

The Temporal Generality Problem

Brian Weatherson

2012

The traditional generality problem for process reliabilism concerns the difficulty in identifying each belief forming process with a particular kind of process. That identification is necessary since individual belief forming processes are typically of many kinds, and those kinds may vary in reliability. I raise a new kind of generality problem, one which turns on the difficulty of identifying beliefs with processes by which they were formed. This problem arises because individual beliefs may be the culmination of overlapping processes of distinct lengths, and these processes may differ in reliability. I illustrate the force of this problem with a discussion of recent work on the bootstrapping problem.

1 Two Kinds of Generality Problem

The generality problem is a well-known problem for process reliabilist theories of justification.¹ Here's how the problem usually gets started. In the first instance, token processes of belief formation are not themselves reliable or unreliable. Rather, it is *types* of processes of belief formation that are reliable or unreliable. But any token process is an instance of many different types. And these types may differ in reliability.

For instance, imagine I read in the satirical newspaper *The Onion* that Barack Obama is the president. On this basis, I come to believe that Barack Obama is president. The process I have used to form this belief is an instance of each of these types.

1. Coming to believe that Barack Obama is the president;
2. Believing something because it was written in *The Onion*; and
3. Believing something because it was written in a newspaper.

The first type of process is very reliable, at least in 2012. The second is highly unreliable, and the third is very reliable. So should we say that the token process I used was reliable or unreliable? More generally, is there a principled way to map token processes to types of process in a way that lets us systematically say whether a particular process is reliable or not? Critics of reliabilism argue that there is not.

As I said, this problem has been around for quite a while, but I don't think the full force of the problem has been appreciated. Reliabilism is a theory about whether a

¹On process reliabilism, see (Goldman 1979). On the generality problem, see (Feldman 1985; Conee and Feldman 1998)

belief is justified or unjustified. But to determine whether the belief is justified, we step back from the belief itself in two respects. First, we look not to the belief, but to the token process of belief formation from which it results. Second, we look not just to that process, but to kinds of processes of which it is an instant. When carrying this out, we need to make the following two mappings.

4. Belief \rightarrow Token process of belief formation;
5. Token process of belief formation \rightarrow type of process of belief formation

The traditional point of the generality problem is that the second of these mappings is one-many, not one-one. Each token process is associated with many, many types of processes. But what hasn't been sufficiently appreciated is that the first mapping is one-many as well. And this generates a new, and potentially harder, form of the generality problem.

That the first mapping is one-many isn't because of any special properties of beliefs. Typically, an event is the conclusion of more than one process. Imagine that I travel from Michigan to New York to see a friend. I conclude this journey by walking to the friend's apartment. With the last step I take, I conclude several processes. These include:

6. Walking from the subway station to the apartment;
7. Travelling by public transit from the airport to the apartment; and
8. Travelling from Michigan to my friend's apartment.

It is possible that one of these is a quite reliable process, while the others are not. If I am good at navigating the Manhattan street grid by foot, but poor at making it to the airport on time, then process one will be a highly reliable process, while process three will not. So should we say that my arrival at my friend's apartment was the result of a reliable process or not? The best reply to that question is to point out that it is ill formed. Given that I made it to the nearest subway station, I used a reliable process to traverse the last few blocks. But the longer process I used was not as reliable.

This raises a conceptual worry for process reliabilist theories. If there is no such thing as the reliability of a conclusion, but only the reliability of a process of getting from one or other starting point to that conclusion, then it seems that in identifying the justifiedness of a belief with *the* reliability of the process used to generate it, we commit a kind of category mistake. Note that this problem would persist even if we had a one-one mapping from token processes to epistemologically relevant types of processes that would let us solve the traditional form of the generality problem. We would still need a way of saying which of the many processes which terminate in a belief is the epistemologically relevant one. I don't think there's any reason to think there is a good answer to this question. I call this the *Temporal* Generality Problem, because the different processes that culminate in a belief are typically of different durations.

2 Can the Problems be Solved Simultaneously?

I've argued in the previous section that in theory the Temporal Generality Problem is distinct from the traditional version of the generality problem. But one might think that in practice a solution to the latter will solve problems to do with the former. Consider the following three step process.

1. I hear an astrologer say that Napoleon Bonaparte will win the 2013 US Presidential election.
2. I form the belief that Napoleon Bonaparte will win the 2013 US Presidential election.
3. I deduce that there will be a US Presidential election in 2013.

The process by which I got from 2 to 3 is, on the face of it, highly reliable. Assuming that I'm a mostly sensible person, coming to believe obvious logical consequences of my prior beliefs is a highly reliable process. Yet clearly the process that runs from 1 to 3, i.e., the process of believing obvious logical consequences of the contents of astrological predictions, is not a reliable process. So, one might ask, is the resultant belief justified, because it is formed by the reliable process that runs from 2 to 3, or unjustified, because it is formed by the unreliable process that runs from 1 to 3?

Clearly, this is a false dilemma. The salient kind of process I'm using between 2 and 3 is not *believe obvious logical consequences of a belief*, but *believe obvious logical consequences of a belief formed by an unreliable process*. Once we identify the kind of process used at the last stage correctly, we can see that the unreliability of the whole process causes the process used at the last stage to be unreliable.

We might even get cases that go the other way. There are plenty of occasions in science where scientists use mathematical techniques which cannot be made rigorous, and idealisations that cannot easily be replaced with approximations, or with any other statement known to be true.² If we looked at such a step in isolation, we would possibly think that it is an unreliable step, even though it is part of a longer, reliable process. But the fact that it is part of a reliable process matters. In particular, it matters to the way we identify the step the scientist is using with a larger kind of inferential processes. That kind won't involve, for instance, all instances of reasoning from false premises, or of reasoning with incoherent mathematical models. Rather, it will just include the kind of reasoning that is licenced by the norms of the science that the scientist is participating in, and that kind might be a very reliable kind of process.

But there is one very special case where I think this kind of solution to the Temporal Generality Problem will not work. It concerns the way in which a reliabilist will try and solve the bootstrapping problem, as developed by Stewart Cohen (2002) and Jonathan Vogel (2000). We'll turn next to that problem.

²On non-rigorous techniques, see (Davey 2003); on idealisations, see (Davey 2011).

3 Generality and Bootstrapping

Hilary Kornblith (2009) has proposed that looking at processes of longer duration generates a reliabilist solution to the bootstrapping problem. I'm going to argue that Kornblith's solution, which I agree is the kind of thing a reliabilist should say, in fact shows that the Temporal Generality Problem is a distinct kind of generality problem, and perhaps a much harder problem than the traditional generality problem.

Let's start with a very abstract version of the problem. Assume device D is highly reliable, and S trusts device D without antecedently knowing that it is reliable. Then the following sequence of events take place.

- At t_0 , S sees that device D says that p .
- At t_1 , S forms the belief that D says at t_0 that p on the basis of this perception.³
- At t_2 , S forms the belief that p , on the basis that the machine says so.
- At t_3 , S forms the belief that the machine is accurate at t_0 , on the basis of her last two beliefs.

What should a reliabilist say about all this? Well, the process that runs from t_0 to t_1 , the process of believing machine readings are as they appear, looks pretty reliable, so the belief formed at t_1 looks pretty reliable. And the process that runs from t_1 to t_2 , i.e., the process of believing that things are as machine D says they are, also looks pretty reliable, so that belief looks pretty reliable. And the process that runs from t_2 to t_3 , i.e., the process of drawing obvious logical consequences from beliefs formed by reliable processes, also looks pretty reliable. It's true that at t_2 , S doesn't know she's using a reliable process. And hence at t_3 , S doesn't know that this is the kind of process that she's using. But none of this should matter to an externalist like the reliabilist, since they think what matters is actual reliability, not known reliability.

But there are two problems lurking in the vicinity. First, many people think that it is very bizarre that S can form a justified belief that D is accurate at t_0 on the basis of simply looking at D . That's the intuition behind the bootstrapping problem. Second, the case looks like an instance of the Temporal Generality Problem. The two problems are related. Kornblith's solution to the bootstrapping problem is to insist that the process used is in fact *unreliable*. What he means to draw our attention to is that the process which runs from t_0 to t_3 is unreliable. And he's right. That looks like a process of determining whether a machine is accurate by simply looking at the machine and trusting it. Of course, there are several other ways we could classify the process used, but Kornblith argues that this is the best classification, and I think he's right. And if he is right, then we have part of a solution to the bootstrapping problem.

But if Kornblith is right, then we pretty clearly also have a nasty instance of the Temporal Generality Problem. Because now it looks like a chain of three reliable processes, those that run from t_0 to t_1 , from t_1 to t_2 , and from t_2 to t_3 , collectively form an unreliable process. The belief that is formed at t_3 is the culmination of two processes; a

³On some theories of perception, it might be that $t_0 = t_1$, since perception involves belief formation. I don't mean to rule those theories out; the notation here is meant to be consistent with the hypothesis that $t_0 = t_1$.

reliable one that runs from t_2 to t_3 , and an unreliable one that runs from t_0 to t_3 . If a belief is justified iff it is the outcome of a reliable process, and unjustified iff it is the outcome of an unreliable process, then the belief is both justified and unjustified, which is a contradiction.

How could the reliabilist escape this problem? I can see only two ways out. One is to say that the process that runs from t_0 to t_3 is in fact a reliable process. But that's to fall back into the bootstrapping problem. And in any case, it seems absurd, since that process really does look like a process of determining whether a machine is reliable by simply looking at it. The other is to say that the process that runs from t_2 to t_3 is unreliable. To do that, we'd need to come up with a natural kind of process which is unreliable, and which this process instantiates. This does not look easy. I'm not going to insist this couldn't be done, but I'll end by noting three challenges that stand in the way of getting it done, and which seem pretty formidable.

First, if we say the process that runs from t_2 to t_3 is unreliable, then we are putting general restrictions on how we can obtain knowledge by deductive inference. As John Hawthorne (2005) argues, any such restrictions will be hard to motivate.

Second, the restrictions will have to be fairly sweeping to cover the range of conclusions that, intuitively, cannot be drawn through this kind of reasoning. Imagine a variant on the above example where at t_3 , S concludes that either D is accurate at t_0 or it will snow tomorrow. That's entailed, obviously, by what she knows at t_2 . And yet the process of getting from t_0 to that conclusion seems unreliable. So we can't simply say that what's ruled out are cases where the agent draws a conclusion that is simply about D .

Third, the classification of the process that runs from t_2 to t_3 must not merely fail to be ad hoc, it must plausibly be the most natural classification available. And yet it seems there is one very natural classification that is not available, namely the classification of the process as an instance of deduction from known premises, or from premises arrived at by highly reliable processes.

So the challenge this problem raises for reliabilism is substantial. I don't mean to say it is a knock-down drawn-out refutation; philosophical arguments rarely are. But it does add a new dimension to the generality problem, and as we've seen in the last few paragraphs, put some new constraints on solutions to the old version of the generality problem.

References

- Cohen, Stewart. 2002. "Basic Knowledge and the Problem of Easy Knowledge." *Philosophy and Phenomenological Research* 65 (2): 309–29. doi: 10.1111/j.1933-1592.2002.tb00204.x.
- Conee, Earl, and Richard Feldman. 1998. "The Generality Problem for Reliabilism." *Philosophical Studies* 89 (1): 1–29. doi: 10.1023/A:1004243308503.
- Davey, Kevin. 2003. "Is Mathematical Rigor Necessary in Physics?" *British Journal for the Philosophy of Science* 54 (3): 439–63. doi: 10.1093/bjps/54.3.439.

- . 2011. "Idealizations and Contextualism in Physics." *Philosophy of Science* 78 (1): 16–38. doi: 10.1086/658093.
- Feldman, Richard. 1985. "Reliability and Justification." *Monist* 68 (2): 159–74. doi: 10.5840/monist198568226.
- Goldman, Alvin. 1979. "What Is Justified Belief." In *Justification and Knowledge*, edited by George Pappas, 1–23. Dordrecht: Reidel.
- Hawthorne, John. 2005. "The Case for Closure." In *Contemporary Debates in Epistemology*, edited by Matthias Steup and Ernest Sosa, 26–43. Malden, MA: Blackwell.
- Kornblith, Hilary. 2009. "A Reliabilist Solution to the Problem of Promiscuous Bootstrapping." *Analysis* 69 (2): 263–67. doi: 10.1093/analys/anp012.
- Vogel, Jonathan. 2000. "Reliabilism Leveled." *Journal of Philosophy* 97 (11): 602–23. doi: 10.2307/2678454.

Published in *Logos and Episteme*, 2012, pp. 117-122.