Section II. Reply to Brian Weatherson

Thanks to the commentators for their stimulating remarks. First, in response to Brian I want to make it clear that there isn't any commitment to foundationalism in this view of knowledge, though the view does have the three virtues that he mentions. It's not foundationalist because the core of the body of knowledge—the things we track—comprises a huge proportion of our knowledge, and contains complicated propositions for example about electrons, and such knowledge is acquired in complicated ways depending on a myriad of things. A vast proportion of those pieces of knowledge in the core are nothing like what the foundationalist wants in "basic knowledge"—like "there is a table in front of me"—because basic knowledge is supposed to be easy to get, and to get without reference to other knowledge. Nothing like that applies here. (This doesn't mean the foundationalist is wrong, just that this view isn't taking a position.) Of course, Brian doesn't think my view works for those complicated propositions, so I'm going to show it does.

About lottery propositions, I want to make it clear that this has long been known to be an area of *strength* for the tracking view. We have the intuition that even though the proposition "My ticket will not win" is highly probable, still I don't *know* my ticket will not win. On the tracking view it's obvious why: if the ticket *were* going to win, I'd still believe it wasn't—I'm not tracking the truth of this claim. No other view has such a good explanation of why we think we don't know. Of course, by imposing closure I've made things a little more complicated on this topic because allowing that we can know what we know follows from something we track opens up the possibility that there's some *other* proposition we track from which the lottery proposition "My ticket will not win" can be known to follow, such as "I won't be able to pay off that big debt next year". This argument was made earlier by John Hawthorne, and I deal with it in the book, arguing on general grounds that it won't happen. (pp. 132-133) In the example Brian uses, the proposition I'm supposed to track and know that it implies that my ticket won't win is: "I won't be able to afford a trip on a spaceship later this year". Call this "-A". Brian says I track -A. I think it's obvious that I don't. To track -A it would have to be the case that P(-B(-A)/A) > .95 so we have to ask what the probable scenarios are in which I have enough money to go to space. Brian mentions that I'm in line for a big new job. Okay, but this doesn't settle things, first because it's not just the *most* probable scenario that matters. When you use conditional probability *all* the scenarios that have greater than 5% probability matter. (This is a very significant effect of the reformulation of tracking in terms of probability rather than counterfactuals, and eliminates several problems in the original view.) Also, granted my winning of the lottery has less than 5% probability, but the requirement for a scenario to be relevant isn't that it has greater than 5% probability simpliciter but greater than 5% probability *given* A. The scenarios *given* A are the scenarios *given* that I will be able to afford the space trip. That means I'm going to have a boatload of money. How could that happen? The job is one way, but given that I also own a lottery ticket, being the winner has a decent probability too. In that scenario, where I will win, I would still believe I'm *not* going to win and hence that I'm *not* going to have enough money to go to space. So the variation, or sensitivity, condition isn't fulfilled.

The general problem Brian mentions: "if we allow that agents can know propositions that are not absolutely certain, and we accept [closure], [then] agents can know propositions that are arbitrarily improbable," isn't a problem on my view since the only way you can know anything, including the conclusion of the inference, is if you know the premises which means they must be true. A probabilistic version of the problem won't work either since the only way you could know the conclusion was probable is if it *was* probable, or partially know the conclusion is if it was probable *enough*, and so forth. That's there in the way the recursion clause is formulated.

However, there are two interesting problems in this neighborhood of worries about closure that I'll call the *horizontal error problem* and the *vertical error problem*. (See slides.) The horizontal problem is that the more conjuncts that you have in the premises, the more potential (though not actual) error you will have when you infer the conjunction. In fact the potential error grows arithmetically with the number (n) of conjuncts. The conjunction will be true when the conjuncts are true, so the closure clause won't allow you to be counted as knowing something that's false or improbable. But what can happen with a fully general closure clause is for you to continue to count as knowing when your modal tether to the truth is becoming increasingly frayed. That is, though your belief in the conjunction is actually true, as n increases the upper bound on the probability that you believe p when it's false is increasing. It seems that at some point we should stop saying you know because the fact that what you're believing is true becomes more and more accidental.

The vertical error problem arises when we imagine many steps of known implication. With a fully recursive closure clause like I've written it you can build knowledge of p from pairwise knowledge of each of the implications from q to r_1 through r_m to p without necessarily knowing the implication from q to p. That's a problem intuitively because every step of inference introduces a new source of error, so the more steps you have the more potential error you'll have, but once again for my view it's not a problem in the actual world. Each of q through p have to actually be true for you to be counted as knowing them. Still there's a problem, which is that with the increasing number of sources of error, the fact that your belief in the conclusion is true is becoming more and more accidental.

Putting the problem in these ways, it seems to me the solution is pretty clear. There's no reason to specify a single answer to the question Should the view be closed under conjunction or not?, or even After how many conjuncts does conjoining them fail to give you knowledge? or After how many known implications does your knowledge peter out? The ability to tolerate various levels of error varies with the context of application of the theory, in particular with changing utilities of those using it. How much does a given kind of error cost you? In these scenarios the potential error grows as n grows and as m grows, so we should let the context determine at what values for n and m a subject stops being counted as having knowledge. Of course, there isn't a context in which we can tolerate potential error of 50%, for in that case we might as well have thrown a coin. This means that if the error threshold you choose for the tracking conditions is 5% (as I have been assuming for illustration), then you will require n and m to be strictly less than 10. How much less depends on your tolerance for the types of errors that come with inferences. This is a refinement of the view in the book, and it fits well with the contextualism there about the error thresholds on tracking. In both cases the context is determined not by what topic is being discussed, but how much different kinds of errors will cost you.

Now the next problem about tracking and conjunction that Brian describes with the playing cards examples can be stated quite simply: You can get cheap knowledge of q by deduction from known $p \wedge q$ if there is a type of case where you can track that conjunction without being sensitive to q (to any degree whatsoever!). I thought (in the paragraph Brian quotes) that the damage was limited to q that had very very high probability given that the conjunction was false. And if it's rarely false, who cares if you do the wrong thing *when* it's false? What Brian's examples ought to be showing is that q can actually be anything. (I don't understand them so I don't know whether they do.) It's easy to see that the restriction I reported that my definition implies on what kind of q this cheap knowledge can be had for isn't right: it's isn't right that $P(-q/-(p \wedge q)) < 0.05$. There's no restriction on how probable q is given that the conjunction is false, because the assumption that p is false given that the conjunction is false, and that you track p are sufficient for sensitivity to the conjunction. This is because $P(-B(p \wedge q)/-Bp)$ is very high: if you didn't believe p you wouldn't believe the conjunction of p with something else. (That is, provided you weren't a victim of the conjunction fallacy, but we can assume someone capable of training herself to avoid that.)

There seem to be at least two ways to avoid this problem. One is a restriction on closure. We would deny closure from conjunction to conjuncts; that is, we would deny that you can know q by inferring it from a conjunction of q and other p_1 through p_n. This would avoid the awkward conclusion that we know any q as long as we infer it from a conjunction of it with some p we track. But it would still allow us to count as knowing that conjunction $p \wedge q$ where we have no handle at all on q. That seems wrong. The other option I can see is to block this at the level of tracking. That is, we would add the restriction that tracking is not sufficient for knowing a conjunction. To know a conjunction you must either track each conjunct and know that the conjunction follows from the collection as premises, or know that the conjunction follows from something else that you know. (Thanks to Michael Strevens for making this suggestion in the session. I used a similar restriction in the rules of application of the tracking conditions, Chapter 3.)