Forthcoming in *Synthese*

Daniel Immerman

Williamson, Closure, and KK

Abstract. *Closure principles say that if you know some proposition which entails a second and you meet further conditions then you know the second. In this paper I construct an argument against closure principles which turns on the idea that knowing a proposition requires that one's belief-forming process be reliable. My argument parallels an influential argument offered by Timothy Williamson against KK principles – principles that say that if you know some proposition and you meet further conditions then you know that you know the proposition. After offering my argument, I provisionally assess its damage to closure principles and also look at how responses to my argument against closure principles can be used to generate responses to Williamson's argument against KK principles.*

Keywords. Closure Principles; Knows-Knows Principles; Timothy Williamson; Internalism; Externalism

This paper focuses on the merits of two types of principle. The first is a "single premise knowledge closure principle" or a "closure principle" for short.[1] Here is a simple schema for this type of principle:[2]

> (Closure): For any subject S and propositions p and q such that p entails q, if S knows p and certain conditions obtain, then S

[1]For discussion of closure principles, see e.g. [Baumann, 2012], [Feldman, 1995], [Kvanvig, 2006], [Lawlor, 2005], [Luper, Winter 2011], [Luzzi, 2010], [Schechter, 2013] and [Vogel, 1990].

[2]I have introduced closure principles via a schema and will talk as if there are multiple closure principles. However, some talk as if there is only one closure principle. For instance, the *Stanford Encyclopedia of Philosophy* article devoted to epistemic closure principles is called "The Epistemic Closure Principle" [Luper, Winter 2011]. I find talk of closure principles in the singular to be somewhat mysterious. Perhaps people who engage in it are assuming there is only one closure principle best suited to a particular role? In any case, I shall continue to refer to closure principles in the plural. Similar things go for KK principles.

knows q.[3]

The second is a "knows knows principle," or a "KK principle" for short.[4] Here is a simple schema for this type of principle:

> (KK): For any subject S and proposition p, if S knows p and certain conditions obtain, then S knows that S knows p.

It is commonly held that there are some true closure principles. For instance, Richard Feldman writes, "the idea that no version of this principle is true strikes me, and many other philosophers, as one of the least plausible ideas to come down the philosophical pike in recent years" [Feldman, 1995, 487].[5]

KK principles have a less good reputation. For instance, Timothy Williamson, writes "It is widely, though not universally, acknowledged that the KK principle is false" [Williamson, 2011, 147]. This is largely due to Williamson himself, who offers an influential argument against a particular KK principle in *Knowledge and its Limits* [Williamson, 2000, 114-6]. In this work, he carefully formulates the KK principle he wishes to criticize so as to avoid counterexamples that had been pressed against other, less carefully formulated principles. As he notes of the KK principle he criticizes, it is "a special case of the general 'KK' principle that if one knows something then one knows that one knows it, but sufficiently restricted to avoid many of the objections to the latter" [Williamson, 2000, 115]. Thus Williamson's argument could be viewed as the final nail in the coffin; even the most carefully formulated KK principle seems doomed to failure.

What does not seem to have been noticed is that a modified version of Williamson's argument can be applied to closure principles.[6] The first part of the paper is devoted to explaining this modified argument against closure and showing how its key premise may be defended. Notably, its key premise rests on a general principle connecting knowledge and reliability that Williamson's argument against KK also appeals to.

---

[3]Stating a more general schema for this type of principle requires allowing for various bells and whistles. These include time indices and modal operators, such as being in a position to know.

[4]For discussion of KK principles, see e.g. [Castañeda, 1970], [Conn, 2001], [Das and Salow, 2018], [Feldman, 1981], [Ginet, 1970], [Goodman and Salow, 2018], [Greco, 2014], [Greco, 2015], and [Hemp, 2014].

[5]For some more quotes in a similar vein, see [Dretske, 2005, 17].

[6]It has been noticed that a modified version can target other sorts of principles; see e.g. [Ramachandran, 2012, 128-31].

The argument I discuss targets a particular closure principle, just as Williamson's argument targets a particular KK principle. So two natural questions arise. First: how good of a job does it do at targeting the closure principle in question? My response: it does a good job; it is pretty persuasive. In particular, the argument turns on a key premise which is in turn is supported by a general principle that is quite plausible.[7] The argument is also significant dialectically, in that if one wishes to deny this general principle, one thereby undermines Williamson's argument against KK.[8]

The fact that my argument targets a *particular* closure principle also gives rise to another natural question, namely: are there other closure principles that can escape my argument? In the later parts of my paper, I will provisionally discuss this question. My provisional conclusions are as follows: one sort of closure principle – a higher-order one – is successfully targeted by my argument. Two others, one related to competent deduction and another related to proper basing, can escape, but each has some limitations. I also provisionally defend the conclusion that for each of the closure principles that escapes my argument, there is a parallel KK principle that escapes Williamson's argument, and that this KK principle inherits some of the motivation that applies to the closure principle it parallels. This proves unfortunate for Williamson. He had hoped, by attacking KK principles, to block certain conclusions they supported. But unfortunately, as we shall see, some of the KK principles that his argument fails to target can be used to support these conclusions.

This paper has six sections. The first section introduces my argument against closure. The second section examines how a key premise of my argument may be defended. The next three sections look, in turn, at three different closure principles and whether they can escape my argument. The sixth section looks at whether there are parallel KK principles, motivated by similar considerations, that escape Williamson's argument.

1 My argument against closure

My argument begins with a story about a character named Lisa. Lisa is working through a series of deduction problems which very gradually increase

[7] As I'll note later on, the principle in question isn't entirely uncontroversial, but this should be unsurprising; arguments for deeply controversial views in epistemology rarely have supports that are entirely uncontroversial.

[8] Thanks to an anonymous referee for suggesting I mention this here.

in difficulty. Each problem gives Lisa two propositions. As Lisa knows, the first of these, which I'll call "the starting proposition" is always a true and surprising fact. So every time Lisa starts a problem, she comes to learn some new thing. Meanwhile, the second proposition, which I'll call "the ending proposition", has an important logical connection to the starting proposition. In particular, for each problem either (i) the starting proposition entails the ending proposition or (ii) the starting proposition entails the negation of the ending proposition. Lisa is not told which; this is one of the things she must figure out. To correctly solve a problem, Lisa must either deduce the ending proposition or its negation from the starting proposition.

For instance, the starting proposition of the first problem is: "Mona Lisa does not have eyebrows." and the ending proposition is: "Either Mona Lisa does not have eyebrows or the brain of an ostrich is the same size as one of its eyeballs." In this case, the starting proposition entails the ending proposition, so to solve this problem, Lisa must deduce the ending proposition from the starting proposition.

Lisa is a competent deducer when it comes to the first problem; she can get this problem and others at this level of difficulty right every time. But as the problems become gradually more difficult, Lisa begins to find them harder and harder. Let us suppose that she makes her first error at problem number 666. Just to be clear, the idea isn't that she has a performance error; she doesn't suddenly become distracted at problem 666 and make a mistake. Rather the problems are sufficiently difficult at that point, containing so many complex steps, that she makes a mistake without noticing. Similar things regarding difficulty have been true for some time. For example, problem 665 was really difficult, contained many complex steps, and was such that Lisa could easily have made a mistake without noticing it.[9]

Let me say a little more about what makes problem 666 and others like it difficult. Sometimes the difficulty in solving a logic problem is in finding the correct answer, rather than verifying that the answer is correct. For instance, there can be cases in which you search for a proof of something for a while, and then in a flash of inspiration you find a proof and once you've found the proof it's blindingly obvious that the proof you've found is right.

[9]Here is a variant story: the problems remain at the same level of difficulty, but Lisa very gradually loses brain function. All my conclusions should go through for this variant story as well. So someone who wishes to challenge what I say here should make sure their challenge works on this variant as well.

But at other times, you can find a proof and still not be sure if it's right or not; it might be sufficiently complex, some of the steps sufficiently difficult to follow, that it's hard to tell if it's right. We should understand Lisa's work on the problems high in the 600's to be of this second type. That is, they're long and messy and complicated and she could easily have made a mistake somewhere without realizing it.

Consider an analogy. Suppose I am supposed to meet someone at the airport. There are a couple of ways in which this task can be difficult. One is the following: I know what the person looks like, but the airport is rather crowded, and so there are a lot of people to search through. For example, suppose I'm looking for my daughter. In such a case, the difficult part of the task is in spotting my daughter; once I spot her, I will recognize her instantly. A second sort of case is the following: I'm supposed to be meeting someone at the airport but I don't really know what the person looks like. For example, suppose I'm looking for a second-cousin I haven't seen for many years. In such a case, the difficult part of the task may well be correctly identifying my second-cousin. Perhaps I keep seeing people who look like him and can't really tell which one is correct.

The way I wish to understand the Lisa case, it is like the second sort of airport case. Even once she's found what she takes to be proofs, they are long and complicated and it's difficult to tell if they're right or not. Just as I am not particularly reliable at identifying someone as my second cousin thanks to the difficulty involved in distinguishing him from someone who looks similar, Lisa is not particularly reliable at identifying something as the correct answer to a problem high in the 600's, thanks to the difficulty in distinguishing real proofs from mistaken ones. This makes her work on problems 665 and 666 very similar; in both cases, they involve lengthy, convoluted logical arguments. She thinks the logical arguments are right in both cases, and even though she's right about one and wrong about the other, her belief-forming process in the two cases is virtually the same.

With this story in place, I'm ready to offer my argument against closure, which makes use of two principles.

First, there's a principle I'll call Set-Up$_{Lisa}$, because it's information provided by the set up to my argument:

> (Set-Up$_{Lisa}$). Lisa knows the starting proposition of each problem and for all but the last problem, the answer Lisa believes to be correct is indeed entailed by the starting proposition of that

problem.[10]

This principle should be uncontroversial; it just follows from details of the story.

My second principle I call (Reliability$_{Lisa}$) because it is connected with the idea that knowledge requires that one's belief-forming process be reliable:

> (Reliability$_{Lisa}$). If Lisa knows the answer to a given problem, she does not get the next problem wrong.[11]

This principle is supported by details in the story; in particular, the fact that Lisa's belief-forming process from one problem to the next is so similar. In such a circumstance, for it to be the case that she really truly knows the answer to a given problem, and is not just luckily guessing, it has to be the case that she does not get the next problem wrong.

This principle is more controversial than the previous one, and I'll devote the next section to defending it. But to briefly preview, it follows from two claims: first, that knowledge requires that one not believe falsely in similar circumstances and second, that the way in which Lisa forms belief regarding the answer to a given problem is very similar to the way in which she forms belief regarding the answer to the next.

Together, my two principles entail the denial of a closure principle, which I'll call (Closure$_{Challenged}$), seeing as my argument challenges it:

> (Closure$_{Challenged}$) For a given problem, if Lisa knows the starting proposition to that problem and believes another proposition entailed by it, then she knows it.[12]

[10] Let me state this formally. Let (Start$_i$) be the starting proposition to problem i and (LisaAnswer$_i$) be the answer that Lisa gives to problem i, that is, the proposition that she believes to be deducible from (Start$_i$). Then (Set-Up$_{Lisa}$) says: for all i from 0 to 666, Lisa knows (Start$_i$) and for all i from 0 to 665, Lisa believes (LisaAnswer$_i$) and (Start$_i$) entails (LisaAnswer$_i$).

[11] Let me state this formally, keeping the same abbreviations as the previous footnote. (Reliability$_{Lisa}$) says: for all i from 0 to 665, if Lisa knows (LisaAnswer$_i$) then (LisaAnswer$_{i+1}$) is true.

[12] Let me state this formally, keeping the same abbreviations as the previous footnote. (Closure$_{Challenged}$) says: for all i from 0 to 666, if Lisa knows (Start$_i$) and (Start$_i$) entails q and Lisa believes q then Lisa knows q.

Let me briefly outline how (Set-up$_{Lisa}$) and (Reliability$_{Lisa}$) entail the denial of (Closure$_{Challenged}$). The argument is a reductio. First, note that (Set-up$_{Lisa}$) and (Closure$_{Challenged}$) together entail, for each of problems 0 through 665, that Lisa knows the conclusion she believes to be the correct answer to that problem. After all, for each of these problems, (Set-Up$_{Lisa}$) tells us that (i) Lisa knows the starting proposition of that problem and (ii) the proposition she believes to be the answer is in fact entailed by that starting proposition. And so, from (Closure$_{Challenged}$) it follows that she knows the proposition she believes to be the answer.

However, Lisa's answer to problem 666 is incorrect. So it follows from (Reliability$_{Lisa}$) that she doesn't know the proposition she believes to be the answer to problem 665. This contradicts the claim that she knows the conclusion of each of problems 1 through 665. So the principles cannot all be true. This means that if (Set-Up$_{Lisa}$) and (Reliability$_{Lisa}$) are true, then (Closure$_{Challenged}$) is false.[13]

Of course, my argument will only be successful if my two principles, (Set-Up$_{Lisa}$) and (Reliability$_{Lisa}$), are correct. As noted above, (Set-Up$_{Lisa}$) should be uncontroversial; it just follows from details of the story. This leaves (Reliability$_{Lisa}$). I have briefly outlined a defense of it, but in the next section, I will develop this defense.

2 Defending my premise linking reliability and knowledge

In developing my defense of (Reliability$_{Lisa}$), it will be helpful to see how Timothy Williamson defends a related principle in an argument he offers against KK principles. This is useful for two reasons: first, I can use what he says in defending his related principle to help defend (Reliability$_{Lisa}$). Second, some readers may wish to reject (Reliability$_{Lisa}$) and it's useful to see how this rejection will affect Williamson's argument against KK.

In order to understand the particular KK principle Williamson is challenging and Williamson's argument against it, we must first understand a story Williamson tells about a character named Mr Magoo.

[13] I would like to briefly note how this argument differs from another argument challenging closure principles, viz. Maria Lasonen-Aarnio's argument in "Single Premise Deduction and Risk." Her argument is explicitly fallibilist, resting on the idea that one can have knowledge if one avoids relevantly similar false belief in most close worlds, even if one doesn't do so in all of them [Lasonen-Aarnio, 2008, 167]. But I make no such assumption. Thanks to an anonymous referee for suggesting I clarify this.

In the story Mr Magoo is looking out his window at a tree that is 666 inches tall and has formed various beliefs about the height of the tree. Mr Magoo's eyesight and ability to judge heights are not good enough to tell to the nearest inch what the height of the tree is. But Mr Magoo can tell some things, like that the tree is not 60 or 6000 inches tall. [Williamson, 2000, 114].

The KK principle that Williamson wishes to challenge runs as follows:

> ($\mathrm{KK}_{Challenged}$): For any given height, if Mr Magoo knows that the tree is not that height, then he knows that he knows the tree is not that height. [Williamson, 2000, 115].

The structure of Williamson's argument is somewhat complicated and irrelevant for my purposes; all that we need to know is that it makes use of a principle involving reliability, which runs as follows:

> ($\mathrm{KReliability}_{Magoo}$) Mr Magoo knows the following: if he knows the tree is not i inches tall, then the tree is not $i+1$ inches tall.

This principle is somewhat complicated. In particular, it is of the form: *Mr Magoo knows X*, where the X in question is the following:

> ($\mathrm{Reliability}_{Magoo}$): If Mr. Magoo knows that the tree is not i inches tall, then the tree is not $i+1$ inches tall.

Because his principle is complex, Williamson works in two stages: first he motivates the simpler principle, ($\mathrm{Reliability}_{Magoo}$), and then he motivates that Mr. Magoo knows that this principle is true. I will focus in this first stage, viz. on his motivation for ($\mathrm{Reliability}_{Magoo}$). Williamson writes:

> To know that the tree is i inches tall, Mr Magoo would have to judge that it is i inches tall; but even if he so judges and in fact it is i inches tall, he is merely guessing; for all he knows it is really $i-1$ inches or $i+1$ inches tall. Equally, if the tree is *i-1* or *i+1* inches tall, he does not know that it is not i inches tall [Williamson, 2000, 115].

Williamson's main defense is thus that if Mr Magoo believed the tree was not i inches tall, when in fact it was $i+1$ inches tall, his success would be due to guesswork.

(As a point of clarification, I should note that Williamson's idea here only applies to the case at hand, in which Mr Magoo formed his belief by looking at the tree from some distance away. Had he been forming beliefs about the tree by some other method, such as carefully measuring the tree, and gotten the measurement right, it would no longer be true to say that his success was merely due to guesswork.)[14]

In order to understand this defense, we must examine what Williamson means by "guess."[15] A plausible interpretation is that by "guess" here Williamson means a belief that could easily have been wrong, given the way it was formed. This meshes with what Williamson does in another place, in which he motivates a similar principle to (Reliability$_{Magoo}$) via the following principle:

> (GeneralizedReliability): If one believes p truly in a case α, one must avoid false belief in other cases sufficiently similar to α in order to count as reliable enough to know p in α [Williamson, 2000, 100].[16]

(GeneralizedReliability) seems to explain Williamson's motivation behind (Reliability$_{Magoo}$). The idea is that if Mr Magoo could easily have falsely

[14] Thanks to an anonymous referee for suggesting I clarify this.

[15] Often, we use the word "guess" to mean a belief arrived at on the basis of a feeling or hunch, as opposed to a process of reasoning. For example, a detective who believed that the butler did it on the basis of a complex reasoning process would not normally be said to have "guessed" that the butler did it, even if the reasoning process employed a number of highly questionable inferences. But Williamson cannot mean "guess" in this sense. For, as the logic of Williamson's argument makes clear, (KReliability$_{Magoo}$) has to apply to a case in which Mr Magoo reasons to the conclusion that the tree is not 665 inches tall, through the employment of some questionable inferences. A similar point is emphasized in [Dokic and Égré, 2009] and [Sharon and Spectre, 2008].

[16] Here is a worry one might have about my interpreting Williamson's argument against (KK$_{Challenged}$) as invoking (GeneralizedReliability): KK principles tend to be defended by internalists, it's not clear that an internalist should accept (GeneralizedReliability), and thus I am interpreting Williamson as invoking a principle which is dialectically suspect. In response, it's worth noting that Williamson explicitly invokes this principle while offering his anti-luminosity argument which is explicitly an argument against internalists so he is clearly happy to make this move that is allegedly dialectically suspect. Also, I argue later in this paper that not all those who accept KK principles need be internalists and that furthermore internalists can endorse an internalist version of KK while accepting (GeneralizedReliability). Thanks to an anonymous referee for pressing this worry.

believed something in a similar case, his belief amounted to guesswork, and this robs him of knowledge.[17]

More formally, we can understand this argument for (Reliability$_{Magoo}$) as running as follows:

> 1. (GeneralizedReliability) If one believes p truly in a case α, one must avoid false belief in other cases sufficiently similar to α in order to count as reliable enough to know p in α.
>
> 2. (SimilarCases$_{Magoo}$) The case in which Mr Magoo believes that the tree is not i inches tall is a sufficiently similar case to the case in which he believes the tree is not $i+1$ inches tall.
>
> –
>
> 3. (Reliability$_{Magoo}$) If Mr Magoo knows the tree is not i inches tall, then the tree is not $i+1$ inches tall.

I can give a parallel argument to defend my reliability principle. First, let me make my task slightly easier. Recall that my reliability principle runs as follows:

> (Reliability$_{Lisa}$). If Lisa knows the answer to a given problem, she does not get the next problem wrong.

It is worth noting that (Reliability$_{Lisa}$) is true regarding all the problems less than or equal to problem 664. This is because in these cases, Lisa gets the next problem right, and thus (Reliability$_{Lisa}$), for those cases, amounts to a material conditional with a true consequent. So all that matters is that I can defend the following:

[17] Here is another piece of evidence that Williamson intends to motivate (Reliability$_{Magoo}$) via (GeneralizedReliability): at one point he writes: "A reliability condition on knowledge was implicit in the argument of section 5.1 [i.e. the argument against KK$_{Challenged}$] and explicit in sections 4.3 and 4.4 [i.e. the sections in which (GeneralizedReliability) is explicitly stated]" [Williamson, 2000, 124]." Williamson clarifies that this reliability condition is (GeneralizedReliability) a page later, writing, "For present purposes, we are interested in a notion of reliability on which, in given circumstances, something happens reliably if and only if it is not in danger of not happening. That is, it happens reliably in a case α if and only if it happens (reliably or not) in every case similar enough to α" [Williamson, 2000, 124].

(Reliability$_{Lisa-Restricted}$) If Lisa knows the answer to problem 665, then she does not get the next problem wrong.

My argument for this principle, which parallels Williamson's, runs as follows:

1. (GeneralizedReliability) If one believes p truly in a case α, one must avoid false belief in other cases sufficiently similar to α in order to count as reliable enough to know p in α.

2. (SimilarCases$_{Lisa}$) The case in which Lisa believes her answer to problem 665 is a sufficiently similar case to the case in which she believes her answer to problem 666.

–

3. (Reliability$_{Lisa-Restricted}$) If Lisa knows the answer to problem 665, then does not get problem 666 wrong.

Williamson's case for (SimilarCases$_{Magoo}$) rests on the details of the story – in particular, on the similarity of Mr Magoo's belief forming processes when he forms the belief that the tree is not i inches tall and when he forms the belief that it is not $i+1$ inches tall. Likewise, my case for (SimilarCases$_{Lisa}$) rests mostly on my story's details; I have sought to make Lisa's belief-forming processes extremely similar from one problem to the next. Recall that both cases involve her looking at long, complicated proofs, proofs that she thinks are right, but where mistakes could easily have escaped her notice.

To further expand: I can specify that Lisa's confidence is roughly the same for both answers. I can also emphasize that she came to believe them them via similar reasoning processes, e.g. that each time she wrote out a series of steps from the problem's starting proposition to the proposition that she believed to be the correct answer, but that both times many steps involved complex rules and that in applying these rules Lisa could easily have made a mistake without noticing. I can further add that the probability of her getting the right answer was about the same for both problems. In short, it seems quite plausible, once certain details are filled in, that the case in which Lisa believes her answer to problem 665 is a sufficiently similar case to the case in which she believes her answer to problem 666.[18] [19]

Given the plausibility of (SimilarCases$_{Lisa}$), probably the best bet for someone who wishes to maintain (Closure$_{Challenged}$) is to reject (GeneralizedReliability).[20] But it's worth noting some implications of doing so.

First, if one rejects (GeneralizedReliability), one thereby rejects Williamson's argument against KK, which relied on it.[21]

[18]One could object to this claim by arguing that in order for one case to be sufficiently similar to another, the proposition believed in each case has to be the same. This idea shows up in the way some people formulate safety conditions, viz:

> S's belief in p is *safe* if and only if S could not easily have falsely believed p in similar cases. [Sosa, 1999, 142].

Here the idea seems to be that the only cases relevant for determining whether a belief in p is safe are ones in the subject believes p. But if we understand similarity in this way, then the case in which Mr Magoo believed that the tree is 665 inches tall would not count as sufficiently similar to the case in which Mr Magoo believed that the tree is 666 inches tall, so (SimilarCases$_{Magoo}$) would be false.

In any case, we can rephrase Williamson's argument, and my own, so as not to need different propositions; in the Magoo case, one can consider a similar situation in which Magoo believes the same proposition but is looking at a tree that has a slightly different height, whereas in my case we can have Lisa deduce the same (rather complex) proposition each time (after first having her memory wiped), with the difference that each time her competence is ever so slightly reduced. Or we can imagine a story in which it is not her competence that is reduced, but rather that we keep the same proposition to be deduced – the true conclusion – and alter the level of difficulty of the problem by altering the starting proposition. Thanks to an anonymous referee for pressing me on this.

[19]One might think that there is a key difference between the two beliefs in the Lisa case, viz. that her belief in her answer to problem 665 is the result of an accurate inference while her belief in her answer to problem 666 is the result of a logical mistake. But there is a similar difference between two beliefs in the Magoo case as well; a belief that the tree is not 665 inches tall is accurate, but a belief that it is not 666 inches results from a mistake. Of course, this is not to say that his belief that the tree is not 665 inches tall amounts to knowledge; Magoo's belief-forming process isn't reliable to pick out the difference in heights. But likewise, it's highly dubious that Lisa's belief in her answer to problem 665 amounts to knowledge; while her inference in this case is logically sound, this is not enough to yield knowledge. Surely not everyone who produces a logically sound inference knows the conclusion, otherwise every belief in a logical tautology I had would amount to knowledge, no matter how complicated the logical tautology was.

[20]I should note that (GeneralizedReliability) is somewhat controversial; see e.g. [Fitelson, 2006], [Ramachandran, 2005]. Thanks to an anonymous referee for suggesting I mention this.

[21]One might worry that, even if Williamson's own motivation for (Reliability$_{Magoo}$) was (GeneralizedReliability), there is another motivation for (Reliability$_{Magoo}$) that cannot be

Second, (GeneralizedReliability) is intuitively plausible; it's plausible that knowledge requires that one's beliefs be reliable, and, in particular, that one not be mistaken in similar cases. So rejecting (GeneralizedReliability) requires rejecting something plausible.

It might be thought that one can soften this blow by replacing (GeneralizedReliability) with a fallibilist version, which says that knowledge requires reliability in most similar cases, rather than all similar cases and thereby capture the intuitions behind (GeneralizedReliability). In other words, one can endorse:

> (GeneralizedReliability$_{Fallibilist}$) If one believes p truly in a case α, one must avoid false belief in *most* other cases sufficiently similar to α in order to count as reliable enough to know p in α.

But this move will not work; I can modify my argument in such a way that I can use this new principle to reach my conclusion. All that I need to do is spell out the story in such a way that Lisa's belief-forming method with regards to problem 665 is not only fallible, but doesn't even produce true belief in most cases. So, for instance, I can add to the story that Lisa is at roughly fifty percent reliability when it comes to the later problems. Then

used to motivate (Reliabilty$_{Lisa}$). In particular, we can motivate (Reliability$_{Magoo}$) as follows: whatever you say about Mr Magoo's epistemic standing with regard to the tree he's observing, he cannot know to the nearest inch (using whatever method he can at the moment) that the tree is not at a certain height when its actual height is one inch from that. This argument relies, then (or so the challenge goes), on the quality of Magoo's evidence not on reliability considerations.

In response, it's worth emphasizing that (Reliability$_{Magoo}$) says more than merely that Mr. Magoo cannot know to the nearest inch what the height of the tree is. To see this, imagine Mr Magoo had a device that precisely measures the tree's height but reports the tree's height in a peculiar way: it either announces that the tree is below 65.5 inches, or it announces that it is between 65.5 and 165.5, or between 165.5 and 265.5, or between 265.5 and 365.5, and so on. Then, if Mr Magoo had this device, he wouldn't be able to tell the height to the nearest inch. But (Reliabilty$_{Magoo}$) would be false, because if the device reported that the tree was between 665.5 and 765.5, and the tree was in fact 666 inches, Mr. Magoo would know that it was not 665 inches, even though it was in fact 666.

Of course, this still leaves open whether there is some way to appeal to facts about evidence in a way that supports (Reliability$_{Magoo}$) without supporting (Reliability$_{Lisa}$). I think the answer is that there is not; in the fifth section of this paper, I give some reason to think Mr. Magoo and Lisa are in a similar evidential position, both with regards to the evidence they have, and their ability to appropriately base their beliefs on that evidence. Thanks to an anonymous referee for pressing this worry.

(Set-Up$_{Lisa}$) plus (Closure$_{Challenged}$) will entail that she knows the answer to problem 665 and (GeneralizedReliability$_{Fallibilist}$) will entail that she doesn't, seeing as, at this point, her belief-forming process isn't even mostly reliable. In short, once these details are added, my argument can use a fallibilist version of (GeneralizedReliability) to generate its conclusion. So trying to avoid my argument through rejecting (GeneralizedReliability) and related principles comes with big intuitive costs.

3 Can a higher-order closure principle escape my argument?

My argument focuses on a particular closure principle – (Closure$_{Challenged}$) – and so a natural question arises: are there other closure principles that can escape my argument? I will now provisionally discuss this question, working, in order, through three closure principles: a higher-order one, an externalist-friendly one, and an internalist-friendly one.

Here are my provisional conclusions: the higher-order principle cannot escape my argument. There are externalist-friendly and internalist-friendly principles that escape my argument, but they have what appear to be significant flaws. And parallel things holds for KK principles.

Recall that my argument challenged the following closure principle:

> (Closure$_{Challenged}$) For a given problem, if Lisa knows the starting proposition to that problem and believes another proposition that is entailed by it, then she knows it.

One might think that (Closure$_{Challenged}$) is false for the following reason: when it came to problem 665, even though Lisa knew the starting proposition and even though the starting proposition entailed Lisa's answer, Lisa didn't *know* that it entailed her answer and this kept her from knowing her answer. This inspires the following thought: perhaps a closure principle whose antecedent required not merely that p entails q, but that the subject knows that p entails q, would be immune to my argument. Here is such a principle:

> (Closure$_{HigherOrder}$): For every proposition p and q, if one knows p and knows that p entails q, then one knows q.

I am doubtful that this strategy can work; I can (arguably) modify the Lisa story so that my argument targets this principle as well. In other words, the premises of my original argument against closure referenced a story regarding Lisa; if these premises were all correct, then the original closure principle I was challenging – (Closure$_{Challenged}$) was false. By altering the story slightly, I can offer a new argument which is like the original one, except that it references this new story, and which shows that (Closure$_{HigherOrder}$) is false.

Recall that my argument is a reductio; it showed that when you combine several principles, one of which is (Closure$_{Challenged}$), you get to an absurd conclusion. And, in more detail, the way the argument worked is that there were two principles, (Set-up$_{Lisa}$) and (Closure$_{Challenged}$) that together entailed, for each of problems 0 through 665, that Lisa knew the conclusion she believed to be the correct answer to that problem. And there was another principle, (Reliability$_{Lisa}$) that entailed that if she knew the answer to problem 665, then her answer to problem 666 was true. But because it was not, one of these principles had to be false.

To generate my new argument, I will modify the story. In particular, for each problem, in addition to "the starting proposition", Lisa will also receive a list of other propositions, which I'll call "the intermediate propositions". She can use these in trying to determine whether the ending proposition is true or false. Each intermediate proposition will be of the form "P entails Q". And, for each problem, there will be one particularly helpful intermediate proposition – although Lisa won't be told which one of the intermediate propositions it is. What makes it particularly helpful is that its antecedent (P) is the starting proposition for that problem and its consequent (Q) is either the ending proposition for that problem or the negation of the ending proposition. Thus, if she selects the correct intermediate proposition, she can use it to solve the problem by either proving the ending proposition true or proving it false. Again, she receives good assurance that the intermediate propositions are true and thus knows them.

I will add one more twist to the new story. In the new story the starting proposition and intermediate propositions – at least for the later problems – will be rather long and complicated. For instance, suppose by the time she gets to the problems in the 600s, the sentences that express the starting propositions and intermediate propositions are each over a page long. In such a case, Lisa might be unreliable at comparing the two sets of formula and confirming that they match, and thus prone to making mistakes. And, as before, the problems gradually get harder and harder and she eventually makes a mistake at problem 666.

The modified argument, then, uses the same structure as the previous one.

First, as before, there's a principle I'll call Set-Up$_{HigherOrder}$, that invokes information provided by the set up to my argument:

> (Set-Up$_{HigherOrder}$). Lisa knows the starting proposition and the intermediate propositions of each problem and for all but the last problem, the answer Lisa believes to be correct is indeed entailed by the starting proposition of that problem.

Again, then, as with the earlier argument, we have a reductio argument. (Closure$_{HigherOrder}$) and (Set-Up$_{HigherOrder}$) together entail that Lisa knew the answers to problems 1 through 665. Thus, (Reliability$_{Lisa}$) entails that she didn't get problem 666 wrong. But she does get it wrong. And thus, assuming the other principles are correct, (Closure$_{HigherOrder}$) is false.

Let me consider an objection to this new argument. The objection starts by noting that the story I just presented requires that two things be true. First, it requires that Lisa know the starting propositions and intermediate propositions for the various problems. Second, it requires that she use a process to solve the problems that is somewhat unreliable, at least when it comes to the problems high in the 600's.

One might challenge that the story I presented is impossible. In particular, one might think that coming to know the starting and intermediate propositions is more challenging than comparing and confirming that the starting proposition matches the antecedent of the intermediate proposition and the ending proposition (or its negation) matches its consequent. And thus, one might challenge the possibility of Lisa's knowing the starting and intermediate propositions while being unreliable at identifying if starting, intermediate and ending propositions match in appropriate ways.

In response to this challenge, an analogy might be helpful. Suppose that I have two pictures of a crime scene. Suppose I know that one is from before the police went through the crime scene and that the other is from after they went through the crime scene. I know that if the two pictures are the same, then the police have not altered the crime scene. But I might mess up the task of confirming that the two pictures are the same. This is especially likely to happen if the pictures contain a lot of information.[22]

[22] I should also note that it's possible for the difficulty of this task to increase gradually. For instance, there can be a series that starts with simple pairs of drawings for which it's easy to tell if there are any differences between them and moves gradually to complex pairs of drawings for which it's really difficult to tell if there's any difference between them.

Likewise, suppose I have two long formulas, the second of which is a conditional, and that I know they're both true. Then I know that if the first is the same as the antecedent of the second, then the consequent of the second is true. But if they're rather complicated, I may well be unreliable at checking that the first is the same as the antecedent of the second.[23]

This reasoning suggests that (Closure$_{HigherOrder}$) faces a similar problem to the problem that faced (Closure$_{Challenged}$) and that my argument – if successful in targeting (Closure$_{Challenged}$) – can also be modified to target (Closure$_{HigherOrder}$).

4 Can principles framed in terms of competent deduction escape my argument?

Another sort of closure principle is one that Williamson himself endorses:

> (Closure$_{CompetentDeduction}$) Knowing p, competently deducing q, and thereby coming to believe q is in general a way of coming to know q.[24]

Similar things hold for the task of comparing formulas of the form "A", "A entails B" and "B" to confirm matches. If A and B are simple, it's easy to compare; if they are somewhat long, it's more difficult, and if they are extremely long it's more difficult still. Thanks to an anonymous referee for pressing me on this.

[23]One might object: "the task of knowing the implications of complex formulas seems more demanding than the task of matching two complex formulas. So if the subject knows A entails B, it seems plausible that she would be able to match A with the occurrence of A in A entails B." Two responses. First, even if one task is more demanding than a second, it doesn't follow that if one is able to competently accomplish the first, one is able to competently accomplish the second. For instance, making waffles from scratch is more demanding than reheating frozen ones. Nonetheless, I may competently accomplish the first and fail to competently accomplish the second. Maybe reheating frozen things causes me great anxiety and leads to my making mistakes. Maybe, for whatever reason, I just never learned how to reheat frozen foods.

Also, it's not obvious that knowing A entails B will always be more demanding than matching formulas. For instance, suppose that first one comes to know that A entails B through being told this by one's math teacher. Next, one has to match A with the occurrence of A in A entails B. Seemingly, in such a case, coming to know A entails B could be an easier task than the matching.

[24]Strictly speaking, Williamson offers a multi-premise version, writing "Knowing p_1, ..., p_n, competently deducing q, and thereby coming to believe q is in general a way of coming to know q." [Williamson, 2000, 117]. I have offered a version restricted to a single premise so as to avoid unnecessary complications.

As just noted, my argument has focused on a different principle:

> ($\text{Closure}_{Challenged}$) For any pertinent propositions p and q such that p entails q, if Lisa knows p then she knows q.

The fact that Williamson doesn't explicitly endorse ($\text{Closure}_{Challenged}$) yields Williamson the following option: he can say that ($\text{Closure}_{Challenged}$) doesn't follow from his preferred closure principle, ($\text{Closure}_{CompetentDeduction}$). If this is right, then even if ($\text{Closure}_{Challenged}$) is false, it does not follow that there is any sort of problem for ($\text{Closure}_{CompetentDeduction}$).

One way Williamson can argue that ($\text{Closure}_{Challenged}$) does not follow from ($\text{Closure}_{CompetentDeduction}$) is by arguing that some instances of ($\text{Closure}_{Challenged}$) do not involve competent deduction. That is, he would need to hold that some of Lisa's work, even with regards to the problems she got right, did not amount to competent deduction. A natural way for Williamson to do this is to say that if one competently deduces something, one could not easily have made an inference that is not truth preserving in a similar situation. This would allow him to say that Lisa's work on problem 665 did not amount to competent deduction, because she did in fact make a mistake in a similar situation, namely her work on problem 666. While Williamson does not spell out competent deduction in this way in *Knowledge and its Limits*, he does do so in a later work, writing:

> If in any case α one believes a conclusion q on a basis b, which consists of competent deduction ... then in any case close to α in which one believes a conclusion q* close to q on a basis b* close to b, b* consists of a truth-preserving deduction ... [Williamson, 2009, 326].[25]

[25]It is worth noting that such a move incorporates a fairly controversial notion of competence, and thus may not appeal to others who are also happy to invoke competence but wish to understand it differently. It holds that in order to competently deduce something, one could not easily have made a mistake in a similar situation. But some think that one can competently come to believe something, even though one could easily have made a mistake in a similar situation. For instance, Ernest Sosa discusses the following case:

> You see a surface that looks red in ostensibly normal conditions. But it is a kaleidoscope surface controlled by a jokester who also controls the ambient light, and might as easily have presented you with a red-light+white-surface combination as with the actual white-light+red-surface combination [Sosa, 2009, 31].

By adding in the word "competent" Williamson can avoid the conclusion of my argument. That said, requiring competent deduction, rather than deduction simpliciter, yields two problems.

First, a more minor one: what it means to be "competent" must be carefully spelled out, especially given, as just noted in a footnote, that its meaning seems to be rather different from the way in which it is normally used. If left intuitive, the principle will be hard to apply. Williamson does spell out a version, but it leads to some fairly unpalatable consequences. For instance, he has to endorse the view that there can be events with the following two features: (i) it's highly physically probable that they'll occur and (ii) they can't easily happen [Williamson, 2011].[26]

Second, a somewhat less minor one: there is an apparent dilemma regarding how competent one must be in order to be "competent": must one be perfectly competent, or merely mostly competent?

If one must be perfectly competent, then, in order for (Closure$_{CompetentDeduction}$) to yield the result that we know things via deduction, it must be the case that some of us are sometimes perfectly competent. That is, it must be the case that some of us, in deducing things, must be such that we couldn't easily have made mistakes. There are some who are willing to endorse this sort of conclusion, such as Williamson himself, but such people take on some

---

Sosa thinks in this case you have an apt belief – that is, one that succeeds through the exercise of that competence – even though your belief could easily have been false. [Sosa, 2009, 35-6].

Similar things are often said about fake barn cases, in which one looks at a real barn in an area full of barn facades. That is, such cases are often alleged to be cases in which one's current belief that there is a barn before one is due to a competence, even though in similar cases one would have had a false belief. See e.g. [Sosa, 2009, 96].

[26] Very briefly, the reason is because Williamson understands competence in terms of not easily having been mistaken, where this in turn is understood in terms of not being mistaken in any nearby worlds. He then gets the controversial commitments through endorsing the following three claims, each of which is very difficult to avoid: (1) Not being mistaken in nearby worlds is closed under disjunction; if there are no nearby worlds in which is one mistaken that P and no nearby worlds in which one is mistaken that Q, then there are no nearby worlds in which one is mistaken that P or mistaken that Q. (2) There are a number of events each of which has low physical probability, is independent from the others, and which you know didn't happen. For instance, one such event is that the interior of my desk just rearranged itself a moment ago in such a way that it's no longer a desk thanks to weird quantum phenomena. (3) The disjunction of a sufficiently large number of physically independent events can be quite probable even if their individual physical probabilities are each quite low.

pretty serious burdens as a result of endorsing this claim. For instance, as just noted, Williamson himself, in order to maintain consistently, is led at a result of endorsing this sort of claim to endorsing certain deeply controversial claims regarding probability and what could easily have happened.

The other horn of the dilemma is to say that competent deduction does not require perfect competence. The trouble with this horn is that, if we understand competently deduction in this way, then the version of (Closure$_{CompetentDeduction}$) that we get will fall prey to my argument against closure.

In particular, suppose we understand my story involving Lisa in such a way that Lisa's deductive process in problem 665 is mostly reliable, but she still gets problem 666 wrong. Then (Closure$_{CompetentDeduction}$) will deliver the result that she knows the answer to problem 665 – after all, her process was mostly reliable, and thus counts as a "competent deduction." At the same time, (GeneralizedReliability) will deliver the result that she fails to know the answer to problem 665, after all, she makes a mistake on problem 666 and thus is not perfectly reliable when it comes to problem 665. So, assuming (GeneralizedReliability) is right, it will follow that (Closure$_{CompetentDeduction}$) – with "competent deduction" understood in terms of being mostly reliable – is false.

Let me put the point a different way. We are considering the possibility of endorsing a fallibilist closure principle – one on which competent deduction is understood in terms of being mostly reliable. But the key principle that I invoked in my argument – (GeneralizedReliability) – is an infallibilist principle; it requires, in order to know, that one not make a mistake in any sufficiently similar circumstance. And this principle is at cross purposes to the fallibilist closure principle.

In other words, I can grant that Lisa's process in problem 665 is mostly reliable. But thanks to the difference between "mostly reliable" and "perfectly reliable", it's possible for her to use a mostly reliable process in problem 665 and then get problem 666 wrong, even if she's using a sufficiently similar process in the two problems. In short: a mostly reliable process can still go awry.

It might seem that the reason I was able to give the argument I just gave is that (GeneralizedReliability) requires perfect reliability in order to have knowledge. So it might be thought that the person endorsing a fallibilist version of closure has another move available: in addition to endorsing a fallibilist version of closure, she could also reject (GeneralizedReliability) and replace it with a fallibilist reliability principle. And this would presumably keep my argument from going through. In other words, one could respond that my response worked because we were combining a fallibilist closure principle with an infalliblist principle connecting knowledge and reliability, and that if one endorsed instead a fallibilist principle connecting knowledge and reliability, there would no longer be an argument available against the fallibilist version of closure in question.[27]

But this move will not work. Even if one relaxes to a falliblist principle linking knowledge and reliability, one still cannot endorse a fallibilist closure principle. Put more formally: if one endorses the claim that knowledge, in general, merely requires that one's belief-forming process be mostly reliable, as opposed to perfectly reliable, then one still cannot endorse (Closure$_{CompetentDeduction}$) – with "competent deduction" understood in terms of being mostly reliable. The problem is the following: even if one is mostly reliable in believing a proposition and mostly reliable in deducing a conclusion from this proposition then one's overall process can still fail to be mostly reliable. To take a toy example: I might be seventy percent reliable in believing p, seventy percent reliable in deducing q from p and only forty-nine percent reliable in believing q. So even if one moves to a fallibilist view regarding the connection between knowledge and reliability, one should still reject (Closure$_{CompetentDeduction}$) – with "competent deduction" understood in terms of being mostly reliable.

In sum, then, if "competence" requires perfect reliability, then there are difficulties applying (Closure$_{CompetentDeduction}$) to humans, whereas if "competence" requires less than perfect reliability, then (Closure$_{ComptentDeduction}$) can still be targeted by my argument.

To summarize: Williamson's alternative to (Closure$_{Challenged}$) avoids my argument against closure, but it faces some problems as a result.

5 Can internalist-friendly principles escape my argument?

I would also like to look at internalist-friendly closure principles and see if they can escape my argument. Typically internalists defend closure principles as follows: they note that certain properties, such as having good evidence for, are closed under entailment.[28] Internalists tend to think that if you know that p, then you have good evidence for p, and thus will have good evidence

[27] Thanks to an anonymous referee for suggesting this response.

[28] I will leave "good" in "good evidence" unanalyzed. There are various ways of analyzing "good evidence" on which having good evidence is closed under entailment. One example: evidence for p is good evidence if p is highly probable given this evidence.

for anything that p entails. Meanwhile, they tend to think knowledge is a matter of basing your belief on the evidence in the right sort of way.[29] Because of this, so long as you believe some proposition entailed by p in the right sort of way, you will know this proposition.[30] This sort of reasoning motivates the following sort of principle:

> (Closure$_{Internalist-friendly}$) If S knows p and p entails q then if S bases S's belief in the right way, S knows q.

Just a quick word about basing one's belief in the right sort of way: it is widely recognized that one can have strong evidence for some proposition, and yet believe it for a different reason, and because of this fail to have knowledge. For example, as Conee and Feldman write in their 1985 article, "Evidentialism":

> Suppose Alfred is justified in believing *p*, and justified in believing if *p* then *q*. Alfred also believes *q*. ... if Alfred's basis for believing *q* is not his evidence for it, but rather the sound of the sentence expressing *q*, then it seems ... clear that there is some sense in which this state of believing is epistemically 'defective' – he did not arrive at the belief in the right way. [Feldman and Conee, 1985, 24-5].

There are various views about what it takes to base one's belief in the right sort of way, and I will not endorse a particular one. But the key idea is that there are two different sorts of reasons one could lack knowledge, the first because one lacks evidence, the second that one has evidence, but fails to base one's belief in the right sort of way. In the second case, unlike the first, one can have the same evidence, but base differently, and thereby know.[31]

As with Williamson's externalist Closure principle, someone who endorses (Closure$_{Internalist-friendly}$) can accept the conclusion of my parallel argument involving Lisa, but then contend that this argument doesn't threaten (Closure$_{Internalist-friendly}$). Just as a reminder, my parallel argument challenged:

[29] Usually internalists add in that one must also avoid Gettier cases, but I will leave out this detail, because it is irrelevant for this discussion. See e.g. [Feldman and Conee, 1985].

[30] For such defenses, see e.g. [Klein, 1995, 219], [Luper, Winter 2011], [Moretti and Shogenji, 2017, 7], [Stine, 1975, 250], [Wang, 2014, 1130].

[31] Again, setting aside Gettier cases.

> ($\text{Closure}_{Challenged}$): For any given problem, if Lisa knows the starting proposition to the problem and believes another proposition entailed by it, then she knows it.

Someone who endorses ($\text{Closure}_{Internalist-friendly}$) can say that ($\text{Closure}_{Internalist-friendly}$) does not entail ($\text{Closure}_{Challenged}$). In particular, she can say that some of Lisa's inferences did not involve her basing her belief in the right sort of way. In particular, the defender of ($\text{Closure}_{Internalist-friendly}$) can allege that Lisa's inference in problem 665 involved her basing her belief in the wrong way because it was not a careful enough deduction, given how difficult the problem was. While she did have good evidence for (Start_{665}) and thus good evidence for (End_{665}), she failed to deduce carefully enough. That said, there was good evidence for (End_{665}) available. (End_{665}) follows from (Start_{665}), which she had good evidence for. So had Lisa deduced more carefully, she would have known her conclusion; someone with her evidence could have known (End_{665}).

It's worth noting that in escaping the argument, this principle faces some costs. For instance, the principle doesn't deliver the result that if one knows a proposition and this proposition entails a second proposition, then one knows the second proposition. Rather, it only delivers this result for people who base their belief in the second proposition in the right way, which may be a rather difficult matter, either because they lack the requisite deduction skills or because they lack the reliability requisite to successfully perform the deduction in other similar cases.[32]

Let me expand on this point a bit. First, let me give an example of a case in which it is difficult to base one's belief in the appropriate way. Here's the example: there are some mathematical axioms that are themselves easy to grasp but which entail some complex and sophisticated theorems, theorems that are very difficult to prove from the axioms, or, indeed, even to understand. The internalist closure principle says that so long as one knows the axioms, then if one bases one's belief in these theorems in the appropriate manner, one will know them. But in this sort of case it seems very difficult to base one's belief in the appropriate way, given how difficult it is to prove or even grasp the theorems.

[32]Note: at this point in the paper, we are assuming that my parallel argument is successful and trying to assess the consequences thereof, thus we are assuming for the sake of argument that (Generalized Reliability) is true and thus that knowledge requires reliability in other similar cases).

With this sort of example in mind, we can perhaps better see the limitations of the internalist-friendly closure principle. The internalist-friendly closure principle says that if you know that p and p entails q then you are part of the way to knowing q; so long as you base your belief in q in the appropriate way, then you will know it. But the cost that I am pointing out is that it can be very difficult indeed to base your belief in the appropriate way. And thus it's not clear whether we should care that much about this intermediate state – the state of being such that if one bases one's beliefs in the right way, then one will know. And thus, furthermore, it's not clear that the internalist-friendly closure principle, which says that if one knows p, then one will have reached this intermediate state, is telling us something of much significance, given that the state it focuses on is itself of dubious significance.

Here's an analogy: suppose someone tells me that he can help me become very rich. In particular, he tells me, he can help me develop a certain very special kind of mindset and can teach me some important principles centered around this mindset. What makes the mindset so special is that if one has developed that mindset, then one is one step away from being very rich. In other words, once one develops the mindset, one need take only one more step, and then one will be very rich. Furthermore, he tells me, if I buy his book and take his courses, I will be able to develop the mindset. So far, things sound really good. But if I then discover that this one additional step – the step that comes after developing the mindset – can be extremely difficult to accomplish, then I might start to wonder why I should care so much about the mindset he recommends I cultivate. And I might further begin to doubt that the principles he talks about – principles that center on this mindset – are of much significance.[33]

[33] I should perhaps briefly relate the point I'm making to Williamson's criticisms of internalism. Internalists tend to claim that there are certain properties, such as phenomenal properties, that are of special significance. In particular, they say one has some kind of special access to facts regarding these properties and thus, if these properties are instantiated, then, so long as one bases one's beliefs regarding these properties in the appropriate way, one's beliefs will amount to knowledge. In other words, when these properties obtain, one is in a special sort of intermediate state and this intermediate state is such that if one bases one's belief in the right way, then one will know.

But Williamson argues that this appropriate basing can be quite difficult. In particular, he argues that these properties can be instantiated without one's being in a position to know that they are [Williamson, 2000, 24]. Given this argument, Williamson thinks it is dubious that the intermediate state internalists focus on and the properties related to

6 Parallel things hold for KK principles

In the last two sections, I looked at Closure principles that escape my argument, but appear to have significant flaws. In this section, I will argue that a parallel thing is true of KK principles and Williamson's argument.

6.1 Can principles framed in terms of competent deduction escape Williamson's argument?

Let us start with competent deduction. We can formulate a KK principle that is parallel to (Closure$_{CompetentDeduction}$). As we shall see, the parallel principle can be motivated in similar ways, it faces similar problems, and it escapes Williamson's argument against KK.

Actually, there are multiple ways to formulate such a principle because (Closure$_{CompetentDeduction}$) makes reference to a particular way of coming to know an entailment: deducing it. There are multiple options for a parallel term to "deduction." I will use "validation." Here is my principle:

> (KK$_{CompetentValidation}$) Knowing p, competently validating that one knows p, and thereby coming to believe that one knows p is in general a way of coming to know that one knows p.

Here are some quick notes about this principle. I am not assuming that in order to competently validate that one knows, one has to know that one

---

it, are of much significance. As Williamson says in another context, while criticizing the views of Earl Conee:

> Conee asserts that phenomenal qualities provide a comfortable cognitive home ... because they are always there among our ultimate evidential resources. But what use to you as evidence is a phenomenal quality when you are not even in a position to know that you have it? Conee reassures us that phenomenal qualities are still 'known to us by acquaintance', but in Conee's sense you can know x by acquaintance even if you are not in a position to know that x exists. [Williamson, 2005, 8].

In short, then, Williamson questions the significance of the sort of state that internalists care about – the sort of state one occupies when one instantiates some of the properties in question. This is so despite its being the case that if one occupies this sort of state then one will have knowledge, so long as one bases one's beliefs in the appropriate way. And the reason Williamson is dubious about the significance of this sort of state is because he thinks one can occupy such a state and still be very far from knowledge, given how difficult it can be to appropriately base one's beliefs.

knows. Rather, the idea is that competent validation, in conjunction with knowledge, yields knowledge that one knows. Compare: suppose I am a lawyer responsible for reviewing some cases and deciding whether our firm should take them on. I sort through a pile of cases, marking down those that I think we should take on, and those I think we shouldn't. I am somewhat shaky on health law but good at immigration law, so after sorting through the cases, I divide them into two piles, one: a pile of the health law cases, which someone else should look at, and another: a pile of the immigration law cases, for which I'm confident in my judgment. If I sort a case into the immigration law pile, and my judgment regarding that case is that we should take it on, then I can be confident we should take it on, even though not every case in the immigration law pile is one we should take on (some are ones I correctly believe that we should reject). Likewise, in general if one competently validates that one knows p, and one does, then one knows that one knows p, even if not all propositions that one competently validates that one knows are in fact known.[34]

Williamson's argument has focused on a different principle:

> ($KK_{Challenged}$) For any given height, if Mr Magoo knows that the tree is not that height, then he knows that he knows the tree is not that height.

Like Williamson, I could argue that this principle doesn't follow from my preferred principle.[35] In particular, I can argue that some applications of $KK_{Challenged}$ do not involve competent validation that Mr Magoo knows some proposition. For instance, I could contend that Mr Magoo's coming to believe that he knew that the tree was not 665 inches tall did not involve competent validation. Like Williamson, I could spell out competent validation in terms of what could happen in similar cases, holding that a validation is competent only if it is such that one could not have made a mistake in a similar case – that is, one could not have mistakenly come to believe that one had knowledge in a case in which one did not.[36] What is especially worth stressing is that

[34]Also, I'm not assuming that validation is an inferential or quasi-inferential process; I'm leaving open that one can validate that one knows p directly, without inference.

[35]For those who employ a similar strategy but with a different KK principle, see e.g. [McHugh, 2010].

[36]More formally, this is to endorse:

> Competent Validation. If in any case α one believes that one knows that p

for both Williamson and me, a lot of work will be done by the notion of "competence" and here we seem to be on equal footing.

It's worth noting that my parallel principle faces parallel problems to those that Williamson's faced. First, what it means to be "competent" must be carefully spelled out. Second, there is an apparent dilemma regarding how competent one must be in order to be "competent." If one must be perfectly competent, then it might seem that this principle never applies to humans – none of us is perfectly competent; each can easily make mistakes. If one must merely be mostly competent, then the principle faces the same problems as ($KK_{Challenged}$): I can be mostly competent and still such that I could easily have made a mistake.

6.2 How an internalist-friendly KK principle can avoid Williamson's argument

A similar parallel holds regarding internalist-friendly closure principles and internalist-friendly KK principles. While Williamson does not endorse an internalist closure principle, there are plenty of internalists who do, and thus unless they wish to give up on this sort of principle, they should be very hesitant about endorsing the idea that Williamson has managed to show that key internalist KK principles are false. Instead, these internalists can note that there are parallel internal KK principles to their preferred internal closure principle, and that these parallel principles get around Williamson's argument. For example:

> ($KK_{Internalist-friendly}$) If S knows p then if S bases S's belief in the right way, S knows that S knows p.

As before, an internalist who endorses ($KK_{Internalist-friendly}$) can accept the conclusion of Williamson's argument, but argue that it doesn't threaten ($KK_{Internalist-friendly}$). Such an internalist can allege that for some i, Mr Magoo's belief that the tree is not i inches tall did not involve his basing his belief in the right way. He was not careful enough. If he was a little more

> on a basis b, which consists of competent validation, then in any case close to α in which one believes that one knows a proposition p* close to p on a basis b* close to b, one knows that p*.

Thanks to an anonymous referee for suggesting that I include this.

careful in how he based his belief, he would have had knowledge. And this is why ($KK_{Challenged}$) is false.[37]

I will spell this idea out a little more fully momentarily. But the case involving Mr Magoo is somewhat subtle because it involves perception; I will first discuss a simpler case regarding introspection and then apply the lessons to the case of Mr Magoo. To do this I will discuss an argument Williamson gives that challenges the idea that one is always in a position to know one's mental states [Williamson, 2000]. I have tweaked his argument slightly so as to avoid some confusing details. First I will state this argument. Then I will show how an internalist can argue that its conclusion is not threatening to an introspective analog of ($KK_{Internalist-friendly}$). Then I will apply the lessons there to the case of Mr Magoo.

First the argument. Take a particular shade of red, call it R_{666}. Suppose that Clifford has a series of phenomenal experiences as of looking at a series of shades, R_0 through R_{666}. Let us suppose that Clifford isn't good enough at introspecting that he can tell the difference between any two of the phenomenal experiences, but he can tell the difference between R_0 and R_{666}. There is some difference between neighboring shades– what it's like to be experiencing neighboring shades is slightly different – but Clifford can't pick up on it. Let us call the relevant principle that one is in a position to know one's mental states ($\text{Introspection}_{Challenged}$). Here it is:

> ($\text{Introspection}_{Challenged}$): If Clifford is having the phenomenal experience as of R_i, then he knows he is having the phenomenal experience as of R_i.

The argument that Williamson can give against this principle will run exactly parallel to his original argument against ($KK_{Challenged}$), but the details don't matter for our purposes.

Suppose an internalist wishes to show that this argument doesn't challenge her preferred introspection principle. She can endorse a parallel of ($KK_{Internalist-friendly}$):

[37] Relatedly: one way of putting Williamson's point from his argument against KK is that safety doesn't iterate; that is, one can safely believe that p while failing to safely believe that one safely believes. The internalist response here is that safety iterates so long as one believes in the right way. That is, if one safely believes that p and one bases one's belief that one safely believes that p in the right way, then one safely believes that one safely believes that p.

(Introspection$_{Internalist-friendly}$): If S is having a phenomenal experience P, then if S bases S's belief in the right way, then S knows that S is having P.

The internalist can allege that (Introspection$_{Challenged}$) does not entail (Introspection$_{Internalist-friendly}$) and thus the argument against (Introspection$_{Challenged}$) doesn't target her preferred principle. In particular, the internalist can allege that Clifford isn't basing his beliefs about phenomenal experiences in the right sort of ways. If he was better at introspecting, and thus based his beliefs more carefully, he would know the relevant proposition. That is, Clifford has the relevant evidence available, thanks to a difference in his phenomenal states. If he based his beliefs on the relevant evidence, he would know the propositions in question.[38]

Now let us apply these lessons to the case of Mr Magoo. Take an i such that Mr Magoo knows the tree is not i inches tall but fails to know that the tree is not $i+1$ inches tall. (There has to be such an i because Mr Magoo knows the tree is not 0 inches tall and fails to know it is not 666 inches tall). Here we have a decision point in how we understand the case. On one understanding, the tree would look differently to Mr Magoo if it were i inches as compared with $i+1$ inches, he's just not good enough at introspecting that he can pick up on this difference. That is, what it's like to be Mr Magoo would shift slightly if the tree were i inches as compared with $i+1$ inches, but so slightly that he wouldn't be able to notice. On this understanding of the case, it can be handled just as above. Mr Magoo has the evidence that he knows the tree is not i inches tell – this phenomenological evidence – he just isn't careful enough in his basing.[39] If he had been better at introspecting his phenomenological states, he would have been able to

[38]This is important because, as I shall briefly argue in this footnote, (Introspection$_{Internalist-friendly}$) can be used to establish conclusions that Williamson wanted to reject. (To fully establish this conclusion would take a more thorough argument than I will give here, but I hope my remarks will at least provisionally establish a difficulty for Williamson.) In particular, Williamson wanted to reject the idea that we have a cognitive home – that one is "guaranteed epistemic access to one's current mental states" [Williamson, 2000, 93]. If (Introspection$_{Internalist-friendly}$) is right, then we do have a special kind of access to our phenomenal experiences – so long as we base our beliefs in the right sort of way, we can know that we are having them. So in failing to attack (Introspection$_{Internalist-friendly}$), Williamson leaves himself vulnerable to those who wish to use it to establish that we do have a cognitive home.

[39]This sort of strategy is taken in [Smithies, 2012].

notice a difference between the two cases, and thus could have based his belief in such a way so as to know that he knew the tree was not i inches tall. Put slightly differently, on this understanding, the difference between Mr Magoo and someone who was better at forming accurate beliefs regarding tree heights is not that Mr Magoo has different evidence from the person, it's just that the other person can use the evidence – the slight differences in phenomenology – to distinguish certain cases while Mr Magoo cannot.

On another understanding, the tree would not look any different if it were i inches as compared with $i+1$ inches. On this understanding, what it's like to be Mr Magoo wouldn't shift slightly if the tree were i inches as compared with $i+1$ inches. His perceptual faculties are sufficiently non-discriminating that they wouldn't pick up any differences. On this understanding of the case, an internalist will deny that Mr Magoo knows that the tree is not i inches tall and fails to know that it is not $i+1$ inches tall. If the tree looks the same in the two cases, an internalist will say that Mr Magoo can only know the tree is not i inches tall if he also can know the tree is not $i+1$ inches tall.

In short, no matter which way we understand the case, an internalist can deny that Williamson's argument challenges her preferred KK principle, ($\mathrm{KK}_{Internalist-friendly}$) in just the way an internalist could deny that my argument challenged her preferred closure principle.

This is important because ($\mathrm{KK}_{Internalist-friendly}$) can be used to establish conclusions that Williamson wanted to reject. In particular, Williamson thinks that the surprise test paradox arises thanks to a failure of KK and other similar principles.[40] Some background: the surprise test paradox involves an ingenious argument that seems to show that if students hear from a teacher that they will have a surprise test, then they can deduce that they will not have the test. Williamson notes that in order for this argument to be successful, not only must the students know that they will have the test (on the basis of the teacher's announcement), they must also know that they know it. But because he thinks KK fails, he thinks there's no reason to hold that they can know this [Williamson, 2000, 135-143]. But if ($\mathrm{KK}_{Internalist-friendly}$) is right, then there is some reason to hold that they can know this – they just

[40]Why do I say "KK and other similar principles?" The reason is that, as Williamson notes, there are multiple ways of formulating the surprise test paradox – all of the versions use principles about iterating knowledge, but there are several closely-related options to choose from, only one of which is KK. [Williamson, 2000, 140]. Thanks to an anonymous referee for pressing me to clarify this.

need to base their beliefs in the right way. So assuming ($KK_{Internalist-friendly}$) is right, there is some reason to believe Williamson has failed to solve the surprise test paradox.

Again, it's worth noting that in escaping the argument, this principle faces some costs. For instance, the principle doesn't deliver the result that everyone who knows something that entails a second proposition knows the second proposition. Rather, it only delivers this result for people who base their belief in the second proposition in the right way, which may be a rather difficult matter, either because they lack the sophisticated introspection skills required or because they lack the reliability requisite to successfully perform such introspection in other similar cases.

7 Conclusion

Just to briefly summarize, then: in this paper I have constructed an argument against a closure principle and defended it via a general principle regarding reliability, a principle that is also invoked in an argument from Timothy Williamson against KK. I then provisionally examined whether there are any closure principles that can escape the argument. I found two that could: one that invoked competent deduction and one that invoked proper basing. I noted some flaws of the two principles. I also provisionally argued that there are parallel KK principles that can be motivated in similar ways to the ways the two closure principles are motivated and which can yield conclusions that Williamson himself wished to avoid.[41]

References

[Baumann, 2012] Baumann, P. (2012). Nozick's defense of closure. In Becker, K. and Black, T., editors, *The Sensitivity Principle in Epistemology*, pages 11–27. Cambridge University Press.

[Castañeda, 1970] Castañeda, H.-N. (1970). On knowing (or believing) that one knows (or believes). *Synthese*, 21(2):187–203.

[41] Thanks for comments to Nevin Climenhaga, Amy Flowerree, Jennifer Jhun, Graham Leach-Krouse, Matthew Lee, Fritz Warfield, audiences at the 2014 Central States Philosophical Association and the 2015 Central APA and some very helpful anonymous referees.

[Conn, 2001] Conn, C. (2001). Chisholm, internalism, and knowing that one knows. *American Philosophical Quarterly*, 38(4):333–347.

[Das and Salow, 2018] Das, N. and Salow, B. (2018). Transparency and the k k principle. *Noûs*, 52(1):3–23.

[Dokic and Égré, 2009] Dokic, J. and Égré, P. (2009). Margin for error and the transparency of knowledge. *Synthese*, 166(1):1–20.

[Dretske, 2005] Dretske, F. (2005). Is knowledge closed under entailment. In *Contemporary Debates in Epistemology*. Blackwell.

[Feldman, 1981] Feldman, R. (1981). Fallibilism and knowing that one knows. *The Philosophical Review*, 90(2):266–282.

[Feldman, 1995] Feldman, R. (1995). In defense of closure. *Philosophical Quarterly*, 45(181):487–494.

[Feldman and Conee, 1985] Feldman, R. and Conee, E. (1985). Evidentialism. *Philosophical Studies*, 48:15–34.

[Fitelson, 2006] Fitelson, B. (2006). Williamson's argument against kk. fitelson.org/williamson/notes_6.pdf.

[Ginet, 1970] Ginet, C. (1970). What must be added to knowing to obtain knowing that one knows. *Synthese*, 21(2):163–186.

[Goodman and Salow, 2018] Goodman, J. and Salow, B. (2018). Taking a chance on kk. *Philosophical Studies*, 175(1):183–196.

[Greco, 2014] Greco, D. (2014). Could kk be ok? *The Journal of Philosophy*, CXI(4):169–197.

[Greco, 2015] Greco, D. (2015). Iteration and fragmentation. *Philosophy and Phenomenological Research*, 91(3):656–73.

[Hemp, 2014] Hemp, D. (2014). The kk (knowing that one knows). *The Internet Encyclopedia of Philosophy*.

[Klein, 1995] Klein, P. (1995). Skepticism and closure: Why the evil genius argument fails. *Philosophical Topics*, 23(1):213–236.

[Kvanvig, 2006] Kvanvig, J. L. (2006). Closure principles. *Philosophy Compass*, 1(3):256–67.

[Lasonen-Aarnio, 2008] Lasonen-Aarnio, M. (2008). Single premise deduction and risk. *Philosophical Studies*, 141(2):157–173.

[Lawlor, 2005] Lawlor, K. (2005). Living without closure. *Grazer Philosophische Studien*, 69:25–49.

[Luper, Winter 2011] Luper, S. (Winter 2011). The epistemic closure principle. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. http://plato.stanford.edu/archives/win2011/entries/closure-epistemic/.

[Luzzi, 2010] Luzzi, F. (2010). Counter-closure. *Australasian Journal of Philosophy*, 88(4):673–683.

[McHugh, 2010] McHugh, C. (2010). Self-knowledge and the kk principle. *Synthese*, 173:231–257.

[Moretti and Shogenji, 2017] Moretti, L. and Shogenji, T. (2017). Skepticism and epistemic closure: Two bayesian accounts. *International Journal for the Study of Skepticism*, 11(1):1–25.

[Ramachandran, 2005] Ramachandran, M. (2005). Williamson's argument against the kk-principle. *The Baltic International Yearbook of Cognition, Logic and Communication*, 1.

[Ramachandran, 2012] Ramachandran, M. (2012). The kk-principle, margins for error, and safety. *Erkenntnis*, 76:121–136.

[Schechter, 2013] Schechter, J. (2013). Rational self-doubt and the failure of closure. *Philosophical Studies*, 163:429–452.

[Sharon and Spectre, 2008] Sharon, A. and Spectre, L. (2008). Mr magoo's mistake. *Philosophical Studies*, 139:289–306.

[Smithies, 2012] Smithies, D. (2012). Mentalism and epistemic transparency. *Australasian Journal of Philosophy*, 90(4):723–41.

[Sosa, 1999] Sosa, E. (1999). How to defeat opposition to moore. *Philosophical Perspectives*, 13:141–153.

[Sosa, 2009] Sosa, E. (2009). *A Virtue Epistemology (Apt Belief and Reflective Knowledge, Volume I)*. Oxford University Press.

[Stine, 1975] Stine, G. C. (1975). Skepticism, relevant alternatives, and deductive closure. *Philosophical Studies*, 29:249–261.

[Vogel, 1990] Vogel, J. (1990). Are there counterexamples to the closure principle. In Roth, M. D. and Ross, G., editors, *Doubting: Contemporary Perspectives on Skepticism*, pages 13–27. Kluwer.

[Wang, 2014] Wang, J. (2014). Closure and underdetermination again. *Philosophia*, 42:1129–40.

[Williamson, 2000] Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.

[Williamson, 2005] Williamson, T. (2005). Reply to conee. *Philosophy and Phenomenological Research*, LXX(2):470–476.

[Williamson, 2009] Williamson, T. (2009). Replies to critics. In *Williamson on Knowledge*. Oxford University Press.

[Williamson, 2011] Williamson, T. (2011). Improbable knowing. In *Evidentialism and its Discontents*. Oxford University Press.