What Does Knowledge-Yielding Deduction Require of Its Premises?

Federico Luzzi

According to the principle of Knowledge Counter-Closure (KCC), knowledge-yielding single-premise deduction requires a *known* premise: if S believes q solely on the basis of deduction from p, and S knows q, then S must know p. Although *prima facie* plausible, widely accepted, and supported by seemingly compelling motivations, KCC has recently been challenged by cases where S arguably knows q solely on the basis of deduction from p, yet p is false (Warfield (2005), Fitelson (2010)) or p is true but not known (Coffman (2008), Luzzi (2010)). I explore a view that resolves this tension by abandoning KCC in the light of these challenges, and which acknowledges their force but also their limits. Adopting this view helps identify the epistemic constraints that operate on the premises of knowledge-yielding deduction, clarifies the epistemic role of deduction, and allows us to distil the kernel of truth in the motivations that are standardly taken to support KCC.

1. Friends and Enemies of Knowledge Counter-Closure

When S competently deduces q on the basis of her belief that p, is S's knowing that p *necessary* for S's knowing q? The principle of Knowledge Counter-Closure (KCC) underpins an affirmative answer to this question:

> (**KCC**): Necessarily, if (i) S comes to believe q solely on the basis of competent deduction from her belief that p and (ii) S knows q, then S knows p.

We can isolate two seemingly compelling (and perhaps related) strands of motivation for this principle:

> (**I**): The principal epistemic role of deductive inference is to *transmit* knowledge and other epistemic goods from premise(s) to conclusion. But it cannot create these goods *ex nihilo*. So if there is no knowledge of the premise to transmit in the first place, no knowledge can be transmitted to the conclusion. Therefore, KCC is true.

> (**II**): In single-premise deduction, the epistemic pedigree of the conclusion can be no better than the epistemic pedigree of the premise. So if the conclusion has the pedigree of knowledge, so must the premise. Therefore, KCC is true.

(I) and (II) are powerful enough to make KCC a well-entrenched view: KCC itself or views that entail it enjoy several important endorsements and are at times treated as orthodoxy. For example, Robert Audi writes in his introductory epistemology textbook:

> [w]e can extend our justification and knowledge by inference, but it appears that if we have none to start with, inference […] can give us none. (Audi (2010: 184))

and

> [o]ne kind [of condition on inferential justification and knowledge] concerns the premise(s) of the inference—its foundations, so to speak […] First there are source conditions […]: one needs justification or knowledge in the first place. (185)

In a similar vein, Robert Nozick (1981: 231) begins his search for an analysis of inferential knowledge by laying down two conditions:

> S knows via inference (from p) that q if and only if:
>> (1) S knows that p
>> (2) q is true, and S infers q from p

Nozick justifies the inclusion of condition (1)—which expresses a version of KCC generalized to cover all kinds of inference—by swiftly noting that proof transmits knowledge, and that condition (1) must be true '[o]therwise, there is no knowledge to transmit' (239).

More recently, Timothy Williamson has uncritically relied on one of KCC's instances in his margin-for-error argument (Williamson (1994: 222)[1]. And in later work he argues in favour of the view that one's evidence is one's knowledge (E = K), which under the plausible assumption that the premise of a single-premise deductive inference constitutes evidence for the conclusion, entails KCC (Williamson (2000)). However, the E = K thesis that Williamson deems worth arguing for is stronger than KCC, and tellingly in Williamson's arguments for this thesis no explicit mention of KCC is made.

That these authors accept KCC, either without providing in-depth discussion of the principle or without singling it out as standing in need of argument, indicates that to some prominent epistemologists KCC is plausible indeed.

Yet despite its first-blush plausibility, the motivating considerations in its favour and its endorsements, KCC has recently come under critical fire. Ted Warfield (2005) has proposed alleged cases of deductive knowledge from falsehood. Assuming factivity, these cases are in tension with KCC.[2] Along parallel critical lines, both Coffman (2008) and Luzzi (2010) discuss cases where a subject deduces an allegedly known conclusion from a premise that is true yet 'Gettiered', and hence unknown.

On the one hand, then, we have a seemingly plausible, widely-endorsed principle supported by natural considerations regarding the role of deduction. On the other, there are a series of seemingly compelling counterexamples to the principle. How should we resolve this tension?

My aim in this paper is twofold. Firstly, I aim to clarify the nature of the two different kinds of challenge to KCC, strengthen one kind of challenge, and argue that these challenges should be taken seriously. Doing so raises the question: if knowledge-yielding deduction does not require the premise to be known, do *any* constraints govern the premise of a knowledge-yielding deduction? My second aim is to shed new light on the role of deduction by answering this question affirmatively. I will identify a plausible alternative to KCC, one that imposes a constraint weaker than knowledge on the premise of knowledge-yielding single-premise deduction. The resulting view offers a conception of the role of deductive inference that is different enough from KCC to avoid the problems that beset it, but similar enough to respect the spirit (if not the letter) of the motivations in favour of KCC.

2. An Overview of the Challenges to KCC

---

[1] This is not to say that the appeal to KCC is problematic in the context of Williamson's argument—it isn't. The point is that Williamson appeals to it without subjecting it to critical scrutiny. This suggests that KCC enjoys a strong *prima facie* appeal.

[2] Other authors have floated or argued for this view (see Hilpinen (1988: 164) and Klein (2008)), but none as forcefully as Warfield. See Fitelson (2010) and Arnold (2011) for further discussion of knowledge from falsehood.

Challenges to KCC come in two varieties: alleged cases of knowledge from a false premise and alleged cases of knowledge from a true-yet-unknown premise. I examine each kind in turn.

2.1 Knowledge from Falsehood

The strongest examples of knowledge from false premise come from Ted Warfield (2005), whose arguments target a *consequence* of KCC, rather than KCC itself:[3]

> True Premises (TP): Necessarily, if (i) S comes to believe q solely on the basis of competent deduction from her belief that p and (ii) S knows q, then S's belief that p is true.

While our primary aim here is to understand how these challenges affect KCC, it is clear how Warfield's critique of TP is relevant: for if TP is undermined, then so is KCC. In other words, any attempt to rescue KCC will be successful only if it rescues its consequence, TP. The following example serves as a representative case of this kind of challenge:[4]

> **Handouts**: Counting with some care the number of people present at my talk, I reason: 'There are 53 people at my talk; therefore my 100 handout copies are sufficient'. My premise is false. There are 52 people in attendance—I double counted one person who changed seats during the count. And yet I know my conclusion. (407-408)

Following Warfield and Arnold (2011), I assume that the correct verdict, supported by what Warfield calls a 'clear and widely held intuition' (408), is that the subject knows the conclusion of the inference; consequently, whether this is a case of knowledge from falsehood turns on whether the agent's knowledge is in the proper sense based on deduction from the false belief that there are 53 people at the talk. Warfield observes that the most natural way of defending TP is to maintain that there is some *true* proposition T that does the epistemic work in the deductive inference in lieu of the false proposition *that there are 53 people at my talk* explicitly stated. This true 'proxy' proposition T has the following features: (a) the agent believes T or is disposed to believe T and (b) T is entailed or justified by the false proposition. For instance, defenders of TP would claim that what really grounds the deduction is the *true*, hence knowable, justified proposition *that there are* approximately *53 people at my talk*.

Warfield argues that this strategy ultimately fails because it commits its endorsers to ascribe knowledge where it is intuitively lacking. In the case he considers, a squirrel sits hidden behind the bush in the yard while a fake toy dog visually indistinguishable from a real dog sits within the agent's sight in the yard. Warfield writes (412):

> I seem to see a dog in the yard. On this occasion I form the belief that there is a dog in the yard and then reason as before to the conclusion that there is at least one animal in the yard. My belief is false (there is no dog, only the toy) and my conclusion though true, because of the squirrel behind the brush, is not known. [Defenders of the view that a false premise cannot ground deductive knowledge] get this clear 'no knowledge'

case wrong. They [must state] that I know that there's at least one animal in the yard. After all, there is a justified and dispositionally believed truth that is *both* evidentially supported by and entailed by my false belief that there is a dog in the yard: the truth 'there is a dog or a squirrel in the yard'.

Whether this is a knock-down criticism is debatable. After all, the objection only goes through if defenders of TP hold that the presence of a true proposition justified or entailed by the evidence is sufficient—rather than merely necessary—for the agent to know the conclusion; and it's unclear that defenders of TP who pursue this strategy are committed to holding that. Nevertheless, it does seem that the burden of proof lies with the defender of TP, who should point to further conditions that a proposition must meet if it is to count as a 'proxy' premise. Without a plausible account of what it takes for a proposition to count as a 'proxy' premise, it is difficult to motivate the TP-rescuing strategy that resists the very natural interpretation according to which the premises of the deduction are simply the propositions explicitly considered by the subject, from which the conclusion of the subject's reasoning is drawn. And without a well-motivated defense of TP, we lack a well-motivated defense of KCC.

2.2. A New Challenge from Knowledge from Falsehood

The point against this pro-TP strategy can be strengthened. There is in fact independent reason for thinking that its endorsers are led, if not to certain failure, at least to a further unwelcome result. Consider this case:[5]

> MARBLES: As they swiftly roll by on the wooden track I have assembled for them, I count a series of marbles. The procedure yields 53 as a result. With some confidence, I come to believe that there are 53 marbles on the wooden track. Recalling that my logic professor told me earlier that day that precision entails approximation, I competently deduce that there are approximately 53 marbles, without any loss of confidence in my belief that there are 53 marbles. But despite my best efforts in the difficult task of counting the rapidly-rolling marbles, I double-counted one marble; there are actually only 52.

Similarly to Handouts, in Marbles the agent's deductive inference intuitively yields knowledge even though it proceeds from a false premise. But a distinctive problem now confronts the defensive strategy canvassed, for the most natural candidate 'proxy' proposition is *identical* to the conclusion—namely, *that there are approximately 53 marbles*. And if it is *this* proposition that is doing the epistemic work in lieu of the false premise explicitly considered by the subject, then defenders of TP seem committed to holding that the inference in Marbles is of the blatantly circular form: 'p; therefore p'. But this claim is clearly odd, as the inference seems to exhibit a different structure and seems like a perfectly legitimate way of acquiring knowledge.[6]

---

[5] I owe this example to Crispin Wright.

[6] It is worth remarking that this knowledge need not be all that banal. It may be a surprising fact to some subjects that precision deductively entails approximation—in fact, saying that there are exactly n Fs gives rise to the implicature that the speaker does *not* believe that there are approximately n Fs. (Thanks to Aidan McGlynn for discussion.) Even if it were deemed somewhat banal, it is certainly not as banal as the knowledge delivered by an inference of the form 'p; therefore p'. (Note that while Sorensen (1991) has discussed non-circular inferences of the form 'p; therefore p', the inference at issue here clearly does not share the relevant features of Sorensen's cases.)

A defender of TP might fall back on the following defense: the epistemically relevant proposition, i.e., the proposition on which the agent's conclusion is epistemically based, is something along the following lines.

E: it seems to me as though there are 53 marbles

According to defenders of this strategy, it is E, rather than the false proposition *that there are 53 marbles*, that provides the basis for the deduction. They would claim that the false proposition may be *causally* relevant to the formation of the belief that there are approximately 53 marbles (after all, it is the proposition explicitly considered in the agent's reasoning), but it is not of primary *epistemic* importance.

However, in the absence of a general account of the distinction between epistemically relevant and merely causally relevant premises, this move seems question-begging,[7] and will not satisfy critics of TP. It is dissatisfying to defend TP by claiming that every false premise that seemingly generates knowledge only plays a causal role, without saying more on what underpins the distinction between premises that are epistemically relevant and those that are only causally relevant, other than that the latter may be false but the former must be true.

Moreover, this defense attributes to cognizers an unpalatable blindness to the character of their reasoning. According to this line of defense, although in Handouts I understand my reasoning to be of a deductive nature, it is actually of an ampliative nature, as it really proceeds from the way things merely *seem* to me to a conclusion about how things are in the world.

The general contours of the problem are these: to avoid knowledge from falsehood, a true proxy must be posited; and since a true proxy must be logically weaker than the false premise, then it will not entail all that the falsehood does. The problem with this rescue strategy is that typically the conclusion is among the propositions entailed by the falsehood but not by the proxy, and therefore, on this line of defense, the inference turns out to be ampliative. Lest we run into the problem Marbles highlights, the deductive nature of the inference in Handouts is only salvaged if the premise of the inference is false.

These considerations provide good reason to think that, at the very least, a convincing defense of TP, and hence of KCC, requires a good deal of further work.

2.3. Knowledge from A True-Yet-Unknown Premise

To make matters worse for KCC, there is a distinctive route to attacking this principle, one that does not work by undermining TP. Cases have been described where a *true* yet *unknown* premise allegedly grounds deductive knowledge (Coffman (2008), Luzzi (2010)). E.J. Coffman's is a representative example:

> **Gettierized Reliable Reporter**: CNN breaks in with a live report. The reliable reporter says: 'The President is in Utah, speaking to supporters'. I then reason: 'The President is now in Utah; therefore he is not attending today's NATO talks in Brussels'. I know my conclusion but my premise lacks warrant.[8] Unbeknownst to me, the President is speaking at a "border rally" at the Utah-Nevada border. He is speaking from a large stage that has a "Utah-part" and a "Nevada-part". The President walked into Utah as the reporter was making his assertion about the President's location, which was just before I formed my belief that the President is in Utah. I do not know that the President is in Utah: the President could very easily have been in *Nevada* when I formed

---

[7] See Arnold (2011) for an airing of this complaint.

[8] Coffman understands 'warrant' as the epistemic property that turns true belief into knowledge.

> my belief that he is in Utah. My true belief that the President is in Utah thus falls short
> of knowledge […] (191)

It is easy to see that the KCC-rescuing strategy of finding a true proposition that can ground the inferential knowledge obtained by the subject is of no use here, since the premise that the subject explicitly relies on in these cases *is* true. What defenders of KCC need to argue is that despite appearances to the contrary, there is a *known* proposition—distinct from the 'Gettiered' premise—that does the epistemic work. The natural strategy is to resort to a known proposition concerning how things appear to the agent, along the lines of the proposition *that it seems to me that the President is in Utah.*

But again, this strategy enforces an odd view of the nature of reasoning and of our relationship to our own inferential procedures. Defenders of KCC will have to claim that if the President had been safely in Utah, then the subject's inference would have been grounded in the known proposition that the President is in Utah. On the other hand, if, as Coffman's case has it, the President could very easily have been in Nevada, then the grounding belief is some different, known proposition (e.g. *that it seems to me that the President is in Utah.*) But it is odd that a fact about the external world whose possibility I need never have entertained would on its own be allowed to determine what premise I base my belief on, and, in turn, whether my inference is deductive or ampliative. To endorse this view is to make the content of our premises overly open to the influence of extra-mental factors and to make us curiously blind to prominent, essential features of the inferences we competently perform.

2.4. *What the Challenges Suggest*

Jointly, the considerations provided by critics of KCC and the distinctive problem raised by Marbles cast sufficient doubt on the principle to warrant seriously considering the possibility that it is false. I propose to do just that here by exploring the consequences of resolving the tension between the challenge cases, on the one hand, and KCC and its supporting motivations (I) and (II), on the other, in favor of the former camp.[9] The questions I will be chiefly concerned with are the following: if KCC is false, what does this tell us about the epistemology of deductive inference? What, if any, constraints exist on the epistemic status of one's belief that p, if it is to ground deductive knowledge of q?

---

[9] Might we reject KCC on the basis of only one kind of challenge case, but not the other? I think so, but there are interesting limitations to this strategy. Acknowledging the force of the challenges from true-yet-unknown premises goes hand in hand with acknowledging the force of the challenges from a false premise, but not vice-versa. Here is why. All challenge cases to KCC suggest that single-premise deductive knowledge can tolerate epistemic deficiencies in the grounding premise. From this perspective, what assimilates the two kinds of challenge—knowledge from falsehood and knowledge from true-yet-unknown belief—is that while the conclusion of the deduction is known, there is, among nearby worlds, a world F where the subject's justified belief in the premise is false, and it is *this* world that makes trouble for the subject's actual-world knowledge of the premise. What differs is where in the modal neighborhood F lies: in cases of knowledge from unsafe belief, F is among nearby worlds, but is not identical to the actual world; in cases of knowledge from falsehood, F is among nearby worlds *in virtue of being the actual world.* In this sense, then, knowledge from falsehood is (loosely speaking) simply a limit-case of knowledge from true-yet-unknown belief. It would therefore be odd to concede that cases of knowledge from falsehood are possible, while maintaining that cases of knowledge from true-yet-unknown premise are not possible. It would be feasible, however, to claim that deductive knowledge can tolerate a premise that is true in the actual world and false in a nearby world, but not a premise that is false in the actual world. So judging cases of knowledge from true-yet-unknown premises to be successful does not force us to judge cases of knowledge from falsehood to be successful.

How can we best account for the failure of KCC while providing an undramatic revision of the mainstream view of deduction that holds it as a key tenet?[10]

I propose approaching these questions by firstly understanding the features essentially shared by putative KCC failures. This in turn will help to identify a more plausible principle that can adequately replace the questionable KCC but that still enjoys support from motivations sufficiently similar to (I) and (II). I will argue that there is an epistemic condition weaker than knowledge that belief in p must meet if it is to ground deductive knowledge of q. By showing how a plausible replacement for knowledge is available as a constraint on premises that ground deductive knowledge, I pave the way for those who find at least some of the challenges to KCC compelling, but who nonetheless want to respect the spirit, if not the letter, of motivations (I)-(II). I suggest a course that makes this a tenable—indeed, a robust—position.

3. Epistemic Aristocracy

Diagnosing the challenges will be simplified by recourse to the notion of an *aristocratic* condition. Consider, quite abstractly, a group of inputs, an operation on those inputs, and an output stemming from that operation. Let us define an *aristocratic condition* as follows:

> Condition C is *aristocratic* with respect to a procedure P, inputs $I_1…I_n$ and an output O just in case: O meets C only if each of $I_1…I_n$ meets C.

Thus, for example, the condition *being gluten-free* is an aristocratic property of a cake (output) with respect to the procedure of baking its ingredients (the inputs). That is to say, a cake is gluten-free only if all of its ingredients are gluten-free. Equivalently, if any of the ingredients of a cake are not gluten-free, then baking those ingredients will yield a cake that is also not gluten-free.

By contrast, the condition *being sweet* generally is not an aristocratic property of a cake (output) with respect to the procedure of baking its ingredients (the inputs). That is to say, it is not true that a cake is sweet only if all its ingredients are sweet: most sweet cakes contain non-sweet ingredients (flour, eggs).[11]

The challenge cases to KCC clearly exploit non-aristocratic conditions on knowledge with respect to the procedure of competent single-premise deduction.[12] Truth is non-aristocratic in this way, since it is possible for the conclusion (output) of such a deduction to be true without its premise (input) being true. By filling out the details of such a case to bolster the epistemic credentials of the agent's belief in the conclusion (by making it safe, reliably and responsibly believed, unaffected by undefeated defeaters, etc.), the case can be made—as Warfield has done—that the conclusion is not only true, but also known. Safety is similarly non-aristocratic: it is possible for the conclusion of a competently performed single-premise deduction to be safely believed without the

---

[10] I am not, nor will be, claiming that the challenge cases conclusively show that KCC is false. (Note that this conclusion is embraced in neither Coffman (2008) nor Luzzi (2010).) I merely wish to argue that those who deny KCC can endorse a principle that places better constraints on the epistemic pedigree of a premise that grounds deductive knowledge.

[11] Typically (yet not without exception), aristocratic titles are only obtained by inheriting them. One is a member of the aristocracy only if one's parents are members of the aristocracy. It is this feature of aristocratic titles that makes the label appropriate for the epistemic conditions I am interested in.

[12] To make exposition less cumbersome, I will not always specify the procedure with respect to which an epistemic condition is aristocratic. In these cases, the relevant deductive inference is the procedure with respect to which the condition is aristocratic.

premise it is based on being safely believed. By filling out the details appropriately, the case can be made—as Coffman (2008) and Luzzi (2010) have done—that the conclusion is not only safe but also known.

Challenges against KCC essentially make play with non-aristocratic epistemic conditions. As one of my aims is to identify a plausible replacement principle that will guarantee that some condition weaker than knowledge apply to the premise of a knowledge-yielding competent deduction, it is wise to search for epistemic conditions that, like truth and safety, are necessary for knowledge, but that, unlike truth and safety, are aristocratic. For if an epistemic condition is both necessary for knowledge and aristocratic, then the failure of the conclusion to meet such a condition will preclude both premise and conclusion from being known. This would guarantee that the condition in question could not be exploited to devise cases where the conclusion of a single-premise deduction is known but the premise is not. In what follows, I suggest that *epistemic responsibility* is a condition that fits the bill and that a principle like KCC phrased in terms of epistemic responsibility is a plausible replacement for KCC.

4. Epistemic Responsibility

Consider the following case:

> Harry comes to believe on the basis of no evidence whatsoever that Sally is not married. In fact, he is correct. Moreover, Sally could not easily have been married, since unbeknownst to Harry she feels marriage to be a completely unnecessary and intrusive sanctioning by the state or Church of the real binding element of a romantic relationship—true love. He then reasons as follows:
>
> (1) Sally is not married.
> (2) (If Sally is not married, then Sally is not married to Sheldon.)
> _____
> (3) Sally is not married to Sheldon.

Does Harry know on the basis of this deduction that Sally is not married to Sheldon? Clearly not, since his belief that (3) is irresponsibly held. I suggest that what accounts for his holding this belief that (3) irresponsibly is that, in general, relying on an irresponsibly held belief in one's reasoning is itself an epistemically irresponsible act, and performing an epistemically irresponsible act tarnishes the epistemic results of that act. Short of ceasing to believe the irresponsibly held belief, there is nothing one can do with an irresponsibly held belief to remove the blemish of epistemic irresponsibility from one's actions (including deduction from that belief) and the product of those actions (belief in the conclusion of the deduction).

Compare: stealing a car is usually illegal.[13] In circumstances where stealing a car is illegal, then doing anything with the illegally obtained car—driving it, selling it, painting it blue, using it as a kitchen table—is also illegal. Once stolen, there is little one can do with a car to restore legality to one's actions involving that car. And if the car is used to obtain something by some procedure—say, the car is traded for a motorcycle—or is given to someone in need, the resulting state of affairs is also marked by the blemish of illegality.

I suggest that there is an obstacle to the construction of plausible challenges to KCC that exploit irresponsibly held belief. The obstacle is the compelling thought that if a subject is not epistemically responsible in believing p, then neither is she epistemically responsible in using it as a premise in deductive reasoning; and as a result, her belief that

---

[13] There may be exceptions—stealing a car if this is the only way of saving an innocent person may not be illegal. Thanks to Aidan McGlynn here.

q, based solely on competent deduction from p, cannot be epistemically responsibly held. Therefore, epistemic responsibility is aristocratic. Given that it is also a necessary condition on knowledge, epistemic responsibility fits our desired profile, suggesting that the following principle of Epistemic Responsibility Counter-Closure (ERCC) holds:

> (**ERCC**): Necessarily, if (i) S comes to believe q solely on the basis of competent deduction from her belief that p, and (ii) S's belief that q is responsibly held, then S's belief that p is responsibly held.

Of course, it will not do to rely simply on the intuitive plausibility of ERCC. A more compelling case for ERCC can be made by arguing two further points: (i) ERCC can handle the kind of challenge that makes trouble for KCC; (ii) there is a plausible explanation available for why no compelling challenges to ERCC should be expected, while challenges to KCC should be. I argue these points in turn.

Let us return to Handouts and modify the case so that the subject's premise is irresponsibly believed. Suppose, for example, that in counting the audience the subject is sloppy: she often double-counts, or frequently skips an audience member, or roughly calculates the total number by estimating the number of delegates per row and the number of rows, and performing (perhaps erroneously) the multiplication. On this basis, she comes to believe that there are 53 people in the audience. And on the basis of this belief, she deduces that her 100 handouts are enough. Is it coherent to suppose that the subject nonetheless knows the conclusion?

I wish to answer this question negatively. But before I do so, let me mention and set aside a clear sense in which it *is* coherent to suppose she knows her conclusion. If the subject were in a position to know before the count that counting sloppily would not make a difference to the truth of her conclusion (that her 100 handouts are enough) then it is clear that she knows her conclusion. Despite the approximate nature of her counting method, she may know in advance that the margin of error of this method is relatively limited so as to render her belief-forming method incapable of yielding a drastically mistaken result. But in this case, it seems that either the subject knows in advance that her 100 handouts are enough, or that this knowledge is not acquired solely via deduction from her belief that there are 53 people in the audience, but at least in part by relying on a premise concerning the limited nature of her margin of error. However, while there are plenty of epistemic routes to believing that her 100 handouts are enough, many of which will yield knowledge of this proposition, ERCC (like KCC) is only falsified if the belief in this conclusion is based *solely* on deduction from the belief that there are 53 people.

I suggest, then, that insofar as the subject knows that her 100 handouts are enough when she counts sloppily, she is not basing her belief solely on deduction from her irresponsibly held belief that there are 53 people in the audience. On the other hand, when deduction from this belief is the only basis for her belief in the conclusion, she does not know her conclusion. The reason she does not know is that this belief is held epistemically irresponsibly, since the premise it is based on is also held irresponsibly.[14]

---

[14] The account of introspective knowledge proposed by Alex Byrne might seem to provide the resources to falsify ERCC (See, e.g. Byrne (2005) and (2011)). For example, suppose I am epistemically irresponsible in holding a belief that p and come to believe by introspection that I believe that p. This second-order belief seems responsibly formed and may constitute knowledge. Is this in tension with ERCC? I don't think so. Firstly, it is controversial whether introspection-based transitions from *p* to *I believe that p* are inferential at all (see e.g., Dretske (1995: 60-62) and Boyle (2011)); but—more relevantly—even if they turned out to be inferential, they are clearly not *deductive*, so they do not fall under the jurisdiction of ERCC and cannot falsify the principle. (I am grateful to Lauren Ashwell and Aidan McGlynn for discussion of this point.)

Why are challenges to KCC to be expected, while challenges to ERCC are not? One way of explaining this has to do with the ascertainability by the subject of these epistemic constraints (or lack thereof) and how ascertainability interacts with epistemic responsibility. Consider first that whether a subject's belief is true, or not Gettiered, is a matter that is not always ascertainable by the subject,[15] and that therefore a belief's being false or being Gettiered is perfectly compatible with the agent's holding that belief responsibly; indeed, it is natural to think that part of the reason why the beliefs held by subjects in Gettier cases are justified is that the subjects are not to be blamed for holding them—rather, circumstances conspire against them to deprive them of knowledge despite their epistemically responsible behaviour. Even the epistemic efforts of the best epistemic agents can be flustered if the environment is sufficiently uncooperative; so despite an agent's perfect policing and rational management of her beliefs in response to the evidence available to her with respect to p—i.e., despite epistemically responsible behaviour with respect to p—her belief that p may nonetheless fail to constitute knowledge.

On this natural picture, there are constraints on knowledge whose obtaining or failure to obtain lie beyond the subject's purview, and hence are not and cannot lie within the scope of her epistemic responsibility. It is unsurprising, then, that it is possible for an agent to deduce q from p *while not knowing p* while acting epistemically responsibly throughout, as long as the failure of knowledge for p is due to the failure of some non-ascertainable constraint (truth, non-Gettierization). The obstacle from epistemic responsibility that prevents responsibly held belief that q from arising via deduction from an irresponsibly held belief that p—thereby ensuring the truth of ERCC—does not rule out the possibility that premise p is not known while conclusion q is responsibly believed. Provided that the agent's belief that q meets external epistemic conditions, a strong case can then be made that q is known. In other words, as the challenges by Warfield, Coffman, and others illustrate, challenges to KCC can be brought compatibly with the idea that responsibly held belief must be grounded in responsibly held belief; but challenges to ERCC, of course, conflict with this very plausible idea.

It is thus unsurprising that challenges to KCC hitherto raised share this common feature: despite not knowing the premise, the agent holds his belief in the premise *responsibly*. In Handouts, for instance, the agent's counting procedure is only wrong by 1 in circumstances where the numbers are large, someone changed seats during the count and the counting was in any case performed carefully; in Marbles, the numbers are large and the marbles roll by very swiftly; in Gettierized Reliable Reporter, the border-straddling position of the stage is highly unusual and the agent has no reason to suspect the stage to be so placed. In all cases, the agent has behaved epistemically appropriately: any shortcoming in the epistemic status of their conclusion cannot be chalked up to any shortcoming in the way the agent appreciated the evidence or formed a belief on its basis.[16]

If I am correct, this shared feature is no philosophical contingency; rather, it is an essential characteristic of any plausible challenge to KCC. For if agents in these examples failed to know their premise on account of believing it in an epistemically irresponsible way, then this fact would compromise the knowability of the conclusion thereby drawn, rendering the case toothless against KCC.

I have advanced my diagnosis without saying much about what epistemic responsibility amounts to. I have worked with a rough notion that excludes beliefs

---

[15] This may be more controversial when the belief concerns the agent's own mental states or internal psychology. I set these cases aside here.

[16] This feature is also shared by Warfield's other cases of alleged knowledge from falsehood.

formed on no evidence and that does not necessarily exclude beliefs that are Gettiered or false. Further remarks can be made to clarify this notion. For example, what it takes for a belief to be responsibly held depends, perhaps non-exclusively, on the evidence available to the subject and on the epistemic capabilities of the subject. Thus, I may be epistemically irresponsible in coming to believe that Dylan did not eat the last cookie, whereas you are responsible in doing the same, if I leap to this conclusion on the basis of wishful thinking, whereas you hold that belief on the basis of appropriately appreciating evidence that rules out Dylan's having eaten the last cookie. A highly-skilled mathematician may be irresponsible in believing on the basis of testimony a complex mathematical result that should strike her as dubious and that she is capable of checking for herself, whereas I may hold the same belief responsibly on the basis of the same testimony, if I lack the mathematical ability to detect its dubiousness and check its credentials.

In general, if I hold a belief irresponsibly, it is appropriate to think that I should have done epistemically better in some respect or other, and that the fault lies with me. The respect in which I should have done better may vary across cases, but in no case will I be able to maintain that my failure was *in no way* due to a mismanagement of the epistemic situation on my part. If, on the other hand, I hold a belief responsibly, then that belief may still be wanting in some respect—it may be false or fail to amount to knowledge for some other reason—but I cannot be legitimately criticized for failing to fulfil my epistemic duties.

These remarks cannot constitute a full analysis of epistemic responsibility—but for our purposes there is no need for one. As long as it is agreed that epistemically responsible belief—whatever that may precisely amount to—is a condition that is (i) necessary for knowledge, (ii) weaker than knowledge, and (iii) aristocratic, then critics of KCC can happily abandon that principle and claim to have found in ERCC a plausible replacement.

That my proposal is not committed to a particular view of epistemic responsibility allows it to remain open on the question of whether epistemic responsibility suffices for justification. Some philosophers take responsible belief not to be sufficient for justified belief on the basis of the following kind of case. Suppose I use a certain statistical method taught to me by my mentor, and responsibly believe the conclusions I draw on the basis of applying it to my carefully collected data. But suppose also that unbeknownst to me the method is based on some deeply misguided assumptions, and systematically delivers results that are nowhere near the correct ones. While I am not irresponsible in believing the results, the faultiness of the method prevents my beliefs formed on its basis to count as justified—or so the diagnosis goes.[17]

I do not wish to judge the merits of this type of case, as my project—that of providing a fallback principle for opponents of KCC—is compatible with either verdict on this case. Those who are not impressed by it and insist that responsibility is coextensive with justification should be happy to endorse the following justification-based principle.

> (**DJCC**): Necessarily, if (i) S comes to believe q solely on the basis of competent deduction from her belief that p, and (ii) S's belief that q is doxastically justified, then S's belief that p is doxastically justified.

By contrast, those who take this kind of case to drive a wedge between epistemic responsibility and justification will be disinclined to accept DJCC. For if a belief's being doxastically justified is not exhausted by its being held epistemically responsibly—that is,

---

[17] Pryor (2001: 114-115). See Conee and Feldman (2004: 63) for an analogous case.

if whether S's belief is justified is not a matter that lies within the subject's epistemic purview—then it is possible to construct cases where (i) S comes to believe q from competent deduction from p, (ii) S's belief that p is not doxastically justified because it fails some non-aristocratic constraint on justification (but is epistemically responsibly held) and (iii) S's belief that q meets that non-aristocratic constraint on justification and is epistemically responsibly held. These cases will be analogues for DJCC of the challenges to KCC discussed above. But even though such philosophers will reject DJCC, they will still be able to accept ERCC.

So I am not committing to either view on the relationship between epistemic responsibility and justification. No matter what our view is—whether we think epistemic responsibility is coextensive with justification, or is a weaker notion—there is a plausible replacement for KCC that allows us to abandon KCC while still placing some epistemic constraints on the premise of single-premise knowledge-yielding deduction.

5. Amending The Motivations For KCC

The orthodox conception of deductive inference was underpinned by motivations (I) and (II). But these motivations can be defused and recalibrated by those who take seriously the challenge cases to KCC. Here is how.

Recall that according to motivation (I), deduction is a process whose fundamental function is extending our body of *knowledge*: on this picture, single-premise deductive inference takes some known proposition as an input and, if competently performed, delivers a known proposition as an output. Metaphorically, competently performed deductive inference *transfers* the status of knowledge from premise to conclusion. Consequently, if there is no knowledge of the premise to transmit in the first place, none can be transmitted to the conclusion. Deduction seemingly cannot create knowledge and other epistemic goods *ex nihilo*.

Taking seriously the challenges to KCC, however, suggests that the function of extending knowledge is at best a derivative one. Unless we are prepared to judge that deductive inference in the challenge cases cannot perform its main role, we should concede that the primary role of deductive inference is rather to extend our body of *responsibly held* beliefs: to take an epistemically responsibly held belief as an input and deliver an epistemically responsibly held belief as an output, transferring the status of *responsibly held belief* from premise to conclusion. This is not to say that extending our body of responsibly held beliefs doesn't also typically bring about an extension of our body of knowledge; assuming the external world is sufficiently cooperative, an agent's responsibly held belief will also meet external conditions required for knowledge. However, as challenges to KCC indicate, there are situations where competent deduction leads us to knowledge, even though the input belief is not known. The orthodox conception of deduction proves to be too narrow here: in requiring deduction to take known belief as an input if it is to deliver known belief as an output, the threat from the challenge cases to KCC is unduly neglected.

If this is right, then deduction *can* 'create' some epistemic goods *ex nihilo*, but not in a way that should suggest epistemological 'rabbit out of hat' trickery: quite simply, there are epistemic goods that a premise may lack but that are nevertheless had by a conclusion believed solely via deduction from that premise. For instance, opponents of KCC can say, it is possible for a false premise to yield a true conclusion, as Warfield's Handouts case shows; and it is possible for an unsafe premise to yield a safe conclusion, as Coffman's Gettierized Reliable Reporter illustrates. Admittedly, there are some epistemic goods that cannot be so 'created *ex nihilo*': namely, the aristocratic constraint of *epistemic responsibility*. So there is something right about the spirit of this proposed

motivation for KCC, since some epistemic goods cannot attach to the conclusion of a deduction unless they attach to the premise; but we would be wrong to include knowledge among these goods, since several non-aristocratic components to knowledge (e.g. truth and some anti-Gettier conditions like safety) can be met by the conclusion even if they are not met by the grounding premise.

Motivation (II) in favour of KCC was that the epistemic pedigree of a conclusion can be no better than that of the premise. Consequently, this train of thought runs, if the conclusion is known, so is the premise. But in the light of the foregoing considerations, it should be clear that those wishing to reject KCC may deem this understanding of 'epistemic pedigree' too broad. This camp will claim that the motivation is mistaken if the epistemic pedigree of a belief includes its non-aristocratic constraints. Only insofar as it includes aristocratic constraints—whose obtaining is a matter that lies within the scope of the agent's epistemic responsibility—is it true that the pedigree of a conclusion cannot surpass the pedigree of the premise that deductively grounds it. Because knowledge entails non-aristocratic constraints, but epistemic responsibility is aristocratic, this restriction in the original motivation ceases to support KCC and supports ERCC instead.

Both these amendments to the original motivations for KCC can be proposed, in a suitably paraphrased form, by those who take epistemic responsibility to coincide with justification. Philosophers of this persuasion will seek to replace the controversial KCC with DJCC and will understand the primary role of deductive inference as the extension of our body of justified beliefs, rather than the extension of our body of knowledge. On this line, competent single-premise deduction serves primarily to transfer the status of *justification* from premise to conclusion.

6. Conclusion

So far, the literature on KCC has focused on whether in proposed challenges an agent's belief that q can be based solely on competent deduction from a false or Gettiered belief in premise p and still amount to knowledge. The repercussions of rejecting KCC, however, have not hitherto been traced. In this paper, I have advanced the debate by showing that those who reject KCC need not demand drastic revisions of a commonplace view of deduction, but need only make minor adjustments to the motivating picture for KCC. Insofar as rejecting KCC comes hand in hand with recognising the truth of ERCC, would-be deniers of KCC can agree that knowledge-yielding single-premise deduction does place some constraint on the epistemic status of the premise, but disagree with the mainstream view about the nature of this constraint: rather than knowledge, only epistemically responsibly held belief is required. Tracing the theoretical commitments of denying KCC, and highlighting their lightness, thus serves to pave the way for a rejection of this principle.[18]

References

Arnold, A. (2011). Some Evidence Is False. *Australasian Journal of Philosophy* 91, 165-172.

Audi, R. (2010). *Epistemology: A Contemporary Introduction to the Theory of Knowledge,* 3rd edition. London: Routledge.

Boyle, M. (2011). Transparent Self-Knowledge. *Proceedings of the Aristotelian Society Supplementary Volume*, 85, 223-241.

Byrne, A. (2005). Introspection. *Philosophical Topics*, 33, 79-104.

Byrne, A. (2011). Transparency, Belief, Intention. *Proceedings of the Aristotelian Society Supplementary Volume*, 85, 201-221.

Coffman, E. (2008). Warrant Without Truth? *Synthese*, 162, 173-194.

Conee, E., & Feldman, R. (2004). *Internalism Defended*. Oxford: Oxford University Press.

David, M., & Warfield, T. (2008). Closure Principles and Skeptical Arguments, in Q. Smith (Ed.), *Epistemology: New Essays* (p. 137-188). Oxford: Oxford University Press.

Dretske, F. (1995). *Naturalizing the Mind*. MIT Press.

Fitelson, B. (2010). Strengthening the Case for Knowledge From Falsehood. *Analysis*, 70(4), 666-669.

Hilpinen, R. (1988). Knowledge and Conditionals. *Philosophical Perspectives*, 2, 157-182.

Klein, P. (2008). Useful False Beliefs, in Q. Smith (Ed.), *Epistemology: New Essays* (p. 25-62). Oxford: Oxford University Press.

Nozick, R. (1981). *Philosophical Explanations*. Oxford: Oxford University Press.

Luzzi, F. (2010). Counter-Closure, *Australasian Journal of Philosophy,* 88, 673-683.

Pryor, J. (2001). Highlights of Recent Epistemology. *British Journal for the Philosophy of Science*, 52(1), 95-124.

Sorensen, R. (1991). 'P, Therefore, P' Without Circularity. *Journal of Philosophy*, 88, 245-266.

Warfield, T. (2005). Knowledge from Falsehood. *Philosophical Perspectives*, 19, 405-416.

Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University

Williamson, T. (1994). *Vagueness*. Oxford: Oxford University Press.