Some Difficulties for the Problem of Unconceived Alternatives

Samuel Ruhmkorff†‡

P. Kyle Stanford defends the problem of unconceived alternatives, which maintains that scientists are unlikely to conceive of all the scientifically plausible alternatives to the theories they accept. Stanford's argument has been criticized on the grounds that the failure of individual scientists to conceive of relevant alternatives does not entail the failure of science as a corporate body to do so. I consider two replies to this criticism and find both lacking. In the process, I argue that Stanford does not provide evidence that there are likely scientifically plausible unconceived alternatives to scientific theories accepted now and in the future.

1. Introduction. Stanford (2006) presents what he takes to be the most fundamental kind of underdetermination problem that threatens eliminative inference and, by extension, scientific realism: the problem of unconceived alternatives. Forber (2008) and Godfrey-Smith (2008) argue that Stanford's historical evidence does not establish that science fails to conceive of all relevant alternatives. I consider here one of Stanford's replies to this objection as well as a reply he might make. I conclude that although the scientific realist has work to do in response to the possibility of unconceived alternatives, Stanford has not established that science fails to conceive of all relevant alternatives, nor has he provided evidence sufficient for the conclusion that there are likely scientifically plausible unconceived alternatives to scientific theories accepted now and in the future.

2. The Problem of Unconceived Alternatives. According to Stanford's problem of unconceived alternatives, consideration of the history of science demonstrates that, at any given time, there are likely theories that

†To contact the author, please write to: Division of Social Studies, Bard College at Simon's Rock, 84 Alford Road, Great Barrington, MA 01230; e-mail: sruhmkorff@simons-rock.edu.

‡I am grateful for helpful conversations with Bob Snyder and Kyle Stanford.

Philosophy of Science, 78 (December 2011) pp. 875–886. 0031-8248/2011/7805-0014$10.00

(1) are at least as well confirmed by the available evidence as the best-confirmed extant theories, (2) are scientifically serious competitors to these extant theories, and (3) have not been conceived by scientists. These competing theories might come to be conceived and differentiated evidentially from those best-confirmed extant theories at some point in the future; however, at that time there will, in all probability, be new unconceived scientifically serious competitors. In other words, we are in a position of "recurrent, transient" underdetermination (Stanford 2006, 17). Accordingly, we should withhold belief from our best-confirmed scientific theories now and in the future. Since Stanford defines scientific realism as "the position that the central claims of our best scientific theories about how things stand in nature must be at least probably and/or approximately true" (6), he takes the problem of unconceived alternatives to support scientific antirealism, specifically, a form of instrumentalism according to which we should not literally believe all the theories we accept.

According to Stanford, the problem of unconceived alternatives is a problem specifically for eliminative inference. Eliminative inference begins with the consideration of the set of competing extant theories in a given domain. All but one of these are refuted or rendered improbable. We then infer that the remaining contender is true or probable. This form of inference is reliable only when "we can be reasonably sure that we have considered all of the most likely, plausible, or reasonable alternatives before we proceed to eliminate all but one of them" (Stanford 2006, 29). Stanford claims that this form of inference is "often . . . perhaps even typically" used in science (28) and, thus, that concerns about its reliability would have widespread consequences for our stance toward the results of scientific inquiry. Stanford's skeptical challenge to science is targeted at the use of eliminative inference in theoretical contexts in which "we have good reason to doubt that we can exhaust the space of plausible alternative possibilities" (37). Because of the pervasiveness of these contexts in science, he thinks we have reason to doubt "virtually all of those fundamental theories concerning remote domains of nature that lie at the heart of the contemporary scientific conception of the natural world" (37).

The bulk of Stanford's defense of the problem of unconceived alternatives lies in his examination of the historical record. He presents an inductive argument from past instances of scientists who have failed to consider relevant alternatives to the theories they have inferred on the basis of eliminative reasoning to current and future instances of scientists' failing to consider relevant alternatives to the theories they have inferred on the basis of eliminative reasoning; he suggests that the failure to consider relevant alternatives in theoretical contexts might be a general property of human cognition (Stanford 2006, 45). The cases he examines lie in the biological sciences—in the examination of generation and heredity by Darwin, Galton, and Weisman. For example, Darwin's defense of pangenesis failed to consider that there might be a common cause (i.e., genetic material) of parents' and offsprings' similar features (rather than parents' features directly causing offsprings' features); Galton failed to attend both to the hypothesis that the expression of units of hereditary information might depend on their relations to other units (rather than being automatically expressed) and to the hypothesis that hereditary material does not itself develop into the organs and features it encodes; and Weisman did not conceive of a process by which full, identical copies of an organism's hereditary information could exist in differentiated cells. The examples of these three scientists (51–140) plus a gesture at a longer list of similar failures to conceive of scientifically plausible alternatives (19–20) are supposed to establish that current and future scientists are not good at conceiving of all relevant scientifically plausible alternatives in fundamental domains of nature and that there are scientifically plausible unconceived alternatives to the best-confirmed scientific theories in such domains now and in the future.

If realists can show that past theories, though surpassed by theories unconceived at the time, contained important elements retained by current theory, then the existence of plausible unconceived alternatives would be less threatening to current theory, as current theory could be expected to have elements preserved by future theory. Preservative realist objections to the pessimistic induction threaten to establish this very means of escape from the problem of unconceived alternatives. Stanford's most common response to preservative realists is that they win, if anything, Pyrrhic victories. These victories would be Pyrrhic because they would leave us skeptical of essential aspects of science. For example, if realists were able to preserve the reference of central theoretical terms of discarded theories, they would be allowing the possibility that we are deeply confused about the referents of central theoretical terms in currently accepted theory. We are thereby left unable to trust the relevant features of current science. Stanford refers to this as the "trust" argument (2006, 157) and claims that it entails that realists must respond to the problem of unconceived alternatives by giving us principled, "prospectively applicable" criteria (168) by which we can distinguish the trustworthy aspects of current theory from the aspects that will turn out to be false from the standpoint of future theory. So far, according to Stanford, no realist has met this requirement.

There are three additional aspects of Stanford's argument that are particularly important for my discussion. First, Stanford is avowedly not interested in Cartesian or other gruesome skeptical hypotheses that motivate some of the literature on underdetermination. He affirms repeatedly that he intends to claim only that eliminative inference is undermined by *scientifically plausible* unconceived alternatives (SPUAs); he argues that we can know that the examples he gives of unconceived alternatives are plausible because they, in fact, came to be accepted by later scientific communities (Stanford 2006, 20–21).

Second, Stanford criticizes Sklar's assumption without argument that "there are vast numbers of perfectly respectable scientific hypotheses . . . we just haven't yet brought to mind" (Sklar 1981, 18; quoted in Stanford 2006, 18), and he criticizes Shimony (1970) and Earman (1992) for claiming that we can make assumptions that limit the space of SPUAs (Stanford 2006, 40–43). We can, therefore, attribute to him agnosticism about the existence of SPUAs prior to the arguments he presents.

Third, according to Stanford, the mere possibility of a competing theory is not sufficient to generate doubt. In his discussion of underdetermination, he writes: "In the absence of any evidence, why should we either assume that such [scientifically plausible empirical equivalent] alternatives exist or let the bare possibility that they might exist prevent us from believing the best confirmed theories we do have?" (Stanford 2006, 11).

3. Blocking the Inference from Scientists to Science. The problem of unconceived alternatives has been critiqued by Forber (2008) and Godfrey-Smith (2008) on the grounds that it attributes properties of individual scientists to the whole of science. Stanford's historical examples are from 25 years of an early period in the search for the biological bases of heredity. While he occasionally presents evidence that contemporaries of Darwin, Galton, and Weisman exhibited similar failings of imagination (e.g., Stanford 2006, 120), Stanford is, in the main, concerned with the individual failures of these three scientists. While Stanford recognizes that the process of peer review might be able to improve upon the failings of individual scientists (129), the relatively small size of the scientific community at that time was significantly different from the situation of contemporary science (Godfrey-Smith 2008, 143). All the putative failures to conceive of important alternative theories that are mentioned by Stanford were remedied by the scientific community in short order, historically speaking. If we allow more scientists more time for exploration of unconceived alternatives, we will be rewarded with more robust investigations than those depicted by Stanford (Forber 2008, 139; Godfrey-Smith 2008, 143).

This has the potential to block the inductive inference at the heart of the problem of unconceived alternatives. Earlier and later scientists might have the same individual dispositions to overlook SPUAs. However, later scientists have had more time for their increasingly robust scientific traditions to uncover alternatives. Therefore, later scientists cannot be inferred to be replicating the failures of earlier scientists.

This is a powerful threat to the problem of unconceived alternatives.

It does not follow from individual scientists' unreliability in conceiving of SPUAs that science as a whole over time is not adept at generating—and even exhausting—scientifically plausible relevant alternatives. Note that this is true even if the unreliability of individual scientists is very high. Indeed, that we have some positive success rate in conceiving of previously unconceived scientifically plausible alternatives is established by the very history of science Stanford cites in his argument for the problem of unconceived alternatives.

4. Stanford's Response and Three Potential Conclusions. In considering Stanford's response to this objection, it is important to become more precise about the conclusions of the problem of unconceived alternatives. Stanford frequently speaks of scientists' "failures to exhaust" the space of serious possibilities (e.g., 2006, 36). This language is ambiguous, as it has both epistemic and ontological implications. One way that humans could repeatedly fail to exhaust the space of serious possibilities is by not being good at *noticing* SPUAs. Another way is by there *being* a large number of such alternatives. And, of course, both factors could be at work.

I take Stanford's epistemic claim (see, e.g., 2006, 21) to be the following:

> **Unreliable Detectors**. Individual scientists (and humans in general) are not reliable detectors of SPUAs such that when such alternatives exist, a given scientist is unlikely to conceive of them.

For the purposes of this discussion, I assume that this claim is established by Stanford's historical investigation.

Unreliable Detectors is not sufficient to undermine Forber's (2008) and Godfrey-Smith's (2008) objection. Stanford recognizes this in his response to the idea that the scientific community over time can overcome the failing of individual scientists:

> Though we are capable of genuine conceptual improvement on past science and therefore can ourselves enjoy the luxury of conceiving of and considering an ever-larger space of serious theoretical alternatives . . . the space of serious theoretical alternatives would seem to have a vague and indefinite character, with members that are difficult if not impossible to individuate sharply or unequivocally. . . . And if (or wherever) the space of serious theoretical possibilities in which we seek to apply eliminative tools of confirmation appears to be indeterminate and unbounded in this way, it seems that we can have little confidence in the power of our eliminative inferences to arrive at the theoretical truth of the matter regardless of how (finitely) long we allow them to operate. (2006, 133)

We cannot reasonably hope that the scientific community over time can overcome the inability of individual scientists to conceive of all scientifically plausible alternatives, because there are simply too many such alternatives. This language suggests the following conclusion for the problem of unconceived alternatives:

> **Plenitude**. No matter how many previously unconceived scientifically plausible alternatives become conceived, there will remain, in all probability, SPUAs to our best scientific hypotheses about fundamental aspects of nature.

It is unclear whether Stanford adheres to Plenitude. His talk about the space of serious theoretical possibilities being "unbounded" suggests that he does; in addition, he seems to consider the problem of unconceived alternatives to be an ongoing problem for scientific inference, as is implied by his phrase "recurrent, transient underdetermination" (17). However, at the same time, he uses the words "appears" and "seems" in this context.

Stanford clearly makes a weaker claim entailed by Plenitude about the existence of SPUAs (see, e.g., 2006, 254). This claim also serves as a reply to Forber (2008) and Godfrey-Smith (2008):

> **Sufficiency**. There are currently, in all probability, SPUAs to our best scientific hypotheses about fundamental aspects of nature.

Stanford accepts Sufficiency; I will assume that he accepts Plenitude. These claims constitute his reply to Forber's and Godfrey-Smith's objection that the scientific community over time might come to uncover all scientifically plausible alternatives. If Plenitude is true, then in all probability the scientific community will never uncover all scientifically plausible alternatives, and we should always be skeptical about current and future scientific theories concerning fundamental domains of nature due to the likely existence of SPUAs. If Sufficiency is true, there is hope that science will one day uncover all scientifically plausible alternatives, but we should remain skeptical about current theories concerning fundamental domains of nature, for there are likely SPUAs to these theories.

5. The Evidence for Plenitude and Sufficiency. The success of Stanford's reply to Forber and Godfrey-Smith depends on whether Plenitude and Sufficiency are true. Moreover, each of these claims seems to be an important component of Stanford's position in its own right. I will argue that Stanford does not provide evidence sufficient to establish either of these claims.

Unreliable Detectors does not entail Sufficiency or Plenitude. To show that individual scientists have trouble identifying SPUAs does not mean that there are an indefinite number—or even a limited number—of SPUAs in a given domain of inquiry, any more than humans' failure to detect tigers reliably means that there are probably always tigers. Accordingly, I assume that Sufficiency and Plenitude are to be established directly by inductive argument from Stanford's historical investigations. As we have seen, Stanford looks in detail at the examples of Darwin, Galton, and Weisman. Stanford also lists a number of other situations in which he expects that there were similar failures to notice SPUAs (2006, 19–20). In each of these cases, there were SPUAs. Therefore, there are likely to be (1) always SPUAs or, at the least, (2) SPUAs to our current theories.

However, there is a problem with this reasoning. To see this, first note that in each of the cases Stanford examines, the theories in question end up being discarded in favor of alternatives not yet conceived at the time. Prima facie, Stanford's inductive basis supports:

> **The Devil You Don't Know** (**DYDK**). It is much more probable that an accepted scientific theory will be replaced by some theory unconceived at the time of the first theory's initial acceptance than by a conceived and previously discounted theory.

DYDK is intended to apply only to the consideration of fully fledged scientific theories—so, for example, the Greeks' dismissal of Democritus's atomic theory does not count as an instance in which scientists ranked a theory incorrectly. DYDK is one way to encapsulate the idea that scientists are better at theory comparison than they are at theory generation.

But now, notice that those who hold DYDK and Unreliable Detectors will expect that, in all probability, each theory from the history of science that is false from the perspective of current science will have been replaced by an alternative as yet unconceived when that theory was first accepted. Thus, Stanford's inductive basis will be fully expected—not only in the cases he cites but also in every case of theories that are false from the perspective of current science. There is no reason to move beyond DYDK in the direction of Sufficiency or Plenitude.

Another way of putting this point is that Stanford's failure to include scientists' consideration of currently accepted scientific theories in his inductive basis biases his sample. It is clear why Stanford does not include these instances: it would be problematic for both parties. Stanford cannot provide examples of SPUAs to current theories (see Stanford 2006, 18), and the realist cannot take scientists' endorsement of current theories to be conclusive instances in which scientists have eliminated the possibility of SPUAs. But this leaves the realist in the position of not being able to adduce any positive instances on her (or his) behalf, creating an asymmetry in the argument. It is as if we were trying to determine the frequency of tigers without ever being able to state definitely that there are situations with no tigers.

Stanford does not tell us much about the way that he selected his sample.

He states that his intention was to look “for evidence of the historical predicament I have claimed we occupy just where we might expect it to be hardest to find” (Stanford 2006, 51). He chose biology because it seems a less likely location of SPUAs than theoretical physics (51–52). He studied a relatively recent time period in order to deal with theories that had been advanced under theoretical assumptions shared by contemporary biology (e.g., the rejection of vitalism; see 60).

It is possible that Stanford’s selection of his sample was random within these constraints (being in the history of science, being in a prima facie unlikely discipline, and sharing basic theoretical assumptions with contemporary science). Accordingly, Stanford might object that by restricting his consideration to the history of science, he has not thereby restricted his consideration to theories that are not currently accepted by scientists, for there are some theories from the history of science that are still accepted by scientists. If this is true—runs the objection—we learned something by discovering that there were SPUAs in the cases Stanford examines. For it could have been that some of the theories selected from the history of science were true from the perspective of current science. Since it did not happen this way, the evidence adduced by Stanford is informative to some degree vis-à-vis the question of the existence of SPUAs. I will say four things regarding this objection.

First, if this objection is on the right track, then whether there is frequency information about SPUAs in the historical record is related to the frequency of theories that are considered to be false among all theories accepted throughout the history of science—in other words, it seems that the problem of unconceived alternatives might reduce to the pessimistic induction. This might not be that surprising, given that the pessimistic induction and DYDK together entail Sufficiency or Plenitude (depending on the strength of the conclusion of the pessimistic induction). However, this would make the problem of unconceived alternatives less novel.

Second, Stanford’s historical explorations deal only with instances in which the scientists in question were mistaken relative to current theory. A fuller examination would critically consider the extent to which these past scientists got things right by the lights of current theory.

Third, a full analysis of whether it is possible to advance a statistical argument based on the history of science for the likelihood of SPUAs now and in the future would have to consider the proper way to individuate theories, the differences in the number of scientists across time and their ability to communicate with each other, the propensity of various structures of the scientific community (e.g., peer review) to support the uncovering of previously unconceived plausible alternatives, the fact that judgments of scientific plausibility are made relative to background theory, and the stability of currently accepted theories (see Magnus 2006, 297).

Fourth, whatever the results of such an exploration, it should be clear that Stanford's historical inquiries constitute a biased sample, for they are instances of successive theorizing concerning the same phenomena. This guarantees that all but one will be seen as false by contemporary science, reducing the effective sample size to one. Again, assuming DYDK and Unreliable Detectors, nothing uncovered in Stanford's historical inquiries is surprising.

I have argued that Plenitude and Sufficiency are not entailed by Unreliable Detectors and are not supported by Stanford's inductive basis. Some might wonder whether much evidence at all needs to be adduced for Plenitude or Sufficiency. Plenitude has been seen to be obvious by a number of philosophers, and Sufficiency is more plausible than Plenitude. As we have seen, Sklar (1981) has assumed Plenitude. Van Fraassen has written in support of Sufficiency: "I believe, and so do you, that there are many theories, perhaps never yet formulated but in accordance with all evidence so far, which explain at least as well as the best we have now" (1989, 146). Yet at the same time, there are philosophers who do not think these are obvious claims. Earman (1992, 172–73) believes that careful surveys of probability space can create a low probability for the claim that there are SPUAs. Boyd (1991) argues that the plausibility of theories is dependent on their fit with the theoretical virtues and already-accepted theories, which are themselves confirmed by the instrumental success of science, with the result that a limited number of theories are plausible. Whatever the merits of these positions, Stanford commits himself—admirably, on my view—to avoiding assumptions about the existence of SPUAs. Accordingly, he needs to provide evidence that there are likely SPUAs. He has not yet done so.

6. Sufficient for Skepticism? For the remainder of this article, I will assume that the problem of unconceived alternatives provides evidence for Unreliable Detectors and DYDK but does not provide evidence for Sufficiency or Plenitude. In this section, I consider a second response to Forber's and Godfrey-Smith's objection: that Unreliable Detectors alone is sufficient to generate skepticism about scientific theories accepted now and in the future.

There is a version of Stanford's trust argument to be had here. Emphasizing the size of the scientific community and the relevant spans of time needed for attrition to overcome individual unreliability creates doubt, for each individual scientist, that she (or he) is in the right community at the right time to believe her (or his) best theories. Stanford could argue that although scientists over time might be able to uncover all scientifically plausible alternatives, they would not know at what point this had been accomplished and, accordingly, should not trust the results of their eliminative inferences about fundamental domains of nature.

Now, it would not be a good objection to working with a particular investor who manages 999 successful portfolios, and one unsuccessful portfolio, that we cannot trust her (or him)—as, for all we know, our portfolio might be the one ruined, and we have no prospectively applicable criteria to determine which portfolios will be successful and which unsuccessful. The difference between cases in which we should trust and cases in which we should not trust is a matter of degree. Thus, the trust argument is really a more vague way of talking about probabilistic reliability.

The investment manager case is one in which we know that the process is reliable. It has not been established by the realist that science over time reliably uncovers all scientifically plausible alternatives. I have argued that Stanford has not provided evidence that science over time is unreliable at uncovering all scientifically plausible alternatives. How does the trust argument apply in cases in which we do not know whether the process is reliable or unreliable?

Stanford has committed himself to avoiding skepticism based on the possibility of the existence of SPUAs (2006, 11). Thus, he needs to give some information about the reliability of the scientific community's attempts to uncover all scientifically plausible alternatives. However, as we have seen, Unreliable Detectors alone is not enough to support such a claim. If we do not know the reliability of the current community of scientists' attempts to uncover all plausible alternatives, we have no positive reason to distrust the results of this community's eliminative inferences in fundamental domains of nature. Given Stanford's commitment about skepticism, Unreliable Detectors does not establish that we ought to be skeptical about the results of current and future science.

Although these observations, if correct, are a victory for scientific realism, there remains the question of how we ought to believe, given that it has not been established that there are unlikely to be SPUAs to currently accepted theories. This is a weaker threat than the original problem of unconceived alternatives, which holds that there are likely SPUAs to our current best theories and (perhaps) that there will always be such theories. Still, eliminative inference to full belief in a theory seems to require some justification for the claim that there are unlikely to be SPUAs—and the realist has not provided such justification. Stanford might choose to reverse his intention to consider only skepticism generated by evidence of the existence of alternatives and to consider also skepticism based on the absence of reasons for belief in the nonexistence of alternatives. This issue needs to be addressed by scientific realism.

7. Conclusion. I have considered two responses to Forber's and Godfrey-Smith's objection that the community of scientists over time can overcome the individual unreliability of scientists in conceiving of SPUAs. The first response, endorsed by Stanford, is to establish Plenitude or, at the least, Sufficiency. But Stanford does not give evidence for Plenitude or Sufficiency over Scarcity:

> **Scarcity.** In each domain of fundamental scientific theorizing, there are relatively few scientifically plausible alternatives.

If Scarcity, DYDK, and Unreliable Detectors are true, we would expect to find exactly the features of the history of science outlined in Stanford's inductive basis. Therefore, his inductive basis does not support Plenitude or Sufficiency over Scarcity.

Note that if Scarcity is true, then the inductive basis would support the claim that current theories are probably true (or, at least, that recent theories are increasingly more likely to be true). This demonstrates how eliminative inference and the pessimistic induction and its relatives—including the problem of unconceived alternatives—play against each other. Past failures of science lend credence to the idea of the present and future failings of science only if there are lots of scientifically plausible ways to fail. If there are fewer scientifically plausible ways to fail, then past failings of science support the idea that current science is on the right track.

The second response is to marshal a version of the trust argument based on Unreliable Detectors. I have argued that the trust argument reduces to an argument about probabilities and that, in the absence of clear information about the probability that current science is neglecting SPUAs, it does not render a definite conclusion in this case. Nevertheless, the scientific realist should address the epistemic problem that arises from the challenge of providing evidence that it is unlikely that there are SPUAs.

Stanford is correct to focus our attention on the problem of unconceived alternatives, but we have not yet been presented with sufficient evidence for skepticism based on this problem. Future research should consider whether a general statistical argument is possible that would give us information about the probability that there are SPUAs; whether there are reasons apart from the historical record for thinking that in a particular domain of scientific theorizing, there are likely—or likely not—SPUAs; and what attitudes we should have toward our best scientific theories in the face of whatever uncertainty there is about SPUAs.

REFERENCES

Boyd, Richard. 1991. "Observations, Explanatory Power, and Simplicity: Toward a Non-Humean Account." In *The Philosophy of Science*, ed. Richard Boyd, Philip Gasper, and J. D. Trout, 349–77. Cambridge, MA: MIT Press. Originally published in *Obser-*

vation, Experiment, and Hypothesis in Modern Physical Science, ed. Peter Achinstein (Cambridge, MA: MIT Press, 1985).

Earman, John. 1992. *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. Cambridge, MA: MIT Press.

Forber, Patrick. 2008. "Forever beyond Our Grasp?" Review of *Exceeding Our Grasp: Science, History, and the Problem of Unconceived Alternatives*, by P. Kyle Stanford. *Biology and Philosophy* 23:135–41.

Godfrey-Smith, Peter. 2008. "Recurrent Transient Underdetermination and the Glass Half Full." *Philosophical Studies* 137:141–48.

Magnus, P. D. 2006. "What's New about the New Induction?" *Synthese* 148:295–301.

Shimony, Abner. 1970. "Scientific Inference." In *The Nature and Function of Scientific Theories: Essays in Contemporary Science and Philosophy*, ed. Robert G. Colodny, 79–172. Pittsburgh: University of Pittsburgh Press.

Sklar, Lawrence. 1981. "Do Unborn Hypotheses Have Rights?" *Pacific Philosophical Quarterly* 62:17–29.

Stanford, P. Kyle. 2006. *Exceeding Our Grasp: Science, History, and the Problem of Unconceived Alternatives*. Oxford: Oxford University Press.

van Fraassen, Bas C. 1989. *Laws and Symmetry*. Oxford: Oxford University Press.