

# Replies

SHERRILYN ROUSH

*University of California, Berkeley*

## Reply to Goldman

I would like to thank Alvin for a spirited, and gentlemanly, debate we've had on these issues, which is extended further here.

Alvin is exactly right that if we make his assumption about maximum specificity and deduceability (which I have doubts about), then on my view of knowledge Sphere Guy doesn't know there's a sphere in front of him. This may sound silly when we focus on his tactile access to the sphere in the actual world, but if we take a broader view we see that there is more at stake than this.

Contrary to Alvin's impression, methods are not at all excised from my view of knowledge. My theory of how to judge whether someone knows requires us to consider everything (probable) that is and would be responsible for the fact that the person believes or not, whether these occur in his head or in the world, which the formulation in terms of probability helps to make very clear. (See Chapter 3.) Ironically, my refusal to relativize to method has us taking into consideration more facts about the subject's method than Alvin's criteria do, for my view takes into account, as appropriate, what process the person would have used and has a tendency to use, and not just the properties of the one he happened in fact to use.

When the fact that a method was used by a subject in coming to belief in  $p$  is independent of the truth of  $p$ , which is actually most of the time in our lives, the conditions of application of the variation condition insure that we evaluate the subject by considering only what he would do and how he would fare in his beliefs were he to use that method he actually used. So, under that condition, my view agrees with Alvin, and Nozick also. But when whether a subject used that method is not independent of the truth value of  $p$ , then the variation condition in my view says we must consider in addition the subject's resulting beliefs in all probable scenarios where he is such that he might well

have used a different method. That is, we ask about whether his tendencies in method choice are any good.<sup>1</sup> Relativizing a criterion of knowledge to the method the subject actually used requires us simply to ignore, in all cases, the subject's tendencies to use good methods or bad. (I will discuss this issue for the adherence condition below.)

My approach picks up other cases where intuitions are on my side. For example, do we know that we are not brains in vats when if we were brains in vats our process of belief formation would systematically hide that from us? The fact that what method we use is not independent of the truth value of our belief is salient here. Most people think that the fact that being a brain in a vat would hide that from us undermines our knowledge unless we have more to say. On Alvin's view it's not a problem: we know as long as we're actually using good perception. On my view we can know we're not brains in vats, but it's in virtue of tracking the table, which on my scheme turns out to require that our method of coming to a belief about that is highly unlikely to be different in case the table isn't there. This is fulfilled in a normal case. However, if though we are not in fact brains in vats the scenarios in which a person like us is a b-i-v are quite *likely*, then this condition is not fulfilled and we don't even know there is a table in front of us. Though mine is, like Alvin's, an externalist view that does not require internal access to an argument about our situation in order for us to know, the condition I set for whether we know or not is more in keeping with our intuitions. Do we know that there are no nuclear missiles headed for the large American cities if were there to be then it is highly likely that American intelligence would hide this from the public in order to avoid panic? I don't think we would say that we do unless we think that such a missile attack is quite unlikely.

The frequent fabricator is actually different, since which method she uses *is* independent of the truth value of p. Thus my variation condition relativizes to her method. Where she fails is in the adherence condition, since for that the rules of application allow anything that is independent of the truth value of p to vary, to see if the subject is good against their (usual) irrelevance. Here her method choice is evaluated for its general tendencies, and since she fails to have the right belief most of the time, she doesn't know this time.

---

<sup>1</sup> Method is not named in the conditions of application, so there is no problem of specificity of its description. What I say here about method is a consequence of a procedure that tells you for any q, of any level of generality, whether it gets held fixed or not. That procedure constitutes a full solution to the generality problem for tracking in that it gives a determinate answer to every question about level of generality. Whether it is a solution that works is of course less easy to be certain about.

I formulated my example in terms of frequency of fabrication because of Alvin's association of frequency with reliability. (Goldman 1979) But my point doesn't depend on this. If we don't care how it was determined that a subject would use the method she actually used, then we will allow as knowledge cases where we have unreliable usage of a reliable method or process. That makes the process of belief formation broadly unreliable. I think this issue is even more awkward than that for Alvin's view, because we can make the question about which method to use live in the subject's head though not necessarily in her accessible consciousness. Consider a case where the subject's brain determines, badly, whether she uses the fortune-teller or perception method. Alvin cannot deny that that choice itself is a part of her process of coming to believe *p*. If it is, then Alvin must say, like me, that her belief isn't knowledge, now because the process through which it is formed is straightforwardly unreliable. But now Alvin's own theory conflicts with his intuitions.

Alvin says I am entranced with power, and I gather I am meant to be discouraged by this. But if I'm entranced then so is the entire human species. Among his unfortunate properties Sphere Guy also has a serious survival disadvantage over his tracking counterpart: if a tiger were to come up and steal his sphere then he would likely still believe the sphere was there because he wouldn't be using his hands anymore but his hallucination-inducing eyes. The tracker, whose method doesn't go kerflooy when the sphere goes, would pick up on the missing sphere, and if he also had reactions appropriate to the fact that tigers liked to steal spheres, he'd be more likely to take the needed measures to survive. If Sphere Guy has those tiger reactions, they wouldn't help him because he would have no belief triggering them. Organisms like Sphere Guy will tend not to survive as long as their counterparts, so if there are many actual cases where knowledge promotes survival then my view is better for explaining them.

The disagreements Alvin and I have about the power-promoting property of knowledge seem to revolve mainly around how we understand the intuitive power side of the relationship. Neither of our views are developed in much detail, and I think this is a topic that deserves more attention and discussion. That said, my starting point is, of course, much better than his. Alvin says that knowing *p* at some later time does not give the subject power at this time. I agree and haven't proposed otherwise. The question is what it is about having the property of knowledge now that gives one power now.

My view is that having power at a given time requires having a property that gives the potential for future successful action. For example, power of certain kinds possessed now pose a threat *now* (typically to

others) without actually being used now. And we can say of someone who had cutting edge weapons but lost the war through incompetence that he had power but didn't use it well. That is, he had a lot of potential to win the war at the time when he started it, but as a matter of fact he didn't use it properly. If he had won the war he would have exhibited the power that he had, that is, the potential that existed from the beginning. Having power doesn't determine what you will do in the future, but power at a given time enhances your potential for future action.

The having of a functioning memory now will of course also be part of what supports your potential now for acting in the future, but this is all included in tracking the world now. If you don't have a functioning memory now you are not now prepared in your belief dispositions for all those possible scenarios in which you would have to decide whether to change or stay with your beliefs. By contrast, having as a matter of fact formed the belief that you have now through a reliable process doesn't tell us what you now have a tendency to do with your beliefs in scenarios of a sort that you are not now actually in. One reason is that for knowledge now Alvin doesn't require that the subject have a tendency to use a good method of belief formation, but only that he actually formed the current belief using a good method. So, I think that Alvin's kind of knowledge does not pose a threat to mine.

Concerning probability, my claim, in context, was that the frequency interpretation is no more controversial than the propensity or other objective interpretations. All interpretations of the axioms have their advantages and drawbacks. I tend to write using objective language for ease of understanding. As I said in the book, the definition of knowledge I am pushing is independent of the interpretation of probability; that is, it can be of service using any interpretation you like. You do get a different tracking theory depending on which interpretation you use. On an objective interpretation, of course, you get a theory of the conditions under which a subject knows. If you apply the theory using a subjective interpretation then yours will be a theory of which propositions the evaluator's belief in which should determine his belief about whether the subject in question knows. My view is that either approach should, or at least is able, to incorporate tracking. Incidentally, the interpretation of probability has nothing to do with externalism or internalism, which is about what is the subject matter of the propositions in the language that are taken to determine whether a person has knowledge. Are propositions about the person's relation to the external world relevant to whether she knows? The question of which interpretation of probability to use is instead about whether the theory takes those claims about the external world to be relevant as claims of fact or as beliefs of the evaluator that he uses to determine what else he

should believe. An internalist is welcome to say that he has a theory of when a person in fact knows, not merely whether he should believe the person knows, and will also say that the facts relevant to that question of fact are those about that person's internal states. On the other side, even using a subjective interpretation of probability the tracking view is externalist because it says that the propositions of the language our confidence in which determines how confident we should be that S knows p involve content about S's relation to the world.

The rules of application for the tracking conditions make them more determinate, yet still independent of the interpretation of probability. Though their formulation must be complicated, they express two simple ideas. 1) We determine which propositions at which level of generality are to be used in the evaluation of the tracking conditions by starting with the p of "S knows p," which itself has a particular level of generality. For any q of the language, of any level of generality, we plug q into a rule, and out pops an answer about the role it will play—fixed or varying—in each of the tracking conditions. This is a solution to the generality problem for tracking, but it is a strategy that is not available to the process reliabilist since his criterion begins with the process rather than the particular belief. 2) The second simple idea is that those propositions whose probabilities should be allowed to vary when we ask what the subject's beliefs do over a variety of situations, as we do in both tracking conditions, are the ones that are more "influenced" by the truth-value of p than the truth value of p is "influenced" by their truth values. The subject should be able to "catch" the possible and likely indications that p is true and ignore others that are irrelevant, and should be able to catch the indicators that p is false and ignore the others. The possible indicators are simply the matters "influenced" by the truth value of p. (Note: I do not claim that the concept of influence can be cashed out probabilistically. Influence is merely an intuitive idea here.)

Finally, all of these rules are relative to not only an interpretation of probability in general—choose the one you like—but also to a particular probability model, that is, a language, and a function from every proposition in that language to a value that is its probability. I call this function  $P_u$ , the  $U_r$ -probability function, and it is given by one who applies this theory, not by me. This is innocuous; it corresponds to the common practice and assumption that we can expect no answer to the question whether S knows p until someone supplies a description of S, the world, and what S's actions, tendencies, history, and circumstances are in the world, and it is not the theorist's job to do that but the job of someone who applies or tests the theory. The function  $P_u$  simply incorporates the evaluator's complete description of the world and the subject, so that these assumptions are available as needed.

## Response to Godfrey-Smith

Thanks to Peter for his remarks. It is always a pleasure.

Peter helpfully identifies three kinds of discussion of knowledge, all of which I approve of, 1) investigation of ordinary peoples' use of ideas and words like "knowledge," 2) how humans and other animals are actually connected to the world, and 3) how those concepts and our dealings with the world do or do not relate. And I agree that finding harmonies tends to show us something useful, and that this is one of the things I've been up to. But I also look for stipulative definitions that carve into the blurriness we get at the level of ordinary talk, which means I'm willing to be more prescriptive if I find that a crisp concept identifies a phenomenon that we are interested in. In other words, I'm willing to be more dictatorial than Peter may be comfortable with.

Peter is skeptical that the adherence condition is reflected in the ordinary concept of knowledge. I would still defend my account of one route to knowledge that Betsy has in the face of Jesse James, since my view of knowledge of implication is different from most in that I do *not* require the subject to even be in any position to make an argument. Although there are plenty of cases where one needs to have an argument to get to knowledge because the implication is complex, that is a contingent matter concerning the means to knowledge, not a matter of the definition of the goal; in simple cases one who knows has the appropriate responsiveness among her beliefs (described above) automatically. If Betsy doesn't have these dispositions between all she knows about the legendary Jesse James and the identity of the face that has been presented to her, then she is more impaired than the example assumes.

However, after the book was published I discovered that on my view those subjects who use a method accidentally also have a more direct way of knowing. For in cases of accidental use of a method, the scenarios where the accident didn't happen are ones where the subject doesn't form a belief at all. Such scenarios are knocked out of consideration in the first domain-limitation in the rule of application for the adherence condition. The adherence conditions thus ask: If she forms a belief at all, does she form the right one? Accidentality of method choice is not the problem for me that it was for Nozick.

Peter is right to point out that the president I described is faulty in too many ways to be probative here. Consider instead the Underconfident Student. In a particular subject matter she's well aware, perhaps too aware, of everything she's ever been taught about the ways that a claim in that area can be wrong, so, let's say, she would never believe or assert a *p* in that subject matter if it were false. However, partly because of this awareness, she often has trouble bringing herself to

confident beliefs. But what if we find her on a day of high spirits; she's confident in a seminar, where she makes a presentation arguing very coherently for  $p$ , where  $p$  is true. But suppose that an aggressive fellow student were to hold forth insistently against  $p$  or her arguments for it, by talking about a matter that used similar words but was only tangentially related to  $p$  or the argument. In the case of this bluster, which is all too realistic, the person we are imagining could easily lose faith. Maybe we should count someone with this pattern of behavior as knowing in those rare cases where she in fact confidently and legitimately asserts  $p$ , but statistics and experience tell us that we actually usually don't. We expect a little rigidity—even in the face of bluster—that is based on a little confidence that such remarks may not be relevant. We expect some degree of adherence.

The Jamesian connection, the Classical Statistics connection, and the signal detection connection to the tracking conditions are all highly appropriate and welcome. The trade-offs between the two conditions are especially worthy of further discussion. Though the contingent conditions organisms find themselves in clearly have a trade-off structure due to finite resources, I have wondered whether the trade-off is due only to empirical facts or also to something about the abstract relation between the tracking conditions. And it turns out that the tracking conditions do impose a trade-off in so far as they place requirements on what you do when  $p$  is true and when  $p$  is false, and those happen to be mutually exclusive and exhaustive of the ways the world might be with respect to  $p$ . So, in principle there's no avoiding a trade-off between them. However, in practice the more resources we have, the more separating indicators we might find, which would increase the distance between the distributions, and which perfectly matches my view of what evidence is. (Chapter 5)

As Peter points out, our choices when confronted with the curves are to use the indicators we have or expend resources to find more discriminating ones. One instance of this is the choices governments make as to whether to fund more fundamental science or to apply the markers we have already found. Scientists are making the effort—and a resource-intensive effort it is—to go off the charts to find more indicators of how the world is. An attempt to balance the trade-off between the two kinds of error can also be seen in a healthy division of labor within intellectual communities along a spectrum from stubborn dreamers who do not want to miss any truth to those who are at least as concerned that we avoid wasting our time on nonsense.

## References

- Barnes, Eric Christian 2008. "Evidence and Leverage: Comment on Roush." *British Journal for the Philosophy of Science* 59: 549–557.
- Carnap, Rudolf 1950. *Logical Foundations of Probability*. Chicago: University of Chicago Press.
- Gettier, Edmund 1963. "Is Justified True Belief Knowledge?" *Analysis* 23: 121–3.
- Goldman, Alvin 1967. "A Causal Theory of Knowing." *Journal of Philosophy* 64: 357–72.
- 1979. "What is Justified Belief?" In G. S. Pappas, (ed.), *Justification and Knowledge*. Dordrecht: D. Reidel, 1–23.
- Lehrer, Keith and Thomas, Paxson, Jr. 1969. *Journal of Philosophy* 66: 225–37.
- Nozick, Robert 1981. *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Stanford, P. Kyle 2006. *Exceeding our Grasp: Science, History, and the Problem of Unconceived Alternatives*. New York: Oxford University Press.
- Roush, Sherrilyn 2005. *Tracking Truth: Knowledge, Evidence, and Science*. Oxford: Oxford University Press.
- 2009. "Optimism about the Pessimistic Induction." *New Waves in Philosophy of Science*. Palgrave MacMillan.