasets.
- But, is this a fair reconstruction of “empiricism”? After all, our empiricists are *metaphysical* realists. So, why can’t they talk about *novel* or *hypothetical* datasets generated by the true distribution?
- It seems to me that once this move is made, there is room for an “empiricist” explanation of the value of novel predictions.
- Moreover, I think the empiricist can also explain why things like simplicity are important, from a purely empirical point of view!
- But, I think empiricists who take *truth* as the aim of science are going to run into trouble in *statistical* contexts. We need a more quantitative notion — something like “fit” or “likelihood” or “verisimilitude” ...
- I will come around to this point again. Now, back to the debate ...

Challenges to Scientific Realism I

- *Is truth an achievable goal?*
 - Are we *capable* of constructing *literally true* theories of the way the world is? All of our past-best scientific theories have turned out to be false, and there is every reason to think that our current best theories are also false (and will be superseded by another theory).
 - Even if we were *capable*, would we ever be in a position to *know* that we had *succeeded*? Thus, should we think that our currently accepted theories are *literally true*? Won’t future science do *better*?
- *Does truth provide a metric of partial success?* When we say Newton’s theory is “better than” Kepler’s, we do *not* mean that Newton’s is true and Kepler’s is false — they are *both* false! So, how in the world can truth furnish a notion of *partial* success in science?
 - This further indicates that we will need to move to a *quantitative* account in order to undergird intuitions about scientific progress.

Challenges to Scientific Realism II

- Popper, a (sometimes) scientific realist, eventually realized that truth (simpliciter) will not do the work we need in an account of scientific progress. For this reason, he moved to a *verisimilitude* account.
- Verisimilitude is “closeness to the truth.” Popper wanted a measure of “closeness to the truth” based solely on the notion of good-old-fashioned *truth*. Popper’s logical proposal was as follows:
 - (*) p is closer to the truth than q iff (i) p has all the true consequences that q has, (ii) q has all the false consequences that p has, and either (iii) p has some true consequences that q lacks, or (iv) q has some false consequences that p lacks.
- As it turns out, (*) has the (absurd) consequence that *no false theory can be closer to the truth than any other false theory*. This is terrible, since those cases are most interesting and crucial (Kepler/Newton)!
- Many attempts to define “verisimilitude” in terms of truth have been made. But, all have failed (see David Miller’s work on verisimilitude).

How About *Probability of Truth*?

- Why not move from *truth* (or truth-based notions of verisimilitude) to *probability of truth*? Essentially, this is the move Bayesians (and other probabilists) make. There is much to recommend this idea.
- We can certainly say that Newton’s theory was *more probable* (given the total evidence) at the time it superseded Kepler’s theory. Nothing prevents one false theory from being *more probable* than another.
- Moreover, our discussion of curve-fitting can be reconstructed in terms of the *average likelihood* of the competing hypotheses — across hypothetical datasets generated by the truth. Indeed, Roger Rosenkrantz (a Bayesian statistician) has given such a reconstruction.
- Non-Bayesians (including Popper!) argue that Bayesianism cannot explain the preference for LIN over PAR (or M_n), since LIN is a *subset* of PAR. They argue that $\Pr(LIN | E) < \Pr(PAR | E)$ for *any* E , since $LIN \models PAR$. Bayesians do have legitimate ways out of this problem.

Stochastic Verisimilitude Accounts

- If we think of our data as being generated by a true stochastic distribution t , then we can try to come up with *approximating distributions* t' on the basis of our data. We can then define a metric μ of the closeness of our approximating distribution t to the true one t' .
- Then, we can let the goal of inference be to obtain the μ -closest approximation t' to t that we can, given our evidence and resources.
- This kind of stochastic or statistical verisimilitude approach has been championed by Forster and Sober in the last few years.
- It has several advantages over other accounts:
 - μ does not have the problems of traditional veris. measures.
 - μ gives us a quantitative measure which can be used to gauge partial success of science (and to compare false theories).
 - There is a well-developed statistical theory which tells us how to *estimate* μ using actual (finite) datasets (there’s an *epistemology!*).

Stochastic Verisimilitude vs Probability of Truth

- “ p is closer to the truth than q , given our best estimate of μ , based on our total evidence E ” is *not* the same as “ $\Pr(p | E) > \Pr(q | E)$.”
- When Bayesians compare Newton’s theory and Kepler’s, they do so on the basis of *the evidence that was available at the time of Newton*. It’s not so clear how the probabilities compare *given what we know now*.
- But, intuitively, the distance of the Newtonian and Keplerian models of planetary motion from the true distribution *has remained the same*. Newtonian models provide a better approximation than Keplerian models, and Einsteinian models are even better approximations, etc.
- Is Newton’s model *more probable* than Kepler’s (they’re both *almost surely false*)? Or, is Newton’s model *closer to the truth* than Kepler’s?
- There’s still room for Bayesianism here (and for epistemic probabilities). But, the probabilities are now probabilities over *distances from the truth* (not probabilities over *truth simpliciter*).