

Rational Self-Doubt and the Failure of Closure*

Joshua Schechter
Brown University

Abstract Closure for justification is the claim that thinkers are justified in believing the logical consequences of their justified beliefs, at least when those consequences are competently deduced. Many have found this principle to be very plausible. Even more attractive is the special case of Closure known as Single-Premise Closure. In this paper, I present a challenge to Single-Premise Closure. The challenge is based on the phenomenon of rational self-doubt – it can be rational to be less than fully confident in one's beliefs and patterns of reasoning. In rough outline, the argument is as follows: Consider a thinker who deduces a conclusion from a justified initial premise via an incredibly long sequence of simple competent deductions. Surely, such a thinker should suspect that he has made a mistake somewhere. And surely, given this, he should not believe the conclusion of the deduction even though he has a justified belief in the initial premise.

1 Introduction

There seems to be a special relationship between logic and reasoning. Hartry Field has recently argued that the standard ways of defining the logical consequence relation don't work, and so the only way to characterize logic is to somehow tie it to reasoning.¹ Even if that isn't right, there is reason to think that there is a close connection between logic and reasoning. We care about logic. Logicians do not merely study some abstract mathematical structure on a par with very many others. Rather, logical consequence is an interesting and important relation. This is a fact that needs to be explained. The natural suggestion is to appeal to a connection between logic and good reasoning: Logic is important because it is intimately tied to rationality.

* Forthcoming in *Philosophical Studies*. The final publication is available at <http://www.springerlink.com/content/b86081k701235219/>. doi:10.1007/s11098-011-9823-1.

¹ See Field (2009b).

What is this connection between logic and rationality? There are two proposals that are initially attractive:

(Naïve Closure) Thinkers are justified in believing the logical consequences of their justified beliefs.

(Naïve Coherence) Thinkers are unjustified in having a logically inconsistent set of beliefs.

These two claims are much too crude to be accepted as they stand. There are familiar counterexamples to these principles.² To be plausible, they must be significantly refined. But there are strong intuitions that suggest we ought to accept something in the ballpark of each of them.

In this paper, I focus on the closure principle for justification. This principle articulates something like the thought that logic – or better, deductive inference – is fully epistemically secure. There is no loss of justification when a thinker engages in a competent deductive inference.

This paper presents a general challenge to closure for justification. The challenge is based on the phenomenon of rational self-doubt – it can be rational to be less than fully confident in one's beliefs and patterns of reasoning.³ In rough outline, the argument is as follows: Consider a thinker who deduces a conclusion from a justified initial premise via an incredibly long sequence of simple competent deductions. Surely, such a thinker should suspect that he has made a mistake somewhere. And surely, given this, he should not believe the conclusion of the deduction even though he has a justified belief in the initial premise. This provides a counterexample to closure for justification.⁴

² See Harman (1986; 1995).

³ I believe that the term “rational self-doubt” is due to Christensen. See, for example, Christensen (2008).

⁴ See Lasonen-Aarnio (2008) for a related objection to closure principles for knowledge.

There are two familiar challenges facing the closure principle for justification – the lottery paradox and the paradox of the preface.⁵ The challenge presented here differs from those two arguments in an important respect. Those two arguments are often used to motivate a broadly Bayesian approach to rational belief. On such a view, while a general closure principle – applying to deductions with any number of premises – is false, a closure principle restricted to single-premise deductions turns out to be true. Thus, the lottery and preface paradoxes simultaneously undermine multi-premise closure and support single-premise closure. By contrast, if my argument works, it works even against closure principles restricted to single-premise deductions. Unlike the lottery and preface paradoxes, it cannot easily be accommodated within a Bayesian framework for rational belief.

The main goals of this paper are to develop the strongest version of this challenge and to evaluate possible ways of responding to it. My conclusion is *not* that we should reject every closure-like principle. Indeed, I suspect that some closure-like principle is correct. However, I take it that my argument demonstrates two important facts. First, any correct closure-like principle must be rather complicated. There is no simple and precise statement of the relationship between deductive inference and justification. Second, there is no such thing as a fully epistemically secure deductive inference. There is no correct (non-trivial) principle stating that if certain conditions obtain, a deductive inference will be guaranteed to preserve justification.

Before proceeding any further, it may be helpful to state two background assumptions that I will be making in this paper. First, I assume that there is such a thing

⁵ See Kyburg (1970) and Makinson (1965), respectively. There are also the familiar objections to closure for knowledge put forward by Dretske (1970) and Nozick (1981) in their discussions of skepticism. So far as I can tell, the issues raised in this paper have no direct connection with the familiar skeptical challenges.

as deductive inference. Some philosophers have denied this claim.⁶ But it is psychologically plausible that we sometimes reason deductively.⁷ Any reader who rejects this assumption is invited to understand the main point of this paper as a conditional: If there is such a thing as deductive inference, it is not fully epistemically secure.

Second, I assume that deduction involves the employment of rules of inference. It is plausible that reasoning is in general rule-governed.⁸ In the case of deduction, the relevant rules are deductive rules of inference. These are rules that are intimately tied to the logical concepts. Most of my discussion could be reframed so as not to rely on the assumption that deduction is rule-governed.⁹ But making use of this assumption provides a convenient way to talk.

This paper will proceed as follows: In the next section, I present and motivate closure principles for justification. In section three, I develop the “long sequence” argument against such principles. In section four, I turn to three objections to this argument. The first – inspired by Timothy Williamson’s discussion of closure for knowledge – is that the argument depends on a confusion of levels. Rational self-doubt can defeat our justification for believing that we are justified but cannot defeat the justification of our ground-level beliefs. The second objection – inspired by David

⁶ See, for example, Harman (1986).

⁷ Most cognitive psychologists seem to agree that there is a distinctively deductive kind of reasoning. The main debate concerns exactly how it is to be characterized. See Evans, Newstead, and Byrne (1993) for discussion of the major views. Some psychologists, including Cheng and Holyoak (1985) and Cosmides (1989), suggest that we do not employ topic-neutral rules of inference, but only domain-specific reasoning mechanisms. This view faces several difficulties. But even if it is correct, there remains the question of whether justification is closed under any domain-specific patterns of inference.

⁸ Employing a rule of inference should not be taken to require reflective appreciation of the rule. There are familiar obscurities in the notion of following a rule. See Kripke (1982) for the classic discussion and Boghossian (1989) for an overview of the resulting debate. See Boghossian (2008) for discussion of the particular difficulties facing a rule-based picture of reasoning. However, I am not aware of any attractive alternative picture of reasoning.

⁹ For my purposes here, all that I really need is that there is a competence/performance distinction for deductive reasoning. Appealing to rules helps to explicate this distinction – thinkers may employ incorrect rules or they may misapply correct ones.

Lewis's discussion of immodest inductive methods – is based on the claim that we cannot rationally believe that any of our basic rules of inference is not to be trusted. There is simply no room for rational self-doubt about basic deductive rules. The third objection is based on an overgeneralization worry. More specifically, the worry is that the long sequence argument can be generalized to yield the absurd conclusion that we should not simply obey any of our belief-forming methods, but rather should hedge all of our reasoning. Finally, in section five, I briefly discuss why we might have found strong closure principles intuitively so appealing.

2 Closure Principles

There is a very natural line of thought motivating closure. This line of thought is nicely articulated in Timothy Williamson's *Knowledge and its Limits*. According to Williamson, the intuitive idea behind closure is that “deduction is a way of extending one's knowledge.”¹⁰ By this, Williamson does not mean that deductive inference is *sometimes* a way of extending one's knowledge. Rather, deductive inference is *always* a way of extending one's knowledge.

Of course, a thinker may already have known the conclusion of a deduction. So strictly speaking, deduction does not always *extend* knowledge. But the intuitive thought is that when a thinker competently deduces a claim from known propositions, the result is always a known proposition.

¹⁰ Williamson (2000), page 117. There are alternative motivations for closure principles. For instance, one could argue that what one's total evidence supports is always closed under (single-premise) logical entailment and then argue for a tight connection between justification and evidential support. This is a less intuitive and more theory-driven motivation than Williamson's.

Like many discussions of closure in the literature, Williamson's discussion focuses on knowledge. But it seems to me that the underlying thought doesn't directly concern knowledge. There is a more basic idea: Deduction is always a legitimate way to extend one's beliefs.

This thesis is most naturally understood as a claim about diachronic rather than synchronic rationality. That is, it is a claim about rational change of belief and not about belief at a time. The thesis is also most naturally understood as a claim about permission rather than obligation. What this suggests is that a more basic closure principle concerns epistemic justification rather than knowledge: When a thinker competently deduces a claim from justified beliefs, the result is always a justified belief.¹¹

What do I mean by "justification" here? One of the lessons of the internalism/externalism debate in epistemology is that there may be several important concepts of epistemic justification. The notion that seems most relevant to the intuitions underlying closure is one that is closely tied to a notion of epistemic responsibility. The central intuition supporting closure is that deduction is a *responsible* belief-forming method. Thinkers are epistemically responsible in believing what they deductively infer from epistemically responsible beliefs.¹²

¹¹ Notice that coherence principles must be stated for justification rather than knowledge. It is trivial that thinkers cannot know every one of a logically inconsistent set of propositions. This provides evidence that the more basic principles linking deductive inference with rationality concern justification rather than knowledge.

¹² Some clarifications about the notion of epistemic responsibility may be helpful here. First, having a responsible belief does not require that the relevant inquiry was carried out in a fully responsible manner. A thinker may have been irresponsible in, for example, not sufficiently gathering evidence and nevertheless count as responsible in forming a belief given the evidence at hand. Second, responsibility should not be identified with blamelessness. In a strict sense of "blame", we do not typically blame thinkers – that is, have Strawsonian reactive attitudes – for their beliefs. While there may be an extended sense of blame on which we do blame thinkers for their beliefs, so far as I can tell we do not do so in any systematic way. Moreover, on this extended sense, thinkers can count as epistemically irresponsible but blameless if they have an appropriate excuse.

There is some reason to think that there is a second closure principle governing epistemic obligations (rather than permissions). In particular, there may be an obligation to believe if the deduction is sufficiently simple and the relevant issue is highly salient or otherwise important to the thinker. However, I suspect that such a principle should not be thought of as an independent closure principle. Rather, it should be seen as following from closure for justification and a general principle governing when thinkers ought to take a stand on an issue. Roughly put, when the inference is sufficiently simple and the relevant issue is highly salient or otherwise important, thinkers epistemically ought to take a stand. In such circumstances, permissible beliefs become required beliefs.

The closure principle for justification is *not* the fundamental normative principle motivated by the Williamsonian line of thought. There is a still more basic epistemic principle. Consider a thinker who forms a belief on the basis of drawing a competent deduction from a confidently held but unjustified belief. Intuitively, the thinker is in an epistemically problematic state – he has an unjustified belief. But he has not done anything wrong in drawing the inference. He has only made *one* mistake in his reasoning. Similarly, consider a thinker who fails to deduce a conclusion concerning a highly salient or otherwise important issue via a simple deduction from one of his confidently held but unjustified beliefs. Such a thinker intuitively has made *two* mistakes – he has an unjustified belief and he has failed to infer a conclusion on a salient or important matter.

What this suggests is that the fundamental closure principle does not concern the preservation of justification. Rather it concerns epistemic error: Roughly put, thinkers do not make any (additional) epistemic errors in competently deducing from their beliefs. Nevertheless, in what follows, I'll focus on the closure principle for justification. This is

a familiar kind of principle. It does not raise difficult questions about the relationship between epistemic error and more familiar epistemic statuses such as justification and knowledge.

For concreteness, we can carefully state a closure principle for justification as follows:

(Closure) Necessarily, if S has justified beliefs in some propositions and comes to believe that q solely on the basis of competently deducing it from those propositions, while retaining justified beliefs in the propositions throughout the deduction, then S has a justified belief that q.

This principle is analogous to the closure principle for knowledge presented in John Hawthorne's *Knowledge and Lotteries*.¹³ The principle is somewhat complicated in order to avoid certain issues that needn't detain us here.¹⁴ But the core idea is straightforward: If a thinker has some justified beliefs and competently deduces a belief from them, then the resulting belief is also justified.¹⁵

3 Closure and Rational Self-Doubt

3.1 Problems for Closure

There are two well-known problems facing Closure – the lottery paradox and the preface paradox. The lottery paradox can be stated as follows:

Suppose that there is a raffle with 1,000 tickets. One ticket, chosen at random, will win. Suppose that I am aware of this. Suppose that I have a lot of time on my hands, and to fill my time I form beliefs about each of the tickets. In particular, for each ticket, I form the belief that it will lose on the grounds that it has a

¹³ See Hawthorne (2004), page 33.

¹⁴ For instance, the “retaining justified beliefs” clause is present because thinkers can lose justification for believing the premises of a deduction once they notice that an implausible conclusion follows from them.

¹⁵ Closure should be distinguished from transmission, as introduced in Wright (1985). It is compatible with Closure that in certain cases a thinker cannot acquire additional justification for a belief on the basis of competently deducing it from justified premises. For example, having justified beliefs in the premises may require antecedently possessing justification for believing the conclusion.

999/1000 chance of losing. Presumably I am justified in having each of these beliefs. But if I were to infer from them that all of the tickets will lose, I would not be justified in this new belief. That is because I am aware that some ticket will win.

In this scenario, I have justified beliefs in many propositions but am not justified in believing their conjunction.¹⁶ Presumably, reasoning by conjunction introduction counts as competent deduction. Presumably, too, I do not lose my justification for believing the individual conjuncts (or for employing conjunction introduction) if I infer the conjunction. So we have a counterexample to Closure.

I find this counterexample convincing. Nevertheless, there are philosophers who have defended the view that I am not justified in believing of each ticket that it will lose.¹⁷ I think this view is difficult to maintain; it requires adopting a kind of skepticism about merely statistical grounds for belief.¹⁸ But in any case, this issue can be sidestepped by focusing on the second familiar problem for Closure – the paradox of the preface. A version of this paradox can be stated as follows:

There are very many propositions that I justifiably believe. Such propositions include simple claims of mathematics and logic; claims about my self, my environment, and my past experiences; and so on. Consider a heterogeneous conjunction of many of these claims. I am not justified in believing this conjunction upon deducing it from my individual beliefs. That is because I am aware that my beliefs are sometimes – perhaps very rarely – wrong. So I should think it likely that the conjunction is false.

In this scenario, I again have justified beliefs in many propositions but am not justified in believing their conjunction. Presumably, reasoning by conjunction introduction counts as competent deduction. Presumably, too, I do not lose my justification for believing the

¹⁶ Strictly speaking, there is a distinction between the proposition that all the tickets will lose and the conjunction of the 1000 conjuncts. But since I may be (nearly) certain of their equivalence, this cannot be used to avoid the counterexample.

¹⁷ For some examples, see Pollock (1983), Evidentiary (1999), and Nelkin (2000).

¹⁸ See Vogel (1990) for helpful discussion.

individual conjuncts (or for employing conjunction introduction) when I infer the conjunction. So we have a second counterexample to Closure.¹⁹

Notice that there is no prospect of handling this counterexample by suggesting that I am not justified in believing the initial conjuncts. The conjuncts were chosen precisely to be a heterogeneous collection of my justified beliefs.²⁰

The lottery and preface paradoxes have a familiar diagnosis: Having a justified belief is compatible with there being a small risk that the belief is false. Having a justified belief is incompatible with there being a large risk that the belief is false. Risk can aggregate over deductive inferences. In particular, risk can aggregate over conjunction introduction.

There are two kinds of risk that ought to be distinguished here. First, there is the objective chance that my belief is false. Second, there is my rational degree of confidence that my belief is false. These two risks are often correlated: As the number of tickets in a raffle increases, the objective chance that a specific ticket will lose and my rational degree of confidence that it will lose both increase (at least assuming that I am aware of the relevant facts). However, the sort of risk at issue in the counterexamples to Closure is the latter kind. Even if the raffle has already taken place or involves a pseudo-random deterministic mechanism, the counterexample will stand. Similarly, even if I am an extremely reliable thinker and there is a negligible objective chance that the

¹⁹ As I've stated the lottery and preface paradoxes, they involve subjects with a bit more cognitive resources – more computational power, better short-term memories, etc. – than we actually possess. This small amount of idealizing does not undermine the use of the counterexamples. Our intuitions about such subjects are as strong as our intuitions about ordinary reasoners. Moreover, there are related cases that don't require even this small amount of idealization. See Christensen (2004), chapter 3.

²⁰ Some philosophers have defended the view that I am justified in believing the conjunction despite also having a justified belief that the conjunction is likely false. This view strikes me as deeply unintuitive.

heterogeneous conjunction is false, so long as I have a high rational degree of confidence that it is false, we will have a counterexample to Closure.

This diagnosis goes some way toward motivating a broadly Bayesian approach to justified belief. On such a view, a thinker's degrees of confidence rationally ought to conform to the probability calculus. Having a justified belief is compatible with being rational in having a large degree of confidence in the relevant claim. Having a justified belief is incompatible with being rational in having only a very small degree of confidence in the claim.

3.2 Single-Premise Closure

On a broadly Bayesian view, Closure is false. However, a weaker version of this principle is true. According to orthodox probability theory, the probability of the conclusion of a valid inference is at least as great as the probability of the conjunction of the premises. For the special case of a valid inference with only a single premise, the probability of the conclusion is at least as great as the probability of the premise. Given a Bayesian picture, this suggests that a thinker has (propositional) justification for believing the conclusion of a valid single-premise inference whenever the thinker has (propositional) justification for believing the premise. Moreover, if the thinker comes to believe the conclusion on the basis of competently deducing it from a (doxastically) justified belief in the premise, the thinker will presumably have a (doxastically) justified belief in the conclusion. Thus, a broadly Bayesian account of justified belief motivates the following weaker version of Closure:

(Single-Premise Closure) Necessarily, 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.²¹

This principle escapes the objections provided by the lottery and preface paradoxes. As we've seen, a natural diagnosis of those counterexamples is that conjunction introduction – and deduction from multiple-premises, more generally – can aggregate risk. But there cannot be any aggregation of risk from the combination of premises if there is only a single premise. If risk conforms to probability theory, the risk of the falsity of the conclusion cannot be any greater than the risk of the falsity of the premise.

Although the lottery and preface paradoxes do not directly provide objections to SPC, there is a related problem with this principle. Before I present this problem, however, it is helpful to consider two warm-up objections.²²

Here is the first warm-up objection: In some cases, I may have misleading evidence that I've applied a rule that is invalid or unjustified. For example, suppose that the world's twenty best logicians each tell me that one of the deductive rules I employ is incorrect. Suppose that I make use of this rule in competently deducing a conclusion from a single justified premise. Presumably, I am unjustified in believing the conclusion of my deduction. This provides a counterexample to SPC.

Here is the second warm-up objection: In some cases, I may have misleading evidence that I've misapplied a rule. For example, suppose that I've been told by a reliable source that the pill I swallowed earlier today has an unusual side effect: It makes thinkers prone to treat certain invalid inferences as instances of some logically valid rule.

²¹ This principle should be generalized to accommodate zero-premise inferences. I'll leave this generalization implicit in what follows.

²² I do not know who first stated these objections. They appear to be part of philosophical folklore.

Suppose that I make use of this rule in competently deducing a conclusion from a single justified premise. Presumably, I am unjustified in believing the conclusion. This provides another counterexample to SPC.

These objections are intuitively compelling. But there is a straightforward fix that handles both of them. Namely, we can add a “no-defeaters” clause to the statement of SPC:

(Single-Premise Closure*) 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.

A thinker has a defeater for the claim that a deduction was competently performed if the thinker believes (or is justified in believing) that she has (or has likely) employed an incorrect rule, misapplied a correct rule, or otherwise made a mistake in the course of the deduction.²³

This modification answers both warm-up objections. In each case, I have a defeater for the claim that the deduction was competent. Moreover, it is not *ad hoc* to add a no-defeaters clause to SPC. One of the lessons of contemporary epistemology is that epistemic principles should generally include such clauses. Thus, if there is a problem for SPC stemming from the two warm-up objections, we should simply move to SPC*.²⁴

²³ It might be suggested that the no-defeaters clause should be built into the definition of competent deduction. However, it is cleaner to keep it distinct. Whether a thinker has made a competent deduction shouldn't depend on her meta-beliefs about her reasoning.

²⁴ Adding a no-defeaters clause may be incompatible with strict forms of Bayesianism. Insofar as it is, this is more of a difficulty with strict forms of Bayesianism than with the no-defeaters clause.

3.3 The Long Sequence Argument

Nevertheless, there is a serious problem for single-premise closure that remains. The problem concerns what we should say about long sequences of deductions. The problem is not, in its essentials, a new one. It is connected to issues in epistemology that were noticed a long time ago, by the early moderns (if not earlier). Indeed, the basic problem is implicit in the following passage from Hume:

In all demonstrative sciences the rules are certain and infallible; but when we apply them, our fallible and uncertain faculties are very apt to depart from them, and fall into error. We must, therefore, in every reasoning form a new judgment, as a check or controul on our first judgment or belief By this means all knowledge degenerates into probability; and this probability is greater or less, according to our experience of the veracity or deceitfulness of our understanding, and according to the simplicity or intricacy of the question. ... There is no Algebraist nor Mathematician so expert in his science, as to place entire confidence in any truth immediately upon his discovery of it, or regard it as anything, but a mere probability.²⁵

This is an excerpt from the start of the section of Hume's *Treatise* entitled "Of Skepticism with regard to Reason." In this excerpt, Hume's point is that, as we are well aware, we are prone to mistakes in our reasoning. Even when we are in fact reasoning correctly, we should take into account the fact that we are fallible. We should not be highly confident in what results even from competent deduction applied to some of our knowledge. Thus, "all knowledge degenerates into probability."

There is a passage in Locke grappling with the very same issues. Indeed, the excerpt from Hume can be read in part as a response to this passage:

It is true, the perception produced by demonstration is also very clear; yet it is often with a great abatement of that evident lustre and full assurance that always accompany that which I call intuitive: like a face reflected by several mirrors one

²⁵ Hume, *A Treatise on Human Nature*, I.IV.I. The main point of Hume's discussion in this section is to provide a (fallacious) argument that the belief in the conclusion of an inference isn't rationally supported by its premises. But the considerations put forward in the excerpt above do not depend on the details of this argument.

to another, where, as long as it retains the similitude and agreement with the object, it produces a knowledge; but it is still, in every successive reflection, with a lessening of that perfect clearness and distinctness which is in the first; till at last, after many removes, it has a great mixture of dimness, and is not at first sight so knowable, especially to weak eyes. Thus it is with knowledge made out by a long train of proof.²⁶

Locke is here defending the claim that one *can* gain knowledge by way of a long train of inferences. However, he is clearly worried about the epistemic status of the resulting knowledge – such knowledge has a “great mixture of dimness”. It is not the best kind of knowledge to have.²⁷

This paper is primarily concerned with justification rather than knowledge. It’s also concerned with single-premise closure rather than closure more generally. But the materials in these passages can be used to provide an objection to SPC and SPC*. From Locke, we can see that there are concerns about the epistemic status of the conclusion of a long sequence of deductions. From Hume, we can see that the principal worry involves our awareness of our own fallibility. Here, then, is the objection:

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.²⁸

²⁶ Locke, *An Essay Concerning Human Understanding*, IV.II.6.

²⁷ There is a puzzle concerning Locke’s view. In Locke, “dimness” is not the absence of light, but the absence of clarity. This is ultimately to be understood in terms of a notion of resemblance with the world. The puzzle is this: If a thinker performs a competent deduction from known premises, the conclusion – no matter how long the inference – should resemble the world just as well as the premises (collectively) do. So why is there any additional dimness in the conclusion? What this suggests is that the real problem with long sequences of deductions fits Hume’s diagnosis. It concerns our awareness of our own fallibility.

²⁸ One might worry that there are not enough interesting single-premise deductions to cause difficulties for SPC. This is not a serious worry. Simple single-premise deductive rules include conjunction elimination, disjunction introduction, double-negation introduction and elimination, as well the following conditional rules: A/if B then A; both A and if A then B/ B; and if A then B/if it is not the case that A then it is not the case that B. There are also rules that allow us to work within embeddings. For instance, the inference from (A and B and (if B then C) and D) to (A and C and D) plausibly counts as a simple single-premise

I should not be very confident in the conclusion of the long deduction. This is because I should think that it's very likely I've made a mistake in there somewhere. For any given inferential step, I shouldn't think that I've made a mistake in that very step. But I should think it likely that I've made a mistake somewhere in the sequence. After all, I have ample evidence that I'm prone to errors in my reasoning.²⁹

This provides a counterexample to SPC.³⁰ 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 SPC 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.³¹

Moving to SPC* doesn't help. Even if I have a defeater for the claim that the long deduction was competent, there is no defeater of the competence of any of the short

deductive inference. These are more than sufficient to allow non-trivial sequences of simple single-premise deductions. Moreover, if we require the long sequence of deductions to have a single (perhaps conjunctive) initial premise, but we allow multi-premise deductions later in the sequence (using the earlier members of the sequence as premises), we will still have a counterexample to Bayesian views.

²⁹ And even if I don't have any positive evidence for the claim that I'm prone to errors in my reasoning, I should presumably take into account the possibility that I've made a mistake.

³⁰ *Mutatis mutandis*, this also provides a counterexample to single-premise closure for knowledge.

³¹ One might worry that the long chain of deductions is a sorites argument and therefore my conclusion that SPC is false is a hostage to the correct treatment of vagueness. In response, I'd like to make two points. First, the long chain of inferences does not resemble a classical sorites argument in that the major premise, SPC, is not primarily motivated by considerations having to do with vagueness or indeterminacy. It is not motivated by some kind of tolerance in the concept of justification. Rather, it is motivated by the thought that deduction is fully epistemically secure. Second, even were the long chain of deductions a sorites argument, the major contemporary solutions to the sorites paradox – supervaluationism, epistemicism, degree theories, and so on – all agree that the major premise in a classical sorites argument is false. Where they disagree is in what they say next. Thanks to Stew Cohen for pressing me on this issue.

simple deductions.³² So there must be a failure of closure for an undefeated short simple single-premise competent deduction.

I find this argument utterly compelling. And there is a natural diagnosis of what's going on: A thinker's rational degree of belief drops ever so slightly with each deductive step.³³ Given enough steps, the thinker's rational degree of belief drops significantly.

To put the point more generally, the core insight is simply this: If deduction is a way of *extending* belief – as the Williamsonian line of thought suggests – then there is some risk in performing any deduction. This risk can aggregate, too.

As before, “risk” here does not stand for the objective chance that I've made some kind of mistake in my reasoning. Rather, it stands for my rational degree of confidence that I've made a mistake.³⁴ What kinds of mistakes should I worry I might have I made in my reasoning? Potential mistakes include the following: (i) I am employing an unreliable rule of inference; (ii) I am employing an unjustified rule of inference; (iii) I have misapplied one of the deductive rules I employ; and (iv) I have incorrectly linked together short deductions. In ordinary situations, the most pressing worries are the latter two. These are the worries that straightforwardly scale with the length of a deduction.

³² Of course, there is a defeater for one of the steps of the deduction in the thin sense that the premise of the deduction is justified and the conclusion is unjustified. However, modifying SPC by adding a clause to rule out this kind of defeater would trivialize the principle. Moreover, this is not a good way to characterize the intuitive notion of a defeater. Roughly put, a step of an argument is defeated only if that step is to blame for the lack of justification for the conclusion. In the long sequence of deductions, none of the individual steps need be defeated in this thicker sense. Thanks to Stew Cohen for helpful discussion of this issue.

³³ Of course, there may be deductive steps at which the thinker's rational degree of belief increases – perhaps, for instance, the inference from *A* to *either A or B*.

³⁴ Lasonen-Aarnio (2008) uses related considerations to argue that multi-premise and single-premise closure for knowledge stand or fall together. Her arguments primarily focus on a safety-based conception of knowledge. But one of her central ideas is similar. Given that (i) knowledge is incompatible with a high objective chance of falsity and (ii) the objective chance that I've made a mistake can aggregate over long chains of inference, knowledge is not closed under competent deduction. A major difference between her argument and the one presented here is that in the case of justification, the appropriate construal of risk concerns rational degree of confidence rather than objective chance.

4 Three Objections

I find the long sequence argument convincing. SPC and SPC* ought to be rejected. But not everyone will be convinced. In this section, I look at ways one might try to resist the argument.

The key principles of the long sequence argument are the following:

(Relevance of Rational Self-Doubt) Having a justified belief that one's deductive reasoning is not fully reliable partially defeats one's justification for the conclusion of a deduction.³⁵

(Existence of Rational Self-Doubt) Thinkers can (and do) have justified beliefs that their deductive reasoning is not fully reliable.

For each of these two principles, there is an objection worth considering. There is also an overgeneralization worry. These three lines of response raise important issues in epistemology, issues worth examining in their own right.

4.1 Can Rational Self-Doubt Defeat Justified Belief?

The first line of response concerns the first key principle – the relevance of rational self-doubt. The idea is that the initial plausibility of this principle is based on a confusion of levels.³⁶ In particular, a thinker may be justified in believing the conclusion of a competent deduction without being justified in believing that she is so justified. Rational self-doubt can defeat the meta-belief without defeating the ground-level belief.

In his discussion of the preface paradox against multi-premise closure (MPC) for knowledge, Williamson puts forward a version of this response:

One does indeed know each premise, without knowing that one knows it. Since one believes the conclusion on the basis of competent deduction from the

³⁵ Plausibly, having justification to believe that one's deductive reasoning is not fully reliable (whether or not one believes it) suffices. So does merely having the relevant belief (whether or not it is justified).

³⁶ See Alston (1980).

premises, by MPC one also knows the conclusion, although without knowing that one knows it. For each premise, it is very probable on one's evidence that one knows it. However, it is very improbable on one's evidence that one knows every premise. Given that one knows the conclusion (the conjunction) only if one knows every premise, it is very improbable on one's evidence that one knows the conclusion.³⁷

Williamson's view is that in the case of the preface paradox, we know the conclusion of the deduction without knowing that we know. Our evidence that we are fallible defeats our knowledge of our knowledge, but not our knowledge itself.

Williamson's discussion concerns knowledge and not justification. But an analogous claim might be made for justification. The suggestion would be that in the case of a long sequence of competent single-premise deductions from a justified initial premise, we have a justified belief in the conclusion of the deduction without being justified in believing that we are so justified. Our evidence that we are fallible defeats our justification for believing that we are justified, but not our justification itself.

This line of response has at least some initial plausibility. But there are several points that jointly serve to reduce its appeal.

The first point is a preliminary observation: The kind of defeat at issue is not specifically tied to deduction. The phenomenon is much more widespread. Very generally, a thinker should be less confident in a belief formed by a cognitive mechanism if the thinker has a justified belief that the mechanism is not operating reliably. For example, if I were to justifiably believe that my visual system is currently unreliable – say, on the basis of a report from my neurologist – this would defeat my justification for many of my visual beliefs. It would be epistemically irresponsible for me to continue to believe what my visual system tells me while also believing that my visual system is

³⁷ Williamson (forthcoming). Also see Williamson (2009) for relevant discussion.

unreliable. Similarly, if I were to justifiably believe that my predictions about the winners of baseball games are unreliable – say, on the basis of my emotional investment in my favored team and my mediocre track record – this would defeat my justification for some of my beliefs about baseball. Such defeat occurs even in cases in which my visual system or my reasoning about baseball is in fact perfectly reliable. So we can think about these other, non-deductive cases when evaluating the general line of response.³⁸

The second point is that there is something like a diagnosis of the thought that rational self-doubt does not defeat ground-level beliefs, only meta-beliefs. In particular, the line of thought would be appropriate if we had a picture on which all defeat is rebutting defeat.

There are two sorts of defeaters that are commonly recognized – rebutting defeaters and undercutting defeaters.³⁹ Very roughly, a rebutting defeater of a belief provides direct evidence that the belief is false. An undercutting defeater either provides direct evidence against the grounds on which the belief is held or direct evidence that the grounds do not genuinely support the belief in this instance. For example, seeing that my desk drawer is empty rebuts my belief that there is a pencil in the drawer. Learning that a wall that looks red is being illuminated by red lights undercuts my belief that the wall is red rather than white.⁴⁰

Discovering that my reasoning is not fully reliable does not rebut the conclusion of a deduction. It does not directly provide evidence that the conclusion is false. It can at

³⁸ There is an important contrast between worries about the reliability of the inputs to our reasoning – for instance, from vision – and worries about the reliability of our reasoning, itself. But I don't see how the line of response on offer could be sensitive to this contrast.

³⁹ See Pollock (1986). My characterization of the distinction between rebutting and undercutting defeat differs from his.

⁴⁰ This example is originally due to Pollock.

best rebut the claim that the conclusion of the deduction is justified. That is, it rebuts the meta-belief and not the ground-level belief. So if we had a picture on which all defeat is rebutting defeat, the Williamsonian line of response would be appropriate.

But this is a bad picture of defeat. There are numerous examples of undercutting defeat. At the very least, the claim that all defeat is rebutting defeat is a surprising one. It is a significant commitment, in need of substantial theoretical support.

To be fair, the kind of defeat at issue in the long sequence argument is not the kind that plays a role in familiar cases of undercutting defeat. In familiar cases of undercutting defeat, we gain direct evidence against the claim that the grounds support the relevant belief. This is not what happens in cases of rational self-doubt. Learning that my reasoning is not fully reliable does not provide direct evidence that the premises of my deduction do not support the conclusion. Rather, it provides direct evidence that I may not be assessing my evidence correctly. So this is a different kind of defeat. It is what we might call “higher-order defeat”.⁴¹

The third point in reply to the Williamsonian line of response is that there are convincing cases of this kind of defeat. Here is a case originally due to Adam Elga:⁴²

I am the pilot of an airplane. I need to make a mathematical calculation about which direction to turn the wheel of the plane. I discover that the plane is at a sufficiently high altitude that I am likely suffering from a case of hypoxia. (Hypoxia is a condition brought upon by high altitudes. It makes sufferers prone to errors in their reasoning, including their mathematical calculations. In mild cases, it is introspectively undetectable.) I’m not actually suffering from hypoxia.

⁴¹ Christensen (2010) uses the term “higher-order evidence” in discussing this kind of defeat. One way to get a grip on the contrast between undercutting defeat and higher-order defeat is in terms of conditional probabilities. The probability that a wall that looks red is red is presumably greater than the probability that a wall that looks red is red given that the wall is illuminated by red lights. In contrast, suppose that some premise entails some conclusion but that seeing this entailment relies on a complex bit of reasoning. The probability that the conclusion is true given that the premise is true is no greater than the probability that the conclusion is true given that the premise is true and given that I’m unreliable in the relevant kind of reasoning.

⁴² See Elga (unpublished).

I claim that I should be much less confident in the result of my directional calculation if I think it likely that I have hypoxia. My justified belief that I may be suffering from a condition that makes me prone to mistakes in my mathematical calculations partially defeats my justification for believing the conclusion of my reasoning. This is so even if I'm in fact reasoning correctly.

Here is a more mundane case originally due to David Christensen:⁴³

I have just been balancing my checkbook and have come up with a figure. Although I've checked my math several times, I am well aware that I have made mistakes repeatedly in the past, even after carefully checking my math. I've not actually made any mistakes in my calculation.

I claim that I should not be fully confident in the result of my checkbook calculation. My justified belief that I am prone to errors in summing long columns of numbers partially defeats my justification for believing the conclusion of my arithmetical reasoning.

The intuition behind these cases concerns epistemic responsibility. It is irresponsible to be highly confident in the conclusion of my directional calculation if I have a justified belief that I am likely suffering from hypoxia and thus am prone to mathematical errors. It is irresponsible to be highly confident in the conclusion of my checkbook calculation if I have a justified belief that I am prone to errors in balancing my checkbook.

One way to make these cases more pressing still is to emphasize that it would be irresponsible to *act* on the relevant beliefs. It would be irresponsible for me to turn the wheel of the plane the relevant number of degrees without carefully checking my directional calculations (and also breathing in some oxygen, talking to air traffic control, and so on). It would be irresponsible for me to make expensive purchases without making

⁴³ See Christensen (2008).

as sure as I can that I didn't screw up my checkbook calculations. A natural diagnosis of these facts is that I am not justified in holding the relevant beliefs.⁴⁴

I can think of two potential responses to these apparent cases of higher-order defeat. Both responses attempt to explain away our intuitions about the cases.

The first response is to say that since I'm actually reasoning correctly, I have a justified belief, but I have an excuse for double-checking my reasoning. (Perhaps, for instance, I have an excuse in part because I am not justified in believing that I am justified.)⁴⁵ The trouble with this suggestion is that it is not merely the case that in the hypoxia scenario, for example, I have an excuse for double-checking my calculations. Rather, I count as irresponsible if I *fail* to double-check my calculations. Appealing to an excuse cannot explain this fact.

The second response is to say that since I'm actually reasoning correctly, I have a justified belief, but I am exhibiting some other kind of failing in not double-checking my reasoning. This may be a moral or a pragmatic failing. Or it may be a kind of epistemic failing that is compatible with having a justified belief. For instance, I may be exhibiting an epistemically bad character. (Again, perhaps I have this failing in part because I am not justified in believing that I am justified.) This suggestion is implausible, too. If I'm actually reasoning correctly in a particular case, why am I exhibiting a failing if I don't double-check my reasoning? Presumably, the thought is that I should double-check my reasoning because it is a matter of luck that I was reasoning correctly. But then it seems close to inescapable to conclude that I am epistemically irresponsible in not double-

⁴⁴ This diagnosis is especially natural for those who endorse a justification or a knowledge norm on action.

⁴⁵ See Williamson (2000), pages 257-8, for an analogous response put forward in defense of a knowledge norm on assertion.

checking. My failing is an epistemic one, not merely a moral or a pragmatic one. And it is a failure of epistemic responsibility, not a different kind of epistemic failure.⁴⁶

It is easy to generate additional examples of higher-order defeat. Such cases strongly suggest that we should accept the first key principle of the long sequence argument. They are not happily described as cases in which we are justified in believing the conclusion of our reasoning (but not in believing that we are so justified). In both the hypoxia and the checkbook cases, it is epistemically irresponsible to believe the conclusion, at least with as much confidence as we have in the relevant premises. Thus, I conclude that having a justified belief that one's deductive reasoning is not fully reliable can partially defeat one's justification for the conclusion of a deduction.

4.2 Is There a Limit to Rational Self-Doubt?

The second line of response to the long sequence argument concerns the second key principle – the existence of rational self-doubt. There isn't any good prospect for arguing that we are *never* justified in believing that our reasoning is unreliable. But there is an interesting argument that suggests that it is impossible for thinkers to have certain kinds of rational self-doubt. The basic line of thought is based on an argument originally due to David Lewis concerning immodest inductive methods.⁴⁷ I've modified Lewis's argument in a several ways. The most important, perhaps, is that Lewis's argument originally concerned empirical rules – indeed, a thinker's entire inductive practice. But it seems to me that the argument applies equally well to deductive rules of inference.

⁴⁶ This response faces a second problem. There are apparent cases of higher-order defeat in which I have double-checked my calculations as much as is possible for me. In such cases, there is no prospect of claiming that I am exhibiting some kind of failing in not further checking my reasoning.

⁴⁷ See Lewis (1971). Also see Field (2000) and Elga (2010) for closely related arguments.

Some terminology will be helpful in presenting the argument. Let a thinker's basic rules be those rules of inference that the thinker employs not in virtue of employing any other rules. Let a thinker's epistemic system be the collection of basic rules that the thinker employs. Let a thinker believe that a rule is comparatively unreliable if the thinker believes that the rule is less reliable⁴⁸ than some specific competing⁴⁹ rule that she could instead employ.⁵⁰ Let a thinker believe that a rule is untrustworthy if the thinker believes that she ought not to employ it.

There are several versions of what I'll call "the fixed-point argument". Here are two:

If a thinker's epistemic system yields the belief that one of its component rules is comparatively unreliable, then it is making two conflicting recommendations: The system implicitly recommends employing the basic rule. It explicitly recommends instead employing an alternative rule. It is incoherent for a thinker's epistemic system to issue such conflicting recommendations. So it is impossible for a thinker to have a justified belief that one of his basic rules is comparatively unreliable.

If a thinker's epistemic system yields the belief that one of its component rules is untrustworthy, it is making two conflicting recommendations: The system implicitly recommends employing the basic rule. It explicitly recommends not employing the rule. It is incoherent for a thinker's epistemic system to issue such conflicting recommendations. So it is impossible for a thinker to have a justified belief that one of his basic rules is untrustworthy.⁵¹

The fixed-point argument can also be put in terms of incoherence among beliefs. Suppose that my epistemic system includes a rule that tells me to believe that p (given my current epistemic state). Suppose that it also tells me that I shouldn't fully trust that very

⁴⁸ There is a complication here due to the fact that there are different measures of reliability. For simplicity, I'll assume that there is a single relevant measure.

⁴⁹ Two rules are competitors if they provide incompatible pronouncements on what to believe on some matter given the very same inputs.

⁵⁰ The restriction to rules that the thinker could employ is intended to exclude from consideration such rules as "believe all and only the truths".

⁵¹ I suspect that the second version of the argument is more fundamental than the first.

rule. In this case, it tells me that I should believe that p and that I shouldn't fully trust the recommendation to believe that p. This is incoherent.

Lewis provides an analogy that may help to clarify the main thrust of the argument. His analogy involves *Consumer Reports*, a magazine that rates consumer products. Suppose that *Consumer Reports* had an issue rating the reliability of magazines. Suppose that it gave itself a bad rating compared to another magazine that also rates consumer products. This would seem to be incoherent. *Consumer Reports* would simultaneously be implicitly telling us to trust its pronouncements (for instance on whether to buy a certain toaster) and explicitly telling us not to put complete credence in its pronouncements.

Presumably, if it is incoherent to believe that one's basic rules are comparatively unreliable or untrustworthy, then a thinker is not justified in so believing. According to this line of thought, then, it is impossible for thinkers to have certain kinds of rational self-doubt about their basic rules.

This conclusion is not enough to save SPC or SPC*. However, it suggests that we may be able to adopt a weakened version of single-premise closure:

(SPC for Basic Rules) Necessarily, if S has a justified belief that p and comes to believe that q solely on the basis of competently deducing it from p *by applying one of S's basic rules*, while retaining the justified belief that p throughout the deduction, then S has a justified belief that q.

If we like, we can also add a no-defeaters clause.

I find the fixed-point argument arresting. I'm not entirely sure what we should say about it. But I do want to make two general points about it here. The first point is that there are reasons to think that the fixed-point argument must go wrong somewhere.

Endorsing the conclusion of the argument leads to many counterintuitive consequences.

One counterintuitive consequence is that the conclusion of the argument suggests that it is never rational to revise one's set of basic rules. This is because we can never be in a position to think that one of our basic rules is not to be trusted. If we have a view on the matter at all, we should believe the rule to be fully trustworthy. This is intuitively objectionable. We seem able to put any part of our epistemic system up for rational debate.⁵²

A second counterintuitive consequence is that the conclusion of the argument yields an asymmetry between first person and third person cases. Suppose that I discover strong evidence that seems to tell against the reliability of someone's rule. I should believe that the relevant thinker employs a comparatively unreliable rule, so long as that thinker is not me.

One way to see this contrast is to imagine a viewer watching a videotape of a subject being experimented on. Suppose the videotape provides good evidence that the subject is employing one of his basic rules. Suppose, too, that the videotape shows that the subject is making a series of blunders in response to the experiment. The viewer would seem to be in a good position to determine that the subject's rule was comparatively unreliable. But if the viewer then discovers that the lab subject was himself, suddenly, he must withdraw this judgment. This is a very strange consequence.⁵³

A third counterintuitive consequence is that combined with plausible claims about explanation, the view leads to something like an *a priori* argument that my rules are extremely reliable and that there is some explanation of this fact. Let me explain. The fixed-point argument suggests that I must believe that my basic rules of inference are

⁵² See Field (2000; 2009a) for responses to this worry.

⁵³ I owe this way of putting the point to Adam Elga.

comparatively reliable (if I have a view on the matter at all). They are at least as reliable as any available alternative. That a thinker employs only comparatively reliable basic rules of inference is a striking phenomenon. It “calls out” for explanation. In general, there is rational pressure to reject any view that posits striking phenomena that are accidental or otherwise unexplained. So there is rational pressure to believe there is some explanation of my reliability. This leads to an *a priori* pressure to believe that my basic rules of inference are comparatively reliable and that this is not accidental.

This is a troubling consequence. But it gets worse still. There are good empirical reasons to think that ordinary thinkers do not employ comparatively reliable basic rules. So I should believe that I am epistemically special. Unlike everyone else, *my* basic rules are comparatively reliable.⁵⁴ This is a striking fact in need of explanation. So there must be some explanation of it – for instance, perhaps some epistemically benevolent being favors me over everyone else. This is an extremely counterintuitive consequence.

Of course, all of this is just to point out that there are reasons to think that there must be a problem with the fixed-point argument. It is not to identify the problem. I suspect that the right diagnosis is that (a) in some cases, we can be justified in believing the pronouncements of an incoherent epistemic system; and (b) we are rationally required to be incoherent. In general, there is a rational ideal to avoid epistemic incoherence – we should avoid commitments that can be seen by us to be incompatible. But in some cases, we are rationally required to violate this ideal. Being an epistemically modest thinker is also a rational ideal, as is apportioning one’s beliefs to the evidence. Epistemic modesty

⁵⁴ Indeed, I should also think that every thinker is justified in having the corresponding belief. Everyone should believe that they are epistemically special. But unlike everyone else, *I* really am epistemically special. Or so I should think. I owe this point to Phil Galligan.

and respecting the evidence each require that we believe that our basic rules are imperfect. This leads to incoherence.⁵⁵

However, I am not fully confident in this diagnosis. The claim that we are rationally required to be incoherent is a radical view.

There is a second important point to make in response to the fixed-point argument. Even if we accept the conclusion of the argument, the conclusion does not support SPC for Basic Rules. Even if thinkers are never justified in believing that one of their basic rules is untrustworthy or comparatively unreliable, this does not rule out all possible risks. Recall that there are several different kinds of potential mistakes that a thinker can make in reasoning. Even if the fixed-point argument shows that I cannot be justified in believing that I am using an unreliable basic rule, it does not show that I cannot be justified in believing that I may have *misapplied* a basic rule in my reasoning. The fixed-point argument does not straightforwardly apply to this risk. I should have at least some credence in the claim that I have misapplied a rule – whether through inattention, fever, being slipped a drug, hypoxia, random error, or whatever. My rational degree of confidence that I have misapplied a rule may be vanishingly small for any particular inferential step. But it is non-zero. And this risk can aggregate.

Indeed, this was the major intuitive risk in the long sequence argument. The principal worry was not that I may be employing a bad deductive rule but simply that I may have made a mistake in applying the rules that I employ (or in chaining together deductions). So the long sequence argument still stands. There are violations of single-premise closure even for basic rules.

⁵⁵ See Christensen (2008) for discussion of whether rational ideals can be jointly incoherent.

One can formulate a version of the fixed-point argument aimed at the risk of misapplying a rule. The argument would go something like this:

If a thinker's epistemic system yields the belief that the thinker may have misapplied a rule, then it is making two conflicting recommendations: The system implicitly recommends employing the rule. The system explicitly recommends not fully trusting what the rule has in fact pronounced. It is incoherent for a thinker's epistemic system to issue such conflicting recommendations. So it is impossible for a thinker to have a justified belief that the thinker may have misapplied the rule.

Yet, this version of the fixed-point argument is not very compelling. The epistemic system does not tell the thinker to employ a rule and to instead employ a competing rule. The system does not tell the thinker to employ a rule and to refrain from employing the rule. Rather, it merely tells the thinker that it should not fully rely upon a claim that is in fact the output of one of its component rules. This is a kind of incoherence, but it is not an intuitively worrisome kind. It is just what someone should think if they thought they'd screwed up somewhere in their reasoning. This version of the fixed-point argument thus has very little bite.

Here, then, is what we should conclude about the fixed-point argument: There are several reasons to think that the argument doesn't work. If it does work, we are forced to say that we cannot justifiably believe our basic rules to be untrustworthy or comparatively unreliable. This would be an interesting result. It would provide a limit to rational self-doubt. But even if it were to slightly lessen the magnitude of violations of closure, it would not address the principal worry underlying the long sequence argument.

4.3 Overgeneralization

The third and final line of response to the long sequence argument is based on an overgeneralization worry. This line of thought does not pinpoint a specific problem with the argument, but rather suggests that it must go wrong somewhere, since analogous considerations lead to an unpalatable conclusion.

Suppose that R is any rule that we employ in our reasoning. Perhaps R is a deductive rule. Perhaps it is an ampliative one. Perhaps it is a rule that tells us how to take into account our own fallibility in our reasoning. We are justified in believing that our employment of R is not fully reliable. So, by analogous reasoning to the above, we are not justified in simply forming beliefs via applications of R to the degree that R licenses. Rather, we should be a bit more cautious in the formation of our beliefs. At the very least, we should take into account the possible ways in which we may be making mistakes in our reasoning. We should rationally police our reasoning with rule R. But then it seems that there is *no* rule that we are justified in simply obeying in our reasoning.⁵⁶

The problem is not that we should rationally police our rules, and then rationally police our rational policing, and so on, which leads to our credences draining away.⁵⁷ The problem doesn't directly concern the magnitude of our credences. Rather, the problem is that there seems to be no rule that we should *simply obey* in our thinking. But then what are we to do? At some point, we have to simply obey some rule in our reasoning. So the

⁵⁶ This line of thought is reminiscent of the excerpt from Hume. It is also reminiscent of the regress argument in Carroll (1895), albeit in a much more general setting.

⁵⁷ This is essentially the argument that Hume goes on to make in the section that the excerpt is taken from. Hume's argument is fallacious. Among other problems, we needn't think we are more likely to have credences that are too high than credences that are too low. So there is no reason to think that our credences will drain away.

conclusion of the argument seems to make rational reasoning impossible. It leads to a form of epistemic nihilism.

I take it that the right thing to say here is that, *ideally*, we should rationally police each of our rules – we should take into account the risk of making a mistake in forming our beliefs. But we cannot do this because of resource limitations and because it is impossible to simultaneously rationally police all of our rules. So we are justified in simply obeying certain rules and ignoring the risk of making a mistake, at least in certain circumstances. Which rules are these and in what circumstances? This is the crux of the issue.

One view is that we needn't rationally police our deductive rules of inference. Another view – defended above – is that we should. We needn't rationally police some of our other rules. So the main dispute comes down to the question of where rational policing comes to an end.

As we have seen, deductive rules provide a bad candidate for the place where rational policing comes to an end. The cases presented above suggest that we are not justified in simply obeying our deductive rules. The only reason one might have thought that deductive rules are exempt from rational policing is that deductive rules cannot lead us astray if they're correctly applied. But, as we've seen, the principal worry underlying the long sequence argument is not that our rules are bad rules, but that we may have misapplied them. Like any other rule, deductive rules are subject to this risk. Thus, there is no reason to think that they are exempt from rational policing.⁵⁸

⁵⁸ Moreover, the thought that deductive rules are exempt from rational policing because they cannot lead us astray seems to depend on an overly reliabilist conception of epistemic responsibility.

A more promising picture is that we are justified in simply obeying higher-order rules that tell us how to rationally police our other rules (or perhaps still higher-order rules that tell us how to rationally police those rules or ...). Such higher-order rules are not immune from potential misapplication. But we should refrain from policing them because it is too costly to do so. On this picture, the amount of rational policing that thinkers should do depends on their resource limitations. Thinkers should police as much as is reasonable, given their cognitive powers.⁵⁹

On this view, justification turns out to be subject-sensitive in the sense that it does not depend only on a thinker's evidence, but also on her available cognitive resources. This strikes me as a plausible view. A thinker's epistemic responsibilities plausibly depend on her cognitive resources. The more cognitive resources a thinker possesses, the more epistemic responsibilities she has.⁶⁰

This picture of rational policing is intuitively plausible. It does a much better job of fitting our intuitive judgments about cases. I conclude, then, that it is the picture we ought to endorse. This answers the overgeneralization worry.

5 Conclusion

There is a final question worth considering: Why are strong closure principles intuitively so appealing? Why are we tempted to think that deductive inference is fully epistemically secure?

⁵⁹ Policing as much as is reasonable should not be identified with policing as much as is possible (given the thinker's cognitive powers). Indeed, it may be problematic to police one's reasoning as much as is possible. Too much policing may introduce more errors than it corrects. Thanks to Gideon Rosen for pressing me on this issue.

⁶⁰ This is one application of what might be called "the Spiderman principle" in epistemology: With greater cognitive power comes greater epistemic responsibility. This is a plausible principle. For instance, it helps to explain why small children have fewer epistemic obligations to check their reasoning than we do.

The reason strong closure principles are intuitively attractive is that our paradigm cases of deductive reasoning include short surveyable deductions as well as mathematical proofs. It is very natural to come to think that a strong closure principle holds when reflecting about these cases. In the case of surveyable inferences, violations of closure are typically minor. The risk of making a mistake is not very large. Mathematical proofs are typically extensively checked. The conclusions of mathematical proofs are typically checked against examples and against expert mathematical intuition. Moreover, there is an interesting phenomenon in mathematics. Namely, mathematical proofs are *robust* – when mathematicians believe they have a proof of some claim, there typically is a proof of it or of some related result. It is for these reasons that we are easily misled into thinking that Closure holds.

Despite their intuitive appeal, however, we should reject strong closure principles. If what I've argued here is correct, we should reject Closure, Single-Premise Closure, Single-Premise Closure*, and related principles.⁶¹ More generally, there is no (non-trivial) principle of the form: If S has a justified belief that p, S competently deduces q from p, and conditions C obtain, then S has a justified belief that q. There is no guarantee that if we're in the right sort of circumstances and we in fact reason correctly, our reasoning preserves justification. No rule of inference, no matter how simple, is safe from potential misapplication. No deduction is fully epistemically secure.

⁶¹ There are three ways of weakening the closure principle for justification to try to address the long sequence argument while still maintaining that deductive inference is, in some sense and in some circumstances, fully epistemically secure: (i) Closure applies to propositional rather than doxastic justification; (ii) Closure only applies to ideal epistemic agents; (iii) Closure is a rational ideal. I think each of these proposals is untenable, but I do not have the space to argue for this here.

Acknowledgements

Earlier versions of this paper were presented at the Basic Knowledge III workshop at the University of St. Andrews, a workshop on epistemology at the University of Geneva, the Theoretical Philosophy forum at Eötvös University, departmental colloquia at the University of Connecticut and at Princeton University, and the Epistemology Reading Group at MIT. I would like to thank the audiences at these events for their questions and comments. I would also like to thank Maria Lasonen Aarnio, Paul Boghossian, David Christensen, Stew Cohen, Dylan Dodd, Adam Elga, David Estlund, Hartry Field, Phil Galligan, Michael Heumer, Chris Hill, Ram Neta, Stephen Read, Gideon Rosen, Nico Silins, Paul Silva, Ralph Wedgwood, Tim Williamson, and Zsotia Zvolensky for helpful discussions at various stages of this project.

References

- Alston, W. (1980). Level confusions in epistemology. *Midwest Studies in Philosophy*, 5, 143–145.
- Boghossian, P. (1989). The rule following considerations. *Mind*, 98, 507–549.
- Boghossian, P. (2008). Epistemic rules. In his *Content and justification*, (pp. 109–134). Oxford: Oxford University Press.
- Carroll, L. (1895). What the tortoise said to Achilles. *Mind*, 4, 278–280.
- Cheng, P. and Holyoak, K. (1985). Pragmatic reasoning schemas. *Psychology*, 17, 391–416.
- Christensen, D. (2004). *Putting logic in its place*. Oxford: Oxford University Press.
- Christensen, D. (2008). Does Murphy’s law apply in epistemology? Self-doubt and rational ideals. *Oxford Studies in Epistemology*, 2, 3–31.
- Christensen, D. (2010). Higher-order evidence. *Philosophy and Phenomenological Research*, 81, 185–215.
- Cosmides, L. (1989). The logic of social exchange. *Cognition*, 31, 187–276.
- Dretske, F. (1970). Epistemic operators. *Journal of Philosophy*, 67, 1007–1023.
- Elga, A. (2010). How to disagree about how to disagree. In R. Feldman and T. Warfield (Eds.), *Disagreement*. Oxford: Oxford University Press.
- Elga, A. (unpublished). Lucky to be rational.
<http://www.princeton.edu/~adame/papers/bellingham-lucky.pdf>.

- Evans, J., Newstead, S., and Byrne, R. (1993). *Human reasoning: The psychology of deduction*. Hillsdale: Lawrence Erlbaum.
- Evnine, S. (1999). Believing conjunctions. *Synthese*, 118, 201–227.
- Field, H. (2000). A prioricity as an evaluative notion. In P. Boghossian and C. Peacocke (Eds.), *New essays on the a priori* (pp. 117–149). Oxford: Oxford University Press.
- Field, H. (2009a). Epistemology without metaphysics. *Philosophical Studies*, 143, 249–290.
- Field, H. (2009b). What is the normative role of logic? *Proceedings of the Aristotelian Society Supplementary Volume*, 133, 251–268.
- Harman, G. (1986). *Change in view: Principles of reasoning*. Cambridge, MA: The MIT Press.
- Harman, G. (1995). Rationality. In E. Smith and D. Osherson (Eds.), *Thinking: An invitation to cognitive science*, vol. 3 (pp. 175–211). Cambridge, MA: The MIT Press.
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Oxford University Press.
- Kripke, S. (1982). *Wittgenstein on rules and private language*. Cambridge, MA: Harvard University Press.
- Kyburg, H. (1970). Conjunctivitis. In M. Swain (Ed.), *Induction, acceptance, and rational belief* (pp. 55–82). New York: Humanities Press.
- Lasonen-Aarnio, M. (2008). Single premise deduction and risk. *Philosophical Studies*, 141, 157–173.
- Lewis, D. (1971). Immodest inductive methods. *Philosophy of Science*, 38, 54–63.
- Makinson, D. (1965). The paradox of the preface. *Analysis*, 25, 205–207.
- Nelkin, D. (2000). The lottery paradox, knowledge, and rationality. *The Philosophical Review*, 109, 373–409.
- Nozick, R. (1981). *Philosophical explanations*. Cambridge, MA: Harvard University Press.
- Pollock, J. (1983). Epistemology and probability. *Synthese*, 55, 231–252.
- Pollock, J. (1986). *Contemporary theories of knowledge*. First edition. Towota, NJ: Rowman and Littlefield Publishers.
- Vogel, J. (1990). Are there counterexamples to the closure principle? In M. Roth and G. Ross (Eds.), *Doubting: Contemporary Perspectives on Skepticism* (pp. 13–27). Dordrecht: Kluwer.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
- Williamson, T. (2009). Reply to Hawthorne and Lasonen-Aarnio. In P. Greenough and D. Pritchard (Eds.), *Williamson on knowledge* (pp. 313–329). Oxford: Oxford University Press.
- Williamson, T. (forthcoming). Very improbable knowing. In T. Dougherty (Ed.), *Evidentialism and its discontents*. Oxford: Oxford University Press.
- Wright, C. (1985). Facts and certainty. *Proceedings of the British Academy*, 7, 429–472.