

Single Premise Deduction and Risk

Almost final draft, published in *Philosophical Studies*, Vol, 141, no. 2, pp. 157-173, 2008

Maria Lasonen-Aarnio

It is tempting to think that multi premise closure creates a special class of paradoxes having to do with the accumulation of risks, and that these paradoxes could be escaped by rejecting the principle, while still retaining single premise closure. I argue that single premise deduction is also susceptible to risks. I show that what I take to be the strongest argument for rejecting multi premise closure is also an argument for rejecting single premise closure. Because of the symmetry between the principles, they come as a package: either both will have to be rejected or both will have to be revised.

1.

Competent deduction from a known proposition is one way of acquiring new knowledge. Few would deny this, but epistemologists disagree about whether the antecedent of the following principle gives sufficient conditions for acquiring new knowledge by deduction:

Single Premise Closure (SPC)

For all propositions P , Q , and all subjects s , if s knows P , and s comes to believe Q solely based on competent deduction from P , while retaining knowledge of P throughout, then s knows Q .¹

Those who reject Single Premise Closure also have to reject Multi Premise Closure, as the latter entails the former:

¹ Cf. Williamson (2002: 117) and Hawthorne (2004) for similar formulations of closure. As should become clear below, my main points don't depend on this particular formulation of closure.

Multi Premise Closure (MPC)

For all propositions P_1, \dots, P_n, Q , and all subjects s , if s knows each of P_1, \dots, P_n , and s comes to believe Q solely based on competent deduction from P_1, \dots, P_n , while retaining knowledge of each of P_1, \dots, P_n throughout, then s knows Q .

However, there is no logical inconsistency in retaining SPC while rejecting MPC. Indeed, it is tempting to think that though paradoxes and problems created by the former – in particular, the sceptical paradox – are also created by the latter, MPC gives rise to its own paradoxes having to do with accumulation of risks.

I will argue that there is a so far unacknowledged reason for thinking that SPC fails, a reason that is distinct from the more well-known complaint that closure fails simply because subjects cannot know denials of sceptical hypotheses.² The problem I present arises from the realisation that in so far as actual subjects can deduce competently, competence should not require infallibility, or zero chance of error. The problem for SPC that this creates is very similar to a more familiar reason for thinking that MPC fails, having to do with the accumulation of risks involved in the premises of a deduction. Indeed, my main claim will be that the two closure principles are symmetrical: there is no special class of paradoxes that MPC creates, since analogues of the paradoxes that have been thought to be specific to MPC in fact also arise for SPC. Consequently, this class of paradoxes cannot be avoided just by giving up MPC.

² See, in particular, Dretkse (1970) and Nozick (1981). Similarly, to avoid a more recent sceptical paradox exploiting “lottery propositions”, one might maintain that while it is possible to know ordinary propositions, such as the proposition that one won’t be able to afford an expensive safari this year, it is not possible to know lottery propositions, such as the proposition that one’s ticket will lose the lottery. For sceptical arguments exploiting the intuition that subjects cannot know such propositions, see Vogel (1990) and Hawthorne (2004). But note that it is not straightforward that an account of why it is possible to know ordinary propositions but not denials of classical sceptical hypotheses will generalise into an account of why it is possible to know ordinary propositions but not the lottery propositions they entail. For instance, for problems having to do with Nozick’s tracking condition, see Hawthorne (2004: 10–14).

Let me begin by stating the probabilistic motivation for thinking that it is possible to retain SPC while rejecting MPC. The probabilistic analogue of SPC holds, but not that of MPC, since single premise entailment preserves high probability, but multi premise entailment doesn't. The axioms of probability entail that

If P entails Q , then $\Pr(Q) \geq \Pr(P)$.

But they do not entail that

If P_1, \dots, P_n entail Q , then $\Pr(Q) \geq \Pr(P_1) \& \dots \& \Pr(Q) \geq \Pr(P_n)$.

Highly likely premises can entail a conclusion that is highly unlikely. For instance, each ticket in a lottery has a high chance of losing, but the chance that every ticket will lose is not high.

Connections between probability and rational belief or rational acceptance, or probability and knowledge, motivate rejecting the idea that rational belief or knowledge is closed under multi premise entailment. For instance, assume that it is rational to accept a proposition if and only if it is rational to assign a high credence to it, where the threshold for high credence is less than 1. It might be rational for a subject to assign a high credence to each proposition in a large set of propositions $\{P_1, \dots, P_n\}$, and to assign a high credence to the proposition that at least one of these is false, i.e., to $\sim(P_1 \& \dots \& P_n)$. Unless it is rational to accept the contradiction $(P_1 \& \dots \& P_n) \& \sim(P_1 \& \dots \& P_n)$, rational acceptance had better not be closed under multi premise entailment.³ If

³ Cf. Kyburg's (1970) Lottery Paradox. See also Makinson (1965) for the Paradox of the Preface. Note that denying that rational acceptance is closed under multi premise entailment won't help escape the fact that it will still be rational to accept each proposition in a set of propositions entailing a contradiction.

rational acceptance or belief is necessary for knowledge, then these considerations will also give rise to failures of MPC. Nevertheless, nothing in these considerations threatens the idea that rational belief is closed under single premise entailment.⁴

Generally, MPC becomes highly suspect if one accepts the following claims: it is possible to know propositions with some non-zero probability or risk of being false, but there is a limit to how improbable or risky known propositions can be.⁵ This conjunction of claims opens up the possibility of counterexamples to MPC, for the conjunctions of known, highly probable propositions can fail to be highly probable. The problem arises whether probability is understood as subjective probability, evidential probability, or objective chance. In short, if the *risk* of proposition P is $1 - \Pr(P)$, then risks can accrue over multi premise entailments, but not over single premise entailments. That is the probabilistic motivation for thinking that a plausible epistemological theory might be able to retain SPC, but give up MPC.⁶

It is clearly logically possible to retain SPC while rejecting MPC. However, of the probabilistic considerations for rejecting MPC, I think the strongest and least contentious arises in connection with objective chance. But I will argue that arguments against MPC making use of the notion of objective chance can be turned into arguments against SPC. Given a notion of risk relevant for knowledge – in particular, a notion that is connected to

⁴ Indeed, Kyburg's (1970) own solution to the paradox is to reject the idea that rational acceptance is closed under conjunction, which is a special case of multi premise entailment. If rational acceptance is assumed to be a necessary condition for knowledge, then these considerations provide a rationale for a view that rejects MPC but retains SPC.

⁵ Cf. Williamson (2000: 117). See also Hawthorne (2004: 47).

⁶ For instance, Hawthorne (2004: 47) seems to think that the problem for closure created by the accrual of risks is specific to MPC: 'Granted, deductive inference from a single premise does not seem like a candidate for risky inference. If p entails q then q must be logically weaker or equivalent to p '.

an epistemically relevant variety of luck – the seeming asymmetry between SPC and MPC breaks down.

In short, the problem with SPC is the following. 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 P *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 . Call these *accumulation of risk* failures of single premise closure.

I should point out at the outset that I will be ignoring frameworks that assume logical omniscience.⁷ The assumption of logical omniscience is created, for instance, by individuating propositions in a very coarse-grained manner, so that there is only one necessarily true proposition. Assuming that knowledge distributes over conjunction, a coarse-grained way of individuating propositions will guarantee that SPC is more or less trivially true, and is compatible with rejecting MPC. But given these assumptions, subjects would never be able to acquire new knowledge by deduction from a single premise in the first place.⁸

In §2 I begin by explaining why I have formulated single premise closure as I did above, instead of including in its the antecedent the condition that the subject knows that

⁷ For such frameworks, see Hintikka (1970), Stalnaker (1999), and Lewis (1996).

⁸ Assume that P entails Q . Then, P is logically equivalent to the conjunction $P \ \& \ Q$. If propositions are individuated in a coarse-grained manner, then these will be one and the same proposition. Any subject who knows P also knows $P \ \& \ Q$. If knowledge distributes over conjunction, then any subject who knows P also knows Q , whether or not she satisfies the condition that she believes Q based on competent deduction from P .

P entails Q . In §3 I give what I take to be the strongest case against MPC. In §4 I argue that SPC faces the same worry, and say why the problem generalises. In §5 I discuss competent deduction within the frameworks of reliabilist and safety theories of knowledge. I also take a theory of knowledge on which MPC fails and construct a concrete case constituting a counterexample to SPC on that theory. In §6 I address some objections, thereby clarifying what I take to be the philosophical significance of my result.

2.

Let me first say why I have formulated closure as I did above.

One often sees closure principles with antecedents that require subjects to have knowledge that a proposition P entails a proposition Q .⁹ Call these *knowledge of entailment* closure principles, as opposed to *competent deduction* principles. The simplest form of such a principle, employed by sceptics of various persuasions, is the following, where ‘ \rightarrow ’ stands for entailment:

For all propositions P , Q , and subjects s , if s knows P , and s knows $(P \rightarrow Q)$, then s knows Q .¹⁰

But of course, this principle fails. Counterexamples can be generated in which a subject s satisfies the antecedent but fails to satisfy the consequent because i) s simply fails to believe Q , ii) s 's belief in Q is wholly based on the wrong reasons, or because iii) s 's

⁹ For instance, Dretske (2005), (1970), Nozick (1981).

¹⁰ Cf. DeRose & Warfield (1998: 13). More formally, using ‘ $K_s(P)$ ’ for ‘ s knows P ’, and ‘ \rightarrow ’ for strict (logical) implication, the principle reads: For all P , Q , s , if $K_s(P)$ and $K_s(P \rightarrow Q)$, then $K_s(Q)$. Alternatively, the principle could be formulated as follows: For all P , Q , s ($K_s(P)$ and $K_s(P \rightarrow Q)$) \rightarrow $K_s(Q)$). Similarly for all of the closure principle I discuss. I won’t fuss over this issue, but prefer to use the potentially weaker, first formulation.

belief in Q is partly based on the wrong reasons. There are also more contentious worries having to do with defeat that I shall ignore for now.¹¹

Here is a principle that at least avoids the obvious, least contentious counterexamples:

Knowledge of Entailment Closure

For all propositions P , Q , and subjects s , if s knows P , and s knows $P \rightarrow Q$, and s believes Q solely based on competent deduction from P , $P \rightarrow Q$, while retaining knowledge of both P and $P \rightarrow Q$ throughout, then s knows Q .¹²

But Knowledge of Entailment Closure doesn't give sufficient conditions for coming to know a proposition Q by deduction from a proposition P . Satisfying the antecedent of the principle requires that a subject knows that the entailment $P \rightarrow Q$ holds, but not that her knowledge is based on deduction from P . And there are a variety of ways other than deduction in which a subject might come to know $P \rightarrow Q$ – for instance, by relying on the testimony of expert logicians. If a subject whose knowledge of $P \rightarrow Q$ is based on the testimony of a logician comes to believe Q solely based on deduction from the premises P , $P \rightarrow Q$, her knowledge of Q won't be based on competent deduction from P .¹³

Someone might argue that to give sufficient conditions for acquiring knowledge of Q by deduction from P , Knowledge of Entailment Closure must be revised by specifying that the subject's knowledge of $P \rightarrow Q$ be based on competent deduction of Q from P . Even if the goal of specifying sufficient conditions were thus achieved, the resulting principle would certainly not be the simplest way of capturing the idea that a

¹¹ See David & Warfield (*forthcoming B*).

¹² Cf. DeRose & Warfield (1998: p. 23, note 14) for a principle otherwise similar but without 'solely', and David & Warfield (*forthcoming A*).

¹³ Of course, such cases don't constitute a counterexample to Knowledge of Entailment Closure.

subject can acquire knowledge by competent deduction from a single premise. We want a principle giving sufficient conditions that has as little in it that is not necessary as possible. And on pain of an infinite regress, knowledge of the entailment $P \rightarrow Q$ cannot be a necessary condition for coming to know Q by competent deduction from P . For if a subject needs the premise $P \rightarrow Q$ in order to competently deduce Q from P , surely the same must hold when she deduces Q from $P, P \rightarrow Q$ – she will need something like the premise $(P \text{ and } P \rightarrow Q) \rightarrow Q$ to do this. And so forth.¹⁴ Generally, competently deducing Q from P cannot require a subject to deduce Q from the two premises $P, P \rightarrow Q$.

There is not even any reason to think that Knowledge of Entailment Closure captures something about how subjects normally, or paradigmatically, acquire knowledge by deduction. Moreover, its antecedent will never be true of creatures that lack the concept of entailment. Finally, one might fear that Knowledge of Entailment Closure faces the same accumulation of risk problems as MPC.¹⁵

Hence, there are reasons for preferring my formulation of single premise closure, SPC, to Knowledge of Entailment Closure, and similarly for corresponding formulations of multi premise closure. However, what I say below does not essentially depend on preferring SPC to Knowledge of Entailment Closure.

¹⁴ Cf. Carroll (1895).

¹⁵ DeRose (1999: 23, note 14) acknowledges a similar point: ‘Just as I think one can just barely know that P and just barely know that Q , yet just barely fall short of knowing the conjunction of P and Q – even when holding constant the standards for knowledge – because of the accumulation of doubt, so also can one know that P and know that P entails Q , while failing short of knowing that Q , even if one’s belief that Q is based on one’s knowledge of P and knowledge of the fact that P entails Q ’.

3.

Any theory allowing subjects to know propositions with some likelihood or risk of being false will struggle with MPC, for to retain this principle, they will have to endorse consequences that seem highly unintuitive. Depending on the relevant notion of likelihood or risk, they will have to concede that it is possible to know propositions with very low objective chances of truth; propositions to which a subject assigns a very low subjective probability; propositions that are highly unlikely on a subject's evidence; or propositions that a subject has a justification to believe are false.

I think the strongest case against MPC arises in connection with objective chance.¹⁶ The argument I have in mind has three premises. The first of these is a non-sceptical premise. It states that it is at least possible to know a lot of probabilistically independent propositions with non-zero chances of falsity.¹⁷ Call this the *knowledge premise*. In particular, subjects know a lot of things about the future, even assuming the actual world to be indeterministic. For instance, assume that there is a very slight chance that if I drop my pen on the table, it will tunnel through. Despite this, I know that if I drop my pen, it will land on the table. Below I discuss other ways in which slight chances of error can enter into a subject's belief.

Second, any theory of knowledge ought to be able to rule out Gettier-cases as cases of knowledge. A rather compelling diagnosis of what goes wrong in cases in which a subject has a justified, true belief but lacks knowledge is that the subject's belief is true as the result of a problematic bout of good luck.¹⁸ Hence, there is a kind of epistemic luck

¹⁶ Cf. Hawthorne (2004: 47–48).

¹⁷ Whether or not the actual world is indeterministic, epistemologists should be wary of allowing the anti-sceptical force of their theories to depend on such empirical matters.

¹⁸ Cf. Unger (1968). See also Pritchard (2005).

that is incompatible with knowledge. Call this the *no luck intuition*. One way of spelling out what such luck involves would be, for instance, to say that it is present just in case a subject could easily have held a (relevantly similar) false belief.¹⁹ But all that really matters for my purposes here is that there be a connection between chance and luck.

The third premise of the argument is the following: if there is a high chance at t that P is false, but P is nevertheless true, then it is lucky at t that P is true. Call this the *high chance intuition*. In particular, if at t a subject truly believes P , and at t there is a high chance that she holds a false belief in P , then it is lucky that her belief in P is true. This connection between chance and luck explains why subjects cannot normally know low-chance propositions.²⁰

The problem should by now be obvious. Assume that subjects in an indeterministic world (such as, I take it, the actual world) can know a lot of things about the future.²¹ Enough suitably independent high-chance propositions will conjoin into a low-chance conjunction. And indeed, the conjunction of everything an ordinary subject takes herself to know about the future will most likely have a low chance. By MPC, as long as a subject knows all of the individual conjuncts, she can come to know such a low-chance conjunction. But because the conjunction has a high chance of being false, by the high chance and no luck intuitions her belief in it cannot count as knowledge. Here is the structure of the argument:

1. The knowledge premise.

¹⁹ Williamson (2000) defends this idea, which has come to be known as a safety condition on knowledge.

²⁰ I say ‘normally’ because of certain instances of knowledge of singular, contingent a priori propositions that have low chances. See Hawthorne & Lasonen-Aarnio (*forthcoming*).

²¹ If one thinks that even deterministic worlds have objective chances other than 0 or 1, then the argument need not be stated by looking at subjects in indeterministic worlds.

2. The no luck intuition.
3. The high chance intuition.

Therefore,

4. MPC is false.

Generally, theories of knowledge that are able to retain the knowledge premise – and in particular, that allow subjects to have knowledge of propositions with non-zero chances of falsity – the no luck intuition, and the high chance intuition, lead to failures of MPC.

For instance, take a reliabilist theory on which knowledge of a proposition *P* entails that the subject's belief in *P* was formed by a belief-forming process that is sufficiently reliable, its reliability depending roughly on the ratio of true beliefs to all beliefs it produces in actual and counterfactual cases. If the reliability of a belief-forming process depends on the objective chances of the beliefs it gives rise to, then the above point will apply directly to reliabilist theories. For instance, take a fair lottery with 10 tickets, exactly one of which will win. Even if the degree of reliability of the belief that Ticket #1 will lose is 0.9 and likewise for the degree of reliability of the belief that Ticket #2 will lose, the same doesn't hold for the conjunction (Ticket #1 will lose & Ticket #2 will lose). The reliability of beliefs is not closed under competent deduction.

Alternatively, take a safety requirement for knowledge according to which a subject's belief in *P* is safe as long as the subject avoids false belief in most close possible worlds with relevantly similar initial conditions.²² In a fair lottery with sufficiently many tickets, each ticket that actually loses does so in most close worlds,

²² Pritchard (2005) initially defends such a position, though he then opts for a more demanding version of safety.

since it has a very high chance of losing. But given any suitably large subset of tickets, there will be a low chance that all of the tickets in the set lose. Hence, it looks as though a subject can be safe in believing each individual proposition P_1, \dots, P_n , but not safe in believing their conjunction. Similar points apply to various other possible views, both internalist and externalist. For instance, if the evidential threshold required for knowledge is set at a level below 1, then a subject might have evidence for each proposition in a set of propositions S that is strong enough for knowledge, but not have sufficiently strong evidence for their conjunction.

Analogous arguments could be restated with any notion of risk replacing objective chance. All that matters is that risks pile up over conjunctions, and that knowledge is tolerant to the relevant kind of risk in small, but only small, quantities. For instance, the notion of chance in the above argument could be replaced with the notion of degrees of reliability: 1. Knowledge doesn't require 100 % reliability; 2. Knowledge is incompatible with belief that is true by luck; 3. Low reliability makes for belief that is, if true, true by luck; Hence, 4. MPC is false.

To avoid failures of MPC, some further condition on knowledge needs to be adopted that would safeguard against them. But anyone convinced by the above argument will want the new constraint to be non-sceptical, and to respect the no luck intuition as well as the high chance intuition. And given any such constraint, we will be able to construct cases in which a subject satisfies the requirement for each proposition P_1, \dots, P_n , but fails to satisfy it for a proposition Q entailed by P_1, \dots, P_n , despite believing Q solely based on competent deduction from P_1, \dots, P_n .

My purpose here is not to defend the above argument, though I do think it minimally shows that retaining MPC in indeterministic worlds means giving up some very plausible assumption. My main aim is to show that if the above argument creates trouble for MPC, it also does so for SPC.

4.

The above problem for MPC relies on very plausible-seeming epistemological principles connecting knowledge, luck, and high chance, together with a non-sceptical assumption. Theories of knowledge that accept the three premises of the argument, or some analogues formulated in terms of a notion of risk other than objective chance, will have to give up MPC. But one might hope that this wouldn't mean giving up SPC.

The argument against MPC assumed that subjects can know propositions with non-zero chances of falsity. For this to be the case, knowledge must be tolerant to slight chances of error. But if knowledge is compatible with a slight objective chance of holding a false belief, then it seems arbitrary to rule that such chances of error can't enter into the process generating a belief. For instance, it isn't at all far-fetched to think that there is always, or at least often, a very slight chance of a subject's epistemic capacities misfiring. Generalising this to the case of deduction, a subject's deduction can be competent even if there is a very slight objective chance of failure. For all we know, there is always a non-zero chance of some bizarre quantum-event occurring in our brains that would cause attempted deductions to fail. Call such risks *deductive*.

Of course, if at a time t a subject has already formed a belief in a true proposition P based on deduction, then at t the chance that her inferential faculties malfunction at a

slightly earlier time will be 0 (assuming a somewhat standard view on past events have chances of either 0 or 1). Nevertheless, it seems clear that the earlier chance of forming a false belief can contribute to whether or not the subject is lucky to hold a true belief in P at t . For instance, assume that a subject's epistemic capacities are such that there is always a high objective chance that when next employed, they will give rise to a false belief. Then, even if the subject in fact ends up forming a true belief, she will lack knowledge, for she is still lucky to have formed a true belief. Similarly, if a fair lottery draw has already taken place, a subject who has no evidence about its outcome but truly believes on the basis of wishful thinking that she has won doesn't know that she won, for there *was* a high chance that her ticket would lose.

Hence, it seems clear that chances of error earlier (or perhaps even later) than the time when a subject holds a belief in a proposition P can threaten the subject's knowledge of P . Being lucky doesn't pass as easily as chances of error do.²³ So here is a more general way of stating the high chance intuition: if at t a subject's belief in P is not true by luck, then there is no relevant time t' at which there is a high chance that the subject holds a relevantly similar false belief.

The problem for SPC is the following. When a subject believes a proposition Q solely based on competent deduction from P , how risky her belief in Q is, or how much danger of error it involves, will be a function of both the chance of her belief in P being false and the chance of going wrong in the deduction (a deductive risk). But this will allow for cases in which a competent deduction brings in the crucial bit of danger that disqualifies a subject's belief from constituting knowledge. Think, for instance, of the

²³ This seems to be true generally. Assume that there is a very high chance that I fall climbing a certain route. Even after I have successfully completed the climb and there is no longer any chance of falling, it is still true that I am now lucky that I didn't fall.

time just before the subject's belief in P was formed. At that time there might have been a high chance that if the subject forms a belief in some proposition X and then infers from it another proposition Y , her belief in Y will be false.

Again, let 'risk' be a placeholder for something that knowledge is tolerant to only in small quantities. Any view that allows beliefs involving some non-zero risk of being false, while setting a limit to how high this risk can maximally be, will struggle with SPC. Beliefs constituting knowledge can have some minor risk of being false. And knowledge can be acquired in ways that involve a sufficiently small element of risk. But then, the risk involved in a subject's belief in P and the risk involved in her deduction of Q from P can pile up so that her belief in Q will be too risky.

In the next section I construct a more detailed case demonstrating a failure of SPC for safety theories on which knowledge requires avoiding error in most close worlds. But the structure of counterexamples to SPC should now be clear.

There are other similar considerations against MPC exploiting accumulation of risks. A forceful problem is created by Preface-type cases such as the following. A subject has ample evidence in support of each proposition in a set of propositions S , but also has strong evidence, say of an inductive kind, for believing that not all propositions in S are true. In so far as knowledge can be defeated, this evidence, it seems, acts as a defeater, preventing the subject from knowing the conjunction of every proposition in S . But unless widespread scepticism is true, it doesn't appear to defeat the subject's knowledge of each individual proposition in S .²⁴ This looks like a counterexample to

²⁴ Cf. Hawthorne (2004: 47–50).

MPC, since the subject cannot come to know the conjunction of every proposition in S by competent deduction from propositions she knows.

But an analogous problem arises for SPC. Assume that a subject knows a proposition P , and that she has come to believe Q by competent deduction from P . However, the subject also has good evidence for thinking that she did not competently deduce Q from P , and hence, that her belief in Q doesn't count as knowledge. Alternatively, she might have evidence for thinking that the risk involved in her belief in P and the risk involved in the deduction add up so that her belief in Q is too risky to count as knowledge. Friends of knowledge defeat would say that such beliefs act as defeaters preventing the subject from knowing Q .²⁵ But because they don't prevent the subject from knowing P , nor her deduction from in fact having been competent, we seem to have a counterexample to SPC.

5.

Any theory of knowledge faces the challenge of giving some account of what competent deduction is. And I think we can construct a workable idea of what competent deduction would involve at least within the context of externalist theories that propose either reliabilist or safety requirements for knowledge. I will focus on these frameworks not only because they provide rather plausible candidates for views on which it might look as though MPC fails but SPC doesn't, but because they have resources for saying something

²⁵ The two cases involve different types of defeat: in the MPC case, the subject has evidence for thinking that the proposition she has deduced is *false*, whereas in the SPC case she merely has evidence for thinking that the proposition she has deduced doesn't count as knowledge, since she didn't come to believe it by competent deduction from P . This is an example of overriding defeat, whereas the latter is an example of undermining defeat.

substantial about competent deduction, and the types of stories about competence that they tell are independently plausible.

On a reliabilist theory of knowledge, the competence of a deduction will depend on the reliability of the subject's deductive process; on a safety theory, it will depend on the safety of the subject's deduction. Because the competence of a deduction doesn't consist in bearing a suitable epistemic relation to any proposition entertained by the subject, competence cannot be cashed out in terms of a process producing a high ratio of true beliefs, or in terms of the truth of some belief across all or most close worlds. After all, a subject whose beliefs are massively false can be competent at deducing, or inferring, propositions from other propositions. Rather than truth, we should look at truth-preservation in a high ratio of cases, or across all or most close possible worlds.

So the suggestion is that on reliabilist views, a subject's deduction of a proposition Q from a proposition P is competent just in case in a high ratio of cases in which the subject infers some relevant proposition Y from some relevant proposition X , a conditional with X as its antecedent and Y as its consequent is true (and, of course, P entails Q). And on safety views, a subject's deduction of Q from P is competent just in case in all/most close worlds in which the subject infers some relevant proposition Y from some relevant proposition X , a conditional with X as its antecedent and Y as its consequent is true (and P entails Q). The conditional in question could either be a logical entailment or a material implication. My point about failures of closure doesn't rest on this choice. One reason for opting for a material implication is that it would enable us to give a unified account of the competence involved in logical deduction and inference more generally.

I will now construct a concrete case demonstrating how SPC fails on a *most worlds safety* view of knowledge. But similar cases could also be constructed for reliabilist views, since a subject's belief in P might be reliably formed, and her deduction of Q from P might be reliable, but because the reliability of Q depends both on how reliable her belief in P is, and on the reliability of the deduction, her belief in Q might fail to be reliably formed. A reliability condition on knowledge won't guarantee the truth of SPC.²⁶

According to the most worlds safety view I will discuss, knowledge requires avoiding relevantly similar false belief in most close worlds or cases. For simplicity, let us assume that 'most' is read as 'at least 90%'. The exact number won't matter. Given the rough and ready criterion for competent deduction given above, on such a view a subject competently deduces a proposition Q from a proposition P just in case her actual deduction is successful, and in most (at least 90% of) close worlds in which the subject infers from a relevant proposition X some relevant proposition Y , it is true that if X , Y (or, on the stronger view, that X logically entails Y). Which propositions ' X ' and ' Y ' range over depends on what the relevant initial conditions that must obtain in all close worlds are, and the details will vary from case to case. For instance, if the initial conditions of a subject's deduction of Q from P have built into them the fact that she deduces a proposition from P (and not some other proposition), then only worlds in which she deduces some proposition from P count as close. And very often, such a restriction is very plausible. Moreover, to avoid trivialising competence, normally ' Y ' will range over propositions other than just Q and $\sim Q$. For despite the fact that a subject infers Q from P

²⁶ Against, for instance, Brueckner (1998).

by a deeply defective process, there might be no close world in which she infers $\sim Q$ from P .

As an example, take the following case. I am applying to join a basketball team, and I am asked which of the following half-open intervals my height is in: [150 cm, 151 cm) or [151 cm, 152 cm) or or [219 cm, 220 cm). I know that I am approximately 175.9 cm tall²⁷, and deduce from this the proposition that my height is in the interval [175 cm, 176 cm). Given that I could not easily have held some other belief about my height, the fact that I use the proposition that I am approximately 175.9 cm tall as the starting point of my deduction is built into the initial conditions of the case. Moreover, the multiple-choice nature of the problem I need to solve restricts the relevant propositions that I come to deduce to those stating which half-open interval my height is in.

So here is the counterexample to SPC.²⁸ Assume that it is now t_1 , and I have evidence for believing that I am approximately 175.9 cm tall, and that I will still be approximately 175.9 cm tall a week from now. So at t_1 , I believe the true proposition P that at a time exactly one week later (time t_3) I will be approximately 175.9 cm tall. Let Q be the proposition that at t_3 my height will fall within the interval [175 cm, 176 cm). Hence,

P : At t_3 I will be approximately 175.9 cm tall.

Q : At t_3 my height will fall within the interval [175 cm, 176 cm).

At a time t_2 slightly later than t_1 (let us say some seconds later) I come to form a belief in Q solely based on competent deduction from P . When I perform the deduction, my

²⁷ By 'being approximately 175.9 cm tall' I mean: my height is closer to 175.90 cm than it is to either 175.80 cm or 176.00 cm.

²⁸ Those already convinced by my argument might want to skip this example and move to section 6.

situation is the following. I am asked which interval my height will fall within at t_3 : [150 cm, 151 cm) or [151 cm, 152 cm) or ... or [219 cm, 220 cm). To decide the question, I deduce Q from P . Assume further that given the way in which my deductive capacities function, just before I deduced Q from P , there was a 10% chance that I would infer from P something not logically entailed by it. Assume that this also means that in 10% of close worlds I do infer from P some proposition not logically entailed by it. Note that whichever of the two criteria for competence given above one selects, my deduction is still competent. In 90 % of close worlds I infer from P a proposition entailed by it (i.e., Q), and in over 90 % of close worlds I infer from P some proposition Y such that the material conditional ($P \supset Y$) is true.

Assume also that my belief in P is true in exactly 90 % of worlds that are close at t_1 , and it continues being true in exactly 90 % of close worlds at any time within the closed interval $[t_1, t_2]$. Here is a possible story of why this is so. Just before t_1 I was given a special hormone. Throughout $[t_1, t_2]$, the hormone has a chance of 0.1 of making me grow at least 1 cm at some time between 2 – 3 days of when it was taken. Throughout $[t_1, t_2]$, the chance that the hormone won't take effect and that P is false is 0. Further, whether or not the hormone takes effect some days after I deduce Q from P is independent of how well I do on the deduction, and vice versa. Because P is actually true and true in 90 % of close worlds at any time within $[t_1, t_2]$, by most worlds safety my belief in P is safe throughout $[t_1, t_2]$. The most world safety theorist might deny that safe belief is a sufficient condition for knowledge, but whatever other conditions for knowledge I might have to fulfil, safety will presumably be the only modal condition. If, for instance, I must

have good internally accessible reasons for my belief, we can assume that I do have good internally accessible reasons for believing P .

Given that at t_2 I satisfy the antecedent of SPC, I ought to know Q at t_2 . However, by most worlds safety, my belief in Q is not safe and hence, I fail to know Q . This assumes that of the worlds that are close at t_2 , 10 % are ones in which I did poorly on the deduction, inferring from P some proposition not entailed by it. But this seems like a very plausible assumption to make, given the points above about earlier risks affecting the epistemic status of a subject's belief.

To see that closure fails, split all close worlds at t_2 into four sets: 1) worlds in which my belief in P is true and I inferred from P a relevant proposition Y entailed by P – this proposition will have to be Q , since Q is the only proposition in the relevant class entailed by P , 2) worlds in which my belief in P is true but I inferred from P a relevant proposition Y not entailed by P , 3) worlds in which my belief in P is false and I inferred from P a relevant proposition entailed by P (i.e., Q), and 4) worlds in which my belief in P is false and I inferred from P a relevant proposition Y not entailed by P . Assuming that how well I do on the deduction – whether I manage to infer a proposition really entailed by P – is independent of what my height will be at t_3 , i.e., the truth of P , the percentages of worlds that belong to these sets out of the set of all close worlds are as follows:

	Truly believes P	Falsely believes P
Inference good	1) 81 % of all close worlds.	3) 9 % of all close worlds.
Inference bad	2) 9 % of all close worlds	4) 1 % of all close worlds.

In worlds belonging to set 1), I hold a true belief in Q . In worlds belonging to set 2), I hold a false belief, since P is true but I believe a proposition Y within the relevant class that is not entailed by P , and Q is the only proposition within the relevant class of propositions that is compatible with P . So at least 9 % of all close world are ones in which I hold a false belief. In worlds within set 3), my belief in P is false, but I inferred from P a proposition Y entailed by P . The only proposition in the relevant set that is entailed by P is Q . But because P is false in these worlds, they must be worlds in which the hormone takes effect, and I will have grown over 1 cm by t_3 . But then, Q will be false, as my height won't fall within the interval [175 cm – 176 cm). In these worlds I hold a false belief. Summing up, even ignoring worlds in class 4), in at least 18 % of close worlds at t_2 I form a false belief about my height at t_3 as the result of a deduction from P . The belief in Q I actually hold at t_2 is not safe.

Before concluding, let me address some objections, which should also make clearer what I take to be the overall significance of my result to be.

6.

I have argued for the following claims:

1. One should either accept both single and MPC or neither. None of the closure-based paradoxes discussed differentiate between these principles.
2. If one wants to accept both (and to avoid widespread scepticism), there is a strong reason for thinking that both need to be revised.

Before concluding, I will look at some objections that question the significance of my argument.

The first objection is the following:

”Your result rides on looking at fallible deducers. This allows you to bring in risks created not only by there being a non-zero chance that the propositions serving as the premises of a deduction are false (*empirical* risks), but also that the deduction itself goes wrong (*deductive* risks). But when we are interested in formulating closure principles, surely we ought to idealise and focus on subjects who are deductively infallible. And as long as we focus on such subjects, your result no longer applies, since the accumulation of risk problems that arise for MPC will no longer arise for SPC.”²⁹

What must be meant here by ‘infallible’ is not just subjects who in fact never go wrong in their attempted deductions, but subjects for whom there is never even any non-zero chance of going wrong. Now it is perfectly true that if we simply stipulate that closure principles are to apply only to subjects who are infallible in this sense, then my paradox for SPC will simply not arise. Unless MPC is also fixed in some way, the accumulation of risk problems will still arise for MPC. If we only focus on infallible deducers, it is not true that all paradoxes generated by MPC are also generated by SPC.

Either subjects with fallible inferential abilities can or cannot acquire knowledge by deduction.³⁰ Saying that they cannot would seem no more plausible than saying that subjects such as ourselves with fallible perceptual abilities are unable to acquire perceptual knowledge of the world. And given that fallible subjects *can* acquire knowledge by deduction, we should expect to be able to say something about the conditions under which this is possible. Our interest in formulating closure principles for knowledge is not that different from our interest in the conditions under which fallible subjects can acquire and retain knowledge by other means such as perception. As an analogy, to solve sceptical puzzles created by the fact that our perceptual abilities are

²⁹ Thanks to Stewart Cohen and Dorothy Edgington for bringing this line of objection to my attention.

³⁰ Here I am setting aside reasons for questioning this that arise from individuating propositions in a very coarse-grained manner.

fallible by giving an account of perception that applies only to perceptually infallible subjects would be utterly implausible.³¹

Here is a related objection.

“What you have shown is that both empirical and deductive risks can enter into deductions. But the following asymmetry between SPC and MPC still remains. To create trouble for MPC, only empirical risks need to be exploited. Your paradox for SPC exploits both types of risks, since it relies on the fact that empirical and deductive risks can accrue even across single premise deductions. Dealing with your paradox by ruling out deductive risk (by focusing on infallible deducers) still leaves us with accrual of risk problem for MPC. But ruling out empirical risks by stipulating that closure is only to apply when the premises subjects deduce from are certain gets rid of both this worry for MPC and your worry.”³²

Even if this asymmetry obtained, as long as our interest in closure principles is an interest in the conditions under which knowledge can be acquired by deduction in the face of both empirical and deductive risks, it is not clear what significance it would have. But in fact, no such asymmetry exists. To give a simple, clear-cut counterexample to SPC, I only focused on a single deduction, taking the risk involved in the premise of the deduction to be empirical. But by considering long chains of deductions, a similar counterexample can be constructed in which the risk involved in the premise of the deduction is a deductive risk inherited from a previous deduction. Hence, an accumulation of risks problem can be raised for SPC by relying wholly on deductive risks.

Assume that I engage in a long chain of single premise deductions, but because I am a fallible deducer, at each inferential step there is an equal, very slight chance that my

³¹ Note that this fix would not solve the paradox for closure arising from Preface-type cases. For even if there is in fact no chance that a subject's attempted deduction go wrong, the subject might have strong reasons for thinking that there is.

³² Thanks to Stewart Cohen and Tim Williamson for prompting me think about this issue.

attempted deduction will fail. Assume that across any n deductions, there is a high chance that I will end up believing a falsehood. I begin with a mathematical truth P_0 that I know and that (obviously) has no chance of being false, and through n steps arrive at what is in fact a mathematical truth P_n . Despite the odds, I perform competently throughout each step of the chain. By applying SPC to each step, I ought to now know P_n . But if the sort of connection discussed above between objective chance of error and a problematic variety of epistemic luck exists, then it seems wrong to say that I come to know P_n . After all, there was a high chance that by the time I actually came to believe P_n , I would have come to believe a falsehood. This puzzle arises even if we idealise away from empirical risks. Hence, ruling out empirical risks is not sufficient to rid us of all paradoxes created by closure.

The final issue I will briefly consider is the following:

“A large part of our interest in closure is related is to the paradoxes it generates. What significance does your result have for solving these paradoxes?”³³

One reaction to my argument is to take closure to fail across the board. If closure simply failed, then there would be no closure-based paradoxes to solve. Another reaction is to fix single premise closure so as to avoid the type of paradox that I raise.

What seems to me to be the most obvious fix is to simply specify that the combined risk involved in the premise or premises of the deduction (whether empirical or deductive) and that of the deduction itself is sufficiently small. This would be to admit that competent deduction from a known proposition is not guaranteed to yield new knowledge. In long chains of deductions such as the case described above, at some point

³³ Thanks to Stewart Cohen.

in the chain deductive risks inherited from earlier deductions will pile up so as to prevent the subject from coming to know a proposition P_i that she has competently deduced from a known proposition P_{i-1} .

Note that unlike revisions that stipulate that closure is not to apply when the deduced proposition is sceptical, the suggested revision is compatible with the intuition that whether or not a subject comes to know a proposition Q by deduction from a proposition P only depends on the competence of the deduction and the subject's epistemic standing with respect to P . It's just that now 'epistemic standing' is construed as depending on more than whether or not the subject knows P . This revision to closure would have the result of getting rid of paradoxes that exploit accrument of risks. Hence, it would at least solve what I have argued is the most serious paradox that has been thought to arise exclusively for MPC, as well as the paradox as applied to single premise cases, including the chain-case. This would at least be some progress.

Conclusions

I presented what I take to be the strongest argument against multi premise closure. This argument shows that at least within indeterministic worlds, retaining multi premise closure will require giving up some seemingly very intuitive claim about knowledge and its connections with chance. I then argued that the considerations in favour of the three premises of the argument also provide trouble for single premise closure. In particular, risks can pile up not only synchronically, as in multi premise deductions, but also

diachronically, when a subject extends her knowledge by competent deduction from just one premise. Competent deduction is susceptible both to empirical and deductive risks.³⁴

Bibliography

Brueckner, A. (1998). Deductive Closure Principles. (In E. Craig (Ed.), Routledge Encyclopedia of Philosophy. London: Routledge.) Retrieved October 21, 2006, from <http://www.rep.routledge.com/article/P011>.

Carroll, L. (1895). What the Tortoise said to Achilles. *Mind*, 4:14, 278–280

David, M. and Warfield, T. (forthcoming A). Closure Principles and Skeptical Arguments. (In Q. Smith (Ed.), *Epistemology: New Essays*. Oxford: Oxford University Press.)

David, M. and Warfield, T. (forthcoming B). Six Possible Counterexamples to One or Two Epistemic Closure Principles.

DeRose, K. (1999). Introduction: Responding to Scepticism. (K. DeRose, & T. A. Warfield (Eds.), *Scepticism – A Contemporary Reader* (pp. 1 – 24). Oxford: Oxford University Press.)

Dretske, F. (1970). Epistemic Operators. *Journal of Philosophy*, 67, 1007–1023.

Dretske, F. (2005). The Case Against Closure. (In M. Steup & E. Sosa (Eds.), *Contemporary Debates in Epistemology* (pp. 13–26). Oxford: Blackwell.)

Hawthorne, J. (2004). *Knowledge and Lotteries*. (Oxford: Oxford University Press)

Hawthorne, J. & Lasonen-Aarnio, M. (*forthcoming*). Knowledge and Objective Chance. (In P. Greenough & D. Pritchard, D. (Eds.), *Williamson on Knowledge*. Oxford: Oxford University Press.)

Hintikka, J. (1970). Knowledge, Belief, and Logical Consequence. *Ajatus*, 32, 32–47.

Kyburg, H. (1970). Conjunctivitis. (In M. Swain (Ed.), *Induction, Acceptance, and Rational Belief* (pp. 55–82.). Dordrecht: Reidel.)

Lewis, D. (1996) Elusive Knowledge. *Australasian Journal of Philosophy*, 74: 4, 549–567

Makinson, D. C. (1965). The Paradox of the Preface. *Analysis*, 25, 205–207

³⁴ I am grateful to discussions and exchanges with Ville Aarnio, Frank Arntzenius, Stewart Cohen, Dorothy Edgington, John Hawthorne, Daniel Morgan, and above all Tim Williamson.

- Nozick, R. (1981). *Philosophical Explanations*. (Oxford: Oxford University Press)
- Pritchard, D. (2005). *Epistemic Luck*. (Oxford: Oxford University Press)
- Stalnaker, R. (1999). *Context and Content*. (Oxford: Oxford University Press)
- Unger, P. (1968). An Analysis of Factual Knowledge. *Journal of Philosophy*, 65, 157–170
- Vogel, J. (1990). Are there Counterexamples to the Closure Principle?. (In M. Roth & G. Ross (Eds.), *Doubting: Contemporary Perspectives on Scepticism* (pp. 13–27). Dordrecht: Kluwer Academic Publishers.)
- Williamson, T. (2000). *Knowledge and its Limits*. (Oxford: Oxford University Press)