

In Defence of Single-Premise Closure

1 Introduction

Deductive reasoning is one way by which we acquire new beliefs. Some of these beliefs so acquired amount to knowledge; others do not. Here are two principles, each of which states a *sufficient* condition for acquiring knowledge through deduction:

Single-premise closure (SPC) For any propositions P and Q , and for any subject S , if S knows P and comes to believe Q solely on the basis of competently deducing it from P , while retaining knowledge of P throughout, then S knows Q .

Multi-premise closure (MPC) For any propositions P_1, P_2, \dots, P_n , and Q , and for any subject S , if S knows each of P_1, P_2, \dots, P_n and comes to believe Q solely on the basis of competently deducing it from P_1, P_2, \dots, P_n , while retaining knowledge of each of P_1, P_2, \dots, P_n throughout, then S knows Q .

Although both principles are *prima facie* plausible, SPC is often thought to be on firmer ground than MPC—it's often thought that the phenomenon of risk aggregation poses a problem for the latter but not the former.¹ But Lasonen-Aarnio (2008) and Schechter (2013) have recently challenged this thought. The former argues that, insofar as risk aggregation poses a problem for MPC, it poses a similar

¹The lottery paradox and the preface paradox are often thought to pose problems for MPC (or versions of MPC) while leaving SPC (or versions of SPC) intact. For more on the lottery paradox, see Kyburg 1970; for an attempt to solve it, see Nelkin 2000. For more on the preface paradox, see Makinson 1965; for an attempt to solve it, see Ryan 1991.

problem for SPC. The latter argues that SPC (or a version of it) succumbs to the problem outright.

In this paper, I defend SPC against Lasonen-Aarnio and Schechter—I argue that they fail to show that the phenomenon of risk aggregation poses a genuine problem for SPC.²

2 Do MPC and SPC stand or fall together?

Lasonen-Aarnio (2008) holds that ‘of the probabilistic considerations for rejecting MPC, ... the strongest and least contentious arises in connection with objective chance’ (159). Further, she thinks that ‘arguments against MPC making use of the notion of objective chance can be turned into arguments against SPC (*ibid.*).³

On the face of it, we may know that a proposition is true even when there’s a non-zero probability or risk of its being false. Now, suppose S knows, of each of P_1 , P_2 , ..., and P_n , that it’s true, even while each proposition has some risk of being false. And suppose that S comes to believe Q solely on the basis of competently deducing it from P_1 , P_2 , ..., P_n , while retaining knowledge of each of P_1 , P_2 , ..., P_n throughout. By MPC, S knows Q . But given the aggregation of risk over P_1 , P_2 , ..., and P_n , the risk of Q being false might be high enough that, *pace* MPC, S doesn’t know Q after all.

For instance, suppose that $n = 20$ and suppose that P_1 , P_2 , ..., and P_{20} are probabilistically independent of each other. Suppose also that each of P_1 , P_2 , ..., and P_{20} has a 0.05 chance of being false and that Q is their conjunction. In such a case, the chance of Q ’s being true is approximately 0.36. The worry is that this value is too low for S to know that Q . Even if S believes that Q and even if Q is

²Following Lasonen-Aarnio (2008) and Schechter (2013), this paper focuses on the problem that the phenomenon of risk aggregation is supposed to pose for SPC. One might wish to deny SPC (or versions of it) for other reasons—for instance, to try to solve the problem of external world scepticism. See, for example, Dretske 1970 and Nozick 1981.

³In discussing Lasonen-Aarnio’s arguments, I’ll focus on objective chance—as she does—but Lasonen-Aarnio (2008) holds that similar points can be made with respect to other notions of probability (159).

in fact true, the low chance of Q being true means that it is a matter of *epistemic luck* that S 's belief that Q is true. This disqualifies the belief from amounting to knowledge, given that a belief's amounting to knowledge is incompatible with its being true by luck. Or so an opponent of MPC might argue (Lasonen-Aarnio 2008, 162–3).

The above poses a worry for MPC. But, at first blush, SPC is safe from a similar worry. The risk of the conjunction of P_1 , P_2 , \dots , and P_{20} being false is a function of the risk of each individual proposition's being false—risk accrues over P_1 , P_2 , \dots , and P_{20} . But suppose S comes to believe Q solely on the basis of competently deducing it from P , while retaining knowledge of P throughout. (Henceforth, the clause about retaining knowledge of the relevant premise or premises will often be left implicit.) At first blush, the risk of Q 's being false is a function of just the risk of P 's being false. So, at first blush, risk aggregation poses a problem for MPC but not SPC.

Lasonen-Aarnio (2008) demurs. She writes:

When a subject comes to believe a proposition Q solely based on competent deduction from P , her epistemic standing with respect to Q will depend *both* on her epistemic standing with respect to Q and on the competence of her deduction. But because competence doesn't require infallibility, the risk involved in her belief in P and the *deductive* risk involved in the deduction itself can add up so that despite satisfying the antecedent of SPC, the subject fails to satisfy its consequent, and fails to come to know Q . (*ibid.*, 159–60; Lasonen-Aarnio's emphases)

Suppose that risk aggregation poses a real worry for MPC. According to Lasonen-Aarnio, it'll pose a real worry for SPC too. In the case of MPC, risk accrues over multiple premises. In the case of SPC, we're concerned with just one premise, but risk accrues over this single premise *and* the deduction itself. In some cases, this risk might be great enough that, even should the relevant single-premise deduction result in a true belief, the chance of the resultant belief's being true would be too

low for it to count as knowledge.⁴ Thus the phenomenon of risk aggregation poses a problem for SPC if it poses a problem for MPC. Or so Lasonen-Aarnio argues.

Let's grant that we may know P when its chance of being true is less than 1, and let's grant that knowing P requires that the chance of P 's being false is not too great. In addition, let's grant that there's such a thing as deductive risk—that there is always, or at least often, a very slight chance of a subject's epistemic capacities misfiring' and that 'there is always a non-zero chance of some bizarre quantum-event occurring in our brains that would cause attempted deductions to fail' (*ibid.*, 164–5). Such risk, according to Lasonen-Aarnio, is compatible with one's deduction being competent—as she puts it, 'a subject's deduction can be competent even if there is a very slight objective chance of failure' (*ibid.*, 164). I'll argue, however, that the risk of a single premise being false and the risk involved in performing a deduction do not accrue in any way that poses a genuine threat to SPC. Admittedly, a subject may perform a competent single-premise deduction even if there's a slight chance of failure. But this chance of failure is benign. To see why, it'll help to ask: What is the kind of failure at issue?

Suppose we know P and come to believe Q by competently deducing it from P . One might suggest that our competent deduction has some chance of failing in that it has some chance of leading us from a true belief that P to a false belief that Q . But this suggestion can't be right. A single-premise deduction is competent only if the premise entails the conclusion—Lasonen-Aarnio (2008) agrees with this claim (167). But when P entails Q , there's no chance that our deduction will lead us from a true belief to a false belief. So any chance of Q being false will be due solely to the chance of P being false. More precisely, the chance that our belief that Q is true, in a situation in which we know P and P entails Q , will be at least as high as the chance that our belief that P is true.

But if there is no chance of a competent deduction failing in the way sug-

⁴Williamson (2009) has challenged the claim that knowledge that P is incompatible with there being a low chance that P is true. But I'll assume that the claim is true. This assumption is unproblematic for the purposes of this paper—if anything, it makes SPC harder to defend.

gested above, in what other sense might there be a slight chance of failure? Well, even if we will go on to perform a competent deduction, there might be a chance that we go on to perform an *incompetent* deduction instead.⁵ Suppose we are in fact going to competently deduce Q from $P \ \& \ (P \rightarrow Q)$. Still, since we are fallible, there's a chance that we might instead deduce $\neg Q$ from $P \ \& \ (P \rightarrow Q)$ or deduce some proposition Q^* that resembles Q but is not entailed by $P \ \& \ (P \rightarrow Q)$. For as Lasonen-Aarnio points out, our epistemic capacities might misfire or there might be some strange quantum event going on in the our brains that will cause our attempted deduction to fail.

I'll argue, however, that the chance of a subject performing an incompetent deduction is *irrelevant* to whether SPC holds.⁶ To begin, I'll give an example in which, intuitively, a competent single-premise deduction yields knowledge even though the chance of performing an incompetent deduction is very high.

Consider Beth, a pre-eminent logician with excellent deductive abilities. One day, her jealous rival, Ivan, decides to sabotage her career. While Beth is asleep, an evil scientist in Ivan's employ inserts a chip into Beth's brain. This chip is designed to mess with Beth's attempts at single-premise deduction. Whenever

⁵Lasonen-Aarnio (2008) is right that 'competence doesn't require infallibility' (159).

⁶My distinction between the two senses in which there might be a chance of failure bears some similarity to Smith's (2013) distinction between an unconditional error risk and a conditional error risk. Suppose a subject S competently deduces Q from P . According to Smith (2013), unconditional error risk has to do with the probability that S 's reasoning process will generate a false belief whereas conditional error risk has to do with the probability that such a process will generate a false belief given that it generates a belief that Q (1073). Further, Smith argues that the former kind of risk is epistemically benign, whereas the latter risk is no greater than the risk of P being false (*ibid.*). I'm hugely sympathetic to Smith's defence of SPC. But it is incomplete in certain respects. First, Smith (2013) argues for his claim that unconditional error risk is benign by way of analogy with other cases that do not involve SPC (1075–8). He also thinks that his argument is 'somewhat speculative' (1076). But while I'll consider an analogy too, I'll offer an example in which it's intuitively clear that a subject may gain knowledge via a competent single-premise deduction even though the chance of her performing an incompetent deduction and ending up with a false belief is very high. This gives us direct evidence for thinking that the chance of a subject performing an incompetent deduction—or what Smith calls 'unconditional error risk'—is benign. Second, Smith thinks that his 'objections to Lasonen-Aarnio's argument [against SPC] apply equally to the argument that Schechter mounts' (1072). But he does not address Schechter's (2013) intuition that, after we have performed a long sequence of competent single-premise deductions, we should think it likely that we have made a mistake somewhere in our chain of reasoning. Schechter takes this intuition to support his claim that SPC (or a version of it) is false. To defend SPC thoroughly, we need to resist the intuition or show that, *pace* Schechter, it does not undermine SPC. I adopt the latter strategy in section 3.

Beth's deductive process kicks into motion, there's a 90% chance that the chip will be activated, causing her epistemic capacities to misfire, and making her attempt at deduction go awry. When the chip is activated, Beth might choose an incorrect deductive rule and apply it, or she might choose a correct rule but misapply it.⁷ In either case, it's unlikely that even if she performs the deduction on a premise she knows to be true, she will end up with a true belief. And if she does happen to end up with a true belief, it will be a matter of luck that she does so.

However, the chip's design is not perfect—there's a 10% chance that when Beth starts performing a deduction, the chip won't be activated. In this case, Beth will go on to perform her deduction as if there is no chip implanted in her brain. In particular, she will go on to perform her deduction competently, which involves her choosing the correct deductive rule *and* applying it correctly. Suppose Beth wakes up the day after the chip implant, blissfully unaware of Ivan's machinations. In particular, suppose she is unaware that she has been implanted with a chip. The first thing Beth does, upon waking up, is to carry out a single-premise deduction on P , which she knows to be true. Now, the chance of a successful deduction is very low. But against the odds, as Beth begins her deduction, the chip fails to be activated, and she goes on to competently deduce Q from P —*as she would do normally*. Throughout the process, she is not aware that there's a chip in her brain. Intuitively, the case is one in which Beth may come to know Q by competently deducing it from P .⁸ But if this intuition is right, then we have reason to think that, pace Lasonen-Aarnio, the chance of performing an incompetent single-premise deduction isn't relevant to whether SPC holds. For suppose it is. Then whenever Beth performs a single-premise deduction on a known premise, the chance of the premise being false and the chance of her performing an incompetent deduction will accumulate. But then, the very high chance of an incompetent deduction will mean

⁷Perhaps, even though she should be reasoning by modus ponens, she chooses to reason by some incorrect rule. Or perhaps she correctly chooses to reason by modus ponens but ends up performing such reasoning poorly. Schechter (2013) calls the distinction between applying an incorrect rule and misapplying a correct rule the 'competence/performance distinction' (431).

⁸This intuition seems robust whether the chance of the chip being activated is 90% or as high as 99.9%.

that *any* true belief she gains by competent deduction falls short of knowledge.

Now, the chip's being activated is akin to Beth's epistemic capacities misfiring or to some bizarre quantum event occurring in her brain that would cause an attempted deduction to fail.⁹ But if the chance of the former happening poses no threat to SPC, then the chance of the latter happening should pose no threat as well. Granted, an incompetent deduction in the former case is induced artificially, whereas an incompetent deduction in the latter case might be due to natural causes. But there is no reason to think that this difference is epistemically significant.

Although the Beth case is merely an example, there's a principled reason for holding that SPC is not affected by the chance of a subject performing an incompetent deduction. As Pritchard (2005) notes, certain forms of luck are compatible with knowledge. So there will be certain cases in which knowledge is compatible with low chance. In particular, there's a principled reason for holding that when one knows P and competently deduces Q from P , one's knowing Q is compatible with a low chance of the deduction being performed competently (as in Beth's case). For the kind of luck that such a low chance engenders is compatible with knowledge.

To see why, it'll help to consider an example from Unger (1968), which Pritchard discusses. As Unger notes, knowledge acquired by luck is knowledge

⁹An anonymous reviewer suggested that, perhaps, the former is not akin to the latter—perhaps there is a sense in which the chip is external Beth's mind, and when the chip is activated, it is not really her who is doing the reasoning. If that's right, then one might grant that in the example above, Beth does come to know Q by competently deducing it from P —but the chance of her performing an incompetent deduction was never 90% to begin with, since the chip's activating 90% of the time does not affect the reliability of *Beth's* reasoning.

Even if the chip is external to Beth's mind, it seems to me that when it activates and causes Beth's epistemic capacities to misfire, it is still Beth who is doing the reasoning. After all, our glasses and hearing aids are external to us, but whether they improve or impair our perceptual abilities, we are still the ones who are doing the perceiving. Now, perhaps things are different when long-term and widespread changes in our ability to reason occur. And perhaps, talk of an implant in the brain tends to bring up images of brainwashing. Let us thus modify the example above slightly. Suppose there's a pill that has a 90% chance of making Beth perform a deduction incompetently when she wakes up, and suppose that the effects of the pill are one-off. Suppose also that Beth accidentally takes the pill and is unaware that she has done so. As luck would have it, the pill doesn't work, and when Beth wakes up, she competently deduces Q from P , as she would have done normally. Intuitively, the case is one in which Beth may come to know Q by competently deducing it from P , despite the high chance of her performing an incompetent deduction. Granted, pills and drugs are also external to our mind. But we are prepared to think that some of these substances boost, while others impair, our cognitive performances without being inclined to think that they turn us into another person—at least given that the effects are short-term and local.

nonetheless: a person might know that he will be fired even if he acquires such knowledge by accident, say, by accidentally overhearing his employer say that he will be fired (159). Even though he is *lucky to acquire the evidence that he does in support of his belief* that he'll be fired—and we may imagine that the chance of him acquiring such evidence is low—such a kind of luck, which Pritchard (2005) calls 'evidential epistemic luck', is compatible with knowledge (136). Also, since the employee is *lucky to form the belief that he does*—and we may imagine that the chance of him forming such a belief is low, given that the chance of him acquiring the relevant evidence is low—such a kind of luck, which Pritchard calls 'doxastic epistemic luck', is compatible with knowledge too (*ibid.*, 137–8).¹⁰

Of course, the above is not meant to suggest that forming a true belief based on very good evidence or justification is sufficient for knowledge. After all, we might be in a Gettier situation. Suppose we enter Fake Barn Country but have no evidence that suggests that we have done so. In Fake Barn Country, there are a few real barns, but the majority of the structures that look like barns are—unbeknownst to us—very realistic-looking barn facades (Goldman 1976). Suppose we come across something that looks very much like a barn. Based on our perceptual experience as of a barn, we come to believe that there's a barn in front of us. Our belief seems to be justified. If, as it happens, we are looking at one of the few real barns in Fake Barn Country, our belief is true as well. Still, intuitively, it does not amount to knowledge. For in such an environment, we could easily have looked at a barn facade and formed the same belief based on a qualitatively similar perceptual experience. *Our belief, though true, is true as a matter of luck.* (And notice that the chance of such a belief being true is low; more precisely, in Fake Barn Country, a belief that there's a barn in front of us, based on our perceptual experience as of a barn, has a low chance of being true). Such luck, which Pritchard calls 'veritic epistemic luck',

¹⁰To be more careful, we might hold that the employee's knowing that he'll be fired is compatible with the claim that he is lucky to form the belief that he'll be fired, *where the belief is properly based on the evidence in question.* For suppose he ignores the evidence and forms a belief at random. If, by luck, he happens to form the belief that he'll be fired, then since this belief isn't properly based on the relevant evidence, it does not amount to knowledge.

is unlike evidential epistemic luck and doxastic epistemic luck—it is *not* benign (*ibid.*, 146).

How does the above bear on whether SPC is affected by the chance of a subject performing an incompetent deduction? Recall Beth, implanted with a chip that has a 90% chance of being activated whenever she starts to perform a deduction. Against the odds, when Beth starts performing a deduction on P , which she knows to be true, the chip isn't activated, and she goes on to competently deduce that Q . Although the chance of Beth's performing an incompetent deduction is high, that is no bar to her knowing Q . For even granting that such a high chance means that it is a matter of luck that Beth performs a competent deduction, such luck is benign. Beth is doxastically epistemically lucky: her belief that Q —based on a competent deduction from P —is formed against the odds, given the high chance of the chip activating. (If one thinks that Beth's competently deducing Q from P gives her evidence for believing Q , then one might think that evidential luck is involved too.) But since doxastic epistemic luck is benign, the high chance of her performing an incompetent deduction—whether by employing an incorrect deductive rule or by misapplying a correct deductive rule—will not deprive her of the relevant knowledge.¹¹

Granted, veritic epistemic luck might still deprive Beth of knowledge: if there's a very low chance that her belief that Q is true, Beth might not know Q . But in the case imagined, no such luck is at play. Beth's deduction of Q from P is competent only if P entails Q . But given that Beth knows P and that P entails Q , it is not a matter of luck that Q is true. Speaking in terms of chance, the chance that Beth's belief that Q is true, in a situation in which she knows P and P entails Q , will be at least as high as the chance that her belief that P is true. But given that Beth knows P only if the chance of P being true is high enough for her belief in P not to be true merely by luck, the chance of Q being true will also be high

¹¹Smith (2013) also observes that when the chance of performing an incompetent deduction is high, but a person goes on to perform a competent deduction, the person is doxastically epistemically lucky.

enough for her belief in Q not to be true merely by luck.

The reasoning with respect to Beth generalises. To produce a counterexample to SPC, we need a case in which the antecedent of SPC is true but its consequent is false. In other words, we need a case in which, first, a subject knows P and comes to believe Q solely on the basis of competently deducing it from P , and second, *even given such conditions*, the subject does not know Q . Speaking in terms of luck, a counterexample to SPC will involve a case in which one knows P and competently deduces Q from P , but one's belief that Q is true as a matter of luck. But veritic epistemic luck is *not* at play when we know P and competently deduce Q from P . It might be a matter of luck that we perform a competent deduction and form a belief that Q instead of some other belief not entailed by P (just as it might be a matter of luck that the person in Unger's example forms a belief that he will be fired). But any luck that leads to our competently deducing Q from P will not in itself deprive us of our inferential justification for Q and our knowledge that Q . For such luck is either evidential epistemic luck or doxastic epistemic luck, and both are benign. It is *not* a matter of (veritic epistemic) luck, however, that our belief that Q , having been formed via a competent deduction from P , is true (just as it is *not* a matter of luck that the person's belief that he will be fired, having been formed by hearing his employer say that he will be fired, is true). The chance that Q is true, in a situation in which we know P and P entails Q , is at least as high as the chance that P is true. But given that one knows P only if the chance of P being true is high enough for one's belief that P not to be true merely by luck, the chance of Q being true should also be high enough for one's belief that Q not to be true merely by luck.

3 The Long Sequence Argument

Lasonen-Aarnio's argument that risk aggregation poses a problem for SPC insofar as it poses a problem for MPC does not work. But there's yet another argument against SPC to consider—namely, what Schechter (2013) calls 'the long

sequence argument', which he deploys against the following version of SPC:

Single-premise closure for justified belief (SPCJ) For any propositions P and Q , and for any subject S , if S has a justified belief that P and comes to believe that Q solely on the basis of competently deducing it from P , while retaining the justified belief that P throughout the deduction, then S has a justified belief that Q .

Schechter (2013) thinks that his argument against SPCJ can be tweaked so that it becomes an argument against SPC too. In discussing Schechter's argument, I'll focus on SPCJ (or rather, as we'll see, a slightly amended version of it). But I take it that a successful response to Schechter's argument against SPCJ will, suitably modified, also be a successful response to a corresponding argument against SPC.¹²

Schechter (2013) considers two warm-up counterexamples to SPCJ that he thinks can be answered by modifying SPCJ slightly. First, we sometimes have misleading evidence that we've applied an incorrect deductive rule (436). Suppose twenty of the world's best logicians each tells us that a particular deductive rule that we've employed in a single-premise deduction is incorrect. Even if the deduction is in fact competent, and we are justified in believing the premise, one might have the intuition that we're not justified in believing the conclusion, and take this intuition to pose a problem for SPCJ. Second, we sometimes have misleading evidence that we've misapplied a correct rule. Suppose we've applied a particular rule to perform a competent deduction on a single justified premise. But suppose that someone we have good reason to believe is reliable tells us that a pill we took earlier tends to make people treat certain invalid inferences as instances of a deductively correct rule. In this case, one might likewise have the intuition that we are not justified in believing the relevant conclusion.

Here's a slightly amended version of SPCJ that Schechter (2013) thinks will get around the objections:

¹²Lasonen-Aarnio (2008) discusses a version of the argument briefly (171). I focus on Schechter's version since he discusses the argument in greater detail than does Lasonen-Aarnio.

(SPCJ*) Necessarily, if S has a justified belief that p, comes to believe that q solely on the basis of competently deducing it from p, while retaining the justified belief that p throughout the deduction, *and S does not have a defeater for the claim that the deduction was competently performed*, then S has a justified belief that q. (437; Schechter's emphasis)

SPCJ* gets around the above counterexamples to SPCJ, since both counterexamples involve us having defeaters for the claim that the relevant deduction was competently performed. Further, Schechter (2013) thinks that the qualification that there be no such defeater is 'not *ad hoc*', since '[o]ne of the lessons of contemporary epistemology is that epistemic principles should generally include such clauses' (437). I agree with Schechter. But the following question arises: Is there a counterexample to SPCJ* that does not involve one having a defeater for the claim that one's deduction was competently performed?

Schechter thinks so. His argument against SPCJ* runs as follows:

Consider a very long sequence of competently performed simple single-premise deductions, where the conclusion of one deduction is the premise of the next. Suppose that I am justified in believing the initial premise (to a very high degree), but have no other evidence about the intermediate or final conclusions. Suppose that I come to believe the conclusion (to a very high degree) solely on the basis of going through the long deduction. I should think it likely that I've made a mistake somewhere in my reasoning. So it is epistemically irresponsible for me to believe the conclusion. My belief in the conclusion is unjustified . . .

[The above] provides a counterexample to [SPCJ*]. The difficulty stems from the fact that the long deduction is built out of short simple deductions. So the problem can't be solved by saying that the long deduction is not a competent deduction. Even if the long deduction doesn't count as competent, it would be implausible to claim that one of the short deductions is not competent. (Which one?) So we can apply [SPCJ*] to

the short simple deductions one at a time. Thus, there must be a failure of closure for some short simple single-premise competent deduction.

(438–9)

Note that Schechter (2013) holds explicitly that, in the example above, ‘there is no defeater of the competence of any of the short simple deductions’ (439). The example is meant to be one in which, for each single-premise deduction, there is no reason for one to think that one’s deduction isn’t competent.

I won’t question two ‘key principles’ of Schechter’s long sequence argument (*ibid.*, 440). First, I won’t question what Schechter (2013) calls the ‘Existence of Rational Self-Doubt’, according to which ‘[t]hinkers can (and do) have justified beliefs that their deductive reasoning is not fully reliable’ (*ibid.*). In particular, I agree with Schechter that there could be a long sequence of competently performed single-premise deductions such that, by the end of the sequence, we are justified in believing that we have probably made a mistake somewhere in our line of reasoning. For even if all the deductions have in fact been competently performed, we know that we are fallible and cognitively limited. We may also know that we have been prone to making mistakes when performing similar chains of single-premise deductions on previous occasions.

Second, I won’t question what Schechter (2013) calls the ‘Relevance of Rational Self-Doubt’, according to which ‘[h]aving a justified belief that one’s deductive reasoning is not fully reliable partially defeats one’s justification for the conclusion of a deduction’ (*ibid.*). In particular, I grant Schechter that being justified in believing that we have probably made a mistake somewhere in a long line of reasoning can deprive us of justification for believing the conclusion of that line of reasoning.

Other than not questioning Schechter’s two key principles, I further grant that we can’t ward off his purported counterexample to SPCJ* by saying that the long deduction isn’t competent or by saying that one of the short deductions isn’t competent (*ibid.*, 439).

However, even given all that has been granted, it does not follow that ‘there

must be a failure of closure for some ... single-premise competent deduction'—Schechter's conclusion that SPCJ* is false is drawn too quickly (*ibid.*). Although there's an incremental loss of justification as we move down the chain of single-premise deductions, such loss admits of a prosaic explanation that has nothing to do with the failure of SPCJ*.

It's uncontroversial that we might be justified in believing a proposition initially, but with the passage of time, we lose our original justification due to our fading memories or to new defeating evidence that comes our way. Such loss of justification does not mean that we lacked justification for believing the proposition to begin with. Suppose we gain justification for believing a proposition via perceptual experience. Once we stop having the experience in question, we might lose some justification for the belief as time passes. And if we do, it'll be too hasty to attribute this loss to the justificatory impotence of perception. Perceptual experiences may grant us justification for our beliefs, but they do not guarantee that such justification will be retained or will not be defeated by new evidence that comes our way.

To illustrate, suppose we visit a friend's house and form a justified belief that there is a statue of a dog in the garden, where the belief is based on a relevant visual experience. Long after we've gone home, we retain the belief in our memory, but we no longer have as clear an impression of the experience as before. In such a case, we may well have less justification for the belief than we had initially. After all, our memories are imperfect and we know this. So it is reasonable for us to have at least some doubt as to whether they are in good order with respect to the belief. We may wonder, for example, whether we really did have the visual experience in question. This does not mean, however, that the belief did not enjoy the amount of justification it had initially.

Similarly, suppose we gain justification for believing a proposition via a competent single-premise deduction. We might subsequently lose some of this justification. And if we do, it'll be too hasty to attribute this loss to the failure of SPCJ*.

After all, SPCJ* doesn't—and shouldn't—guarantee that we'll retain any justification gained via a competent single-premise deduction. Although this is a mundane point, it bears on the validity of the long sequence argument, as the following example will show.

At time t , Sue—a typical human being with less than perfect memory—has a justified belief that A and competently deduces B from A . Ten years later, seeming to remember that B , she competently deduces C from B . Ten years after that, seeming to remember that C , she competently deduces D from C . And so on, until seventy years after t , she deduces H from G . Here, we have a sequence of competently performed single-premise deductions, where the conclusion of one deduction serves as the premise of the next. We may suppose that Sue is justified in believing the initial premise A , but has no other evidence about the intermediate or final conclusions. We may also suppose that Sue comes to believe the final conclusion solely on the basis of going through the sequence of deductions. Should Sue think it likely that she's made a mistake somewhere in her reasoning? Is her belief in the conclusion unjustified?

Even if we answer 'yes' to both questions, it would be too quick to conclude that the foregoing is a counterexample to SPCJ*. Let's grant that Sue is not justified in believing H even though she has a justified belief that A , and her belief that H is based on a valid chain of inferences from A to H . And let's grant that there is an increasing loss of justification as we move through Sue's chain of deductions. It would be too quick, however, to pin this loss on SPCJ*'s failure. A more prosaic explanation is available.

Although memory preserves justification, it does not preserve it perfectly—at least for non-ideal beings like Sue. Suppose that immediately upon performing the first deduction from A to B , Sue's belief that B is justified. But when she performs her second deduction from B to C ten years after t , it's unlikely that her belief that B will be as justified as it was at t . First, just before Sue performs her second deduction, her memory of B being true may well be less vivid than it

was immediately after the first deduction (just as our memory of the statue in our friend's garden is less vivid than it was before). Second, given the long passage of time that has elapsed and a less than perfect memory, all she might remember is that she deduced B from A at one point—if she even remembers that at all—without remembering the details involved in the performance of the deduction. Given that she knows that she is fallible, short of performing the deduction again, she may even doubt that she performed it correctly ten years ago (just as we may doubt whether we really had a certain visual experience on which our current belief about the statue was originally based). Such doubt may then lead to her losing some justification for believing B . Similar points hold with respect to any of Sue's subsequent deductions.¹³

Note that the loss in justification in the example above does not impugn SPCJ*. To show that SPCJ* is false, Schechter needs to show that there are some cases in which we have a justified belief that P , competently deduce Q from P , and yet, at the very point when the deduction has been completed, enjoy less justification for the belief that Q than for the belief that P . But Sue's loss of justification is compatible with SPCJ* holding true. Her imperfect memory, which fails to preserve justification perfectly across time, leads to such loss after a deduction or in between deductions. But the failure of memory to *preserve* justification perfectly across time is no bar to one's *acquiring* justification via single-premise deductions (just as it would not affect her acquiring justification via perception, although any such acquired justification may subsequently grow weaker). At the end of each deduction, her belief in the relevant conclusion may well enjoy as much justification as her belief in the relevant premise. It is only after each deduction that, with the passing of the years, justification is lost.¹⁴ For example, when Sue deduces B from A , she

¹³Cf. Burge (1995). To be clear, I'm not claiming that successful deductions must appeal to premises about one's memory. As Burge writes, '[i]t is one thing to rely on memory in a demonstration, and another to use premises about memory'; and while '[a]ny reasoning in time must rely on memory . . . not all reasoning must use premises about memory or the past'. (*ibid.*, 276–7)

¹⁴One might think that, even as Sue is performing a single-premise deduction competently, and the deduction is fresh in her mind, her having reason to believe that she is fallible means that she has reason to believe that there's a possibility that she'll make a mistake this time, which in turn

may have as much justification for believing B as she has for believing A . But after ten years, as her memories of B being true and of having performed a competent deduction degrade, her justification for believing B decreases to a certain amount. Now, at that point, when she deduces C from B , she may enjoy this same amount of justification for believing C . But again, as the years go by, this amount of justification for believing C will decrease further, and so on.

The example above is constructed partly for dramatic effect. Similar points hold even if the relevant deductions take place a much shorter time apart, as is probably intended to be the case in Schechter's example. Suppose that Sue's deductions take place at one-second intervals. Given the shorter duration of time between deductions, the effects of her imperfect memory as she goes from one deduction to another might be hardly noticeable. But if we increase the length of the sequence of deductions, even tiny losses of justification can accumulate—to the extent that we would not be justified in believing the final conclusion. But as before, such losses of justification do not impugn SPCJ*.

One might wish to construct an example in which a subject has a perfect memory, or in which all the deductions are done instantaneously, so that any loss in justification cannot be attributed to the subject's having an imperfect memory. Consequently, one might think that, in such a case, the intuition that we are not justified in believing the conclusion of a long sequence of competent single-premise deductions is to be explained by the failure of SPCJ*. But such an example would be

means that she should entertain some doubt as to the competence of her deduction. One might then take this to be a problem for SPCJ*—one might take such doubt to show that Sue can have enough justification for believing a premise without having enough justification for believing the relevant conclusion. Note, however, that the foregoing argument is crucially different from the long sequence argument. The latter depends largely on the intuition that, at the end of a long sequence of competent single-premise deductions, one is not justified in believing the final conclusion. The failure of SPCJ* is then taken to be the best explanation of this intuition. The current argument, however, takes aim at SPCJ* directly. Now, a discussion of the argument's merits will take us beyond the aim of this paper, which is to defend SPC and SPCJ* against specific arguments from Lasonen-Aarnio and Schechter. But it is at least worth noting that the relevant intuition with respect to the long sequence argument is reasonably strong, whereas it is less obvious whether the following claim is true: Even while one is performing a deduction competently, merely having reason to believe that one is fallible should deprive one of some justification for believing the conclusion of the deduction. In fact, one might think that such a claim is precisely what the long sequence argument aims to support.

unrealistic. It would not involve non-ideal beings with limited cognitive abilities like you and me—beings who do not have perfect memories and who take a relatively fair amount of time to perform even simple single-premise deductions. Now, I take it that Schechter intends his original example to involve a typical non-ideal human being. For the example derives much of its initial force from our imagining *ourselves* in a similar situation. When we imagine ourselves performing a long sequence of deductions, we have the intuition that we should think it likely that we have made a mistake in our reasoning and that we are not justified in believing the final conclusion. But if the example is tweaked, and the subject involved has a perfect memory or can perform a long sequence of deductions instantaneously, it's no longer clear what the relevant intuition should be—the example loses its initial force.¹⁵

One might also object that although it's possible to appeal to our imperfect memories in order to offer an explanation of the aforementioned intuition, I haven't completely ruled out an explanation that postulates the failure of SPCJ*. Perhaps the loss of justification that occurs as we go through a chain of deductions is overdetermined. Or perhaps some justification is lost due to our imperfect memories, and some is lost due to the failure of SPCJ*. But here, it is useful to recall the aim of this paper. As mentioned, conventional wisdom has it that the phenomenon of risk aggregation poses a problem for MPC but not for SPC. Lasonen-Aarnio and Schechter have sought to challenge conventional wisdom, whereas I've sought to defend it. My aim is not to offer arguments for believing SPC or SPCJ*, but to defend such principles against arguments that have been levelled at them. Schechter's long sequence argument is one such argument, and the onus is on Schechter to provide

¹⁵What if, instead of considering a subject with perfect memory, we consider human beings with better or worse memories? (Thanks to an anonymous reviewer for raising this question.) It seems that no matter who we pick, we'll have the intuition that if a relevant sequence of competent deductions is long enough, he or she she will not be justified in believing the final conclusion. Does this show that an appeal to memory does not help explain the relevant intuition after all? I don't think so. For each human being with a less than perfect memory, there will be a sequence of competent deductions that is long enough that he or she will not be justified in believing its conclusion. But it seems that, all things being equal, the worse (or better) a person's memory, the shorter (or longer) such a sequence will need to be.

a compelling counterexample to SPCJ*. But I've argued that he has not succeeded in doing so. The intuition that we're not justified in believing the final conclusion of a long sequence of competent single-premise deductions admits of a prosaic explanation. In the absence of reason to think that such an explanation cannot be right, it'll be hasty to hold that the intuition has to be explained by the failure of SPCJ*.

4 Conclusion

Lasonen-Aarnio (2008) argues that insofar as risk aggregation poses a problem for MPC, it poses a problem for SPC too, whereas Schechter (2013) argues that risk aggregation poses a problem for SPC outright. I've argued that neither has succeeded in creating any genuine problem for SPC.¹⁶ Risk aggregation does seem to pose a problem for MPC. But as things stand, it leaves SPC (or SPCJ*) intact.¹⁷

¹⁶For more discussion of SPC (that has a different focus from this paper), see Hawthorne (2004), Wedgwood (2012), Williamson (2009).

¹⁷Thanks to an anonymous reviewer and to the audience at the 2015 Australasian Association of Philosophy Conference for valuable comments. I've also benefitted from discussion with Michael Pelczar, Ben Blumson, and the honours students in my epistemology class in 2015. Thanks, in particular, to Bernadette Chin for feedback on a draft of the paper.

References

- Burge, T. (1995). Content preservation. *Philosophical Issues* 6, 271–300.
- Dretske, F. (1970). Epistemic operators. *The Journal of Philosophy* 67, 1007–23.
- Goldman, A. I. (1976). Discrimination and perceptual knowledge. *The Journal of Philosophy*, 771–91.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford University Press.
- Kyburg, H. E. (1970). Conjunctivitis. In M. Swain (Ed.), *Induction, Acceptance, and Rational Belief*, pp. 55–82. Dordrecht, Holland: D. Reidel Publishing Company.
- Lasonen-Aarnio, M. (2008). Single premise deduction and risk. *Philosophical Studies* 141, 157–73.
- Makinson, D. C. (1965). The paradox of the preface. *Analysis* 25, 205–7.
- Nelkin, D. K. (2000). The lottery paradox, knowledge, and rationality. *The Philosophical Review* 109, 373–409.
- Nozick, R. (1981). *Philosophical Explanations*. Oxford: Oxford University Press.
- Pritchard, D. (2005). *Epistemic Luck*. Oxford: Clarendon Press.
- Ryan, S. (1991). The preface paradox. *Philosophical Studies* 64, 293–307.
- Schechter, J. (2013). Rational self-doubt and the failure of closure. *Philosophical Studies* 163, 429–452.
- Smith, M. (2013). Two notions of epistemic risk. *Erkenntnis* 78, 1069–79.
- Unger, P. (1968). An analysis of factual knowledge. *The Journal of Philosophy* 65, 157–70.
- Wedgwood, R. (2012). Justified inference. *Synthese* 189, 273–295.

Williamson, T. (2009). Probability and danger. *The Amherst Lecture in Philosophy* 4. URL: <<http://plato.stanford.edu/archives/win2012/entries/probability-interpret/>>, 1–35.