

Knowledge from Blindspots¹

Rhys Borchert, Juan Comesaña, and Timothy Kearl

Abstract

No False Lemmas (NFL) says: necessarily, S's belief that p is knowledge only if it is not inferred from any falsehood. Its proponents argue that alleged counterexamples to NFL are really cases of knowledge despite falsehood, wherein the false premise is inessential to the inference; perhaps some nearby truth does the justificatory heavy lifting. This chapter argues that there can be cases of inferential knowledge from a blindspot premise. Given that in such cases the relevant falsehood is essential to the inference, one cannot employ the knowledge despite falsehood strategy to defend NFL. The chapter concludes by discussing how cases of knowledge from blindspots exemplify one way in which we can gain inferential knowledge defectively or non-ideally.

1. Introduction

At least some Gettier cases involve inference from a false belief. This prompted the “No False Lemmas” (NFL) definition of propositional knowledge, according to which knowledge is incompatible with inference from a false belief. Some have argued that the NFL view fails in the same way in which the justified true belief definition of knowledge failed—namely, but not providing sufficient conditions for knowledge, for there allegedly are cases where subjects have justified true beliefs not inferred from a falsehood that nevertheless do not amount to knowledge. But, interestingly, it has also been claimed that the NFL view fails to provide necessary conditions for knowledge, because it is possible to know a proposition on the basis of inferring it from a false belief. This claim has sparked debate, with some arguing that, in those cases where knowledge is genuinely present, the falsehood in question is not essential to it—that they are cases of knowledge despite falsehood, rather than cases of knowledge from falsehood.

In this note, we argue for the existence of cases of knowledge from blindspots—i.e., knowledge from unknowable propositions. An interesting feature of those cases is that, if they are genuinely possible, then they are cases of knowledge from falsehood where the falsehood is essential, and so the aforementioned objection to cases of knowledge from falsehood does not apply. The rest of the chapter develops as follows: Section 2 briefly recounts the sort of Gettier cases that gave rise to the NFL view, Section 3 argues for the existence of knowledge from blindspots, Section 4 defends the existence of knowledge from falsehoods from objections, Section 5 puts the existence of knowledge from blindspots in the larger context of different kinds of defective knowledge, and Section 6 concludes.

2. Gettier Cases, NFL, and Counter-Closure

Recall Gettier's (1963) infamous case of the man with ten coins in his pocket:

¹ Penultimate draft of a forthcoming chapter in [Illuminating Errors: New Essays on Knowledge from Non-Knowledge](#). Please cite published version.

Suppose that Smith and Jones have applied for a certain job. And suppose that Smith has strong evidence for the following conjunctive proposition:

(d) Jones is the man who will get the job, and Jones has ten coins in his pocket.

Smith's evidence for (d) might be that the president of the company assured him that Jones would in the end be selected, and that he, Smith, had counted the coins in Jones's pocket ten minutes ago. Proposition (d) entails:

(e) The man who will get the job has ten coins in his pocket.

Let us suppose that Smith sees the entailment from (d) to (e), and accepts (e) on the grounds of (d), for which he has strong evidence. In this case, Smith is clearly justified in believing that (e) is true.

But imagine, further, that unknown to Smith, he himself, not Jones, will get the job. And, also, unknown to Smith, he himself has ten coins in his pocket. Proposition (e) is then true, though proposition (d), from which Smith inferred (e), is false. In our example, then, all of the following are true: (i) (e) is true, (ii) Smith believes that (e) is true, and (iii) Smith is justified in believing that (e) is true. But it is equally clear that Smith does not *know* that (e) is true; for (e) is true in virtue of the number of coins in Smith's pocket, while Smith does not know how many coins are in Smith's pocket, and bases his belief in (e) on a count of the coins in Jones's pocket, whom he falsely believes to be the man who will get the job.

This case and others like it have engendered two, related lines of inquiry. First, Gettier's own remarks suggest that the source of the problem is the *falsity* of the lemma, (d), from which Smith infers (e). Even today, there are many ardent defenders of this "No False Lemmas" solution to the Gettier Problem; one cannot gain knowledge from falsehood, or so say its proponents.²

No False Lemmas: Necessarily, S's belief that p is knowledge only if it is not inferred from any falsehood.

Second, consider what Federico Luzzi (2010) has called "Counter-Closure":³

Counter-Closure: Necessarily, if (i) S knows that p entails q and (ii) S comes to believe q solely on the basis of competently deducing it from p, and (iii) S knows q, then S knows p.

Counter-Closure offers another way to explain why Smith fails to know (e) by way of the following argument:

² For support of the No False Lemmas thesis, see Harman (1973) and, more recently, Coffman (2008), Schnee (2015), Montminy (2014), and Lee (2021). For arguments against it, see Warfield (2005) and Fitelson (2010). Some of these views will be discussed in more detail below. This idea has taken on a life of its own, even if it is widely agreed to fail to provide a general solution to Gettier-style examples; not all Gettier-style examples are inferential, so not all Gettier-style examples will turn on inference via a false lemma.

³ To be fair, Luzzi was not offering a full-throated endorsement of Counter-Closure, but rather a critical discussion of that principle in light of problem cases, alongside a menu of solutions. See also Warfield (2005) and Klein (2008), and Ball and Blome-Tillman (2014) for a recent defense of Counter-Closure.

1. It's not the case that Smith knows (d).
2. Smith knows that (d) entails (e).
3. Smith comes to believe (e) solely on the basis of competently deducing it from (d).
4. Necessarily, if Smith knows that (d) entails (e) and Smith comes to believe (e) solely on the basis of competently deducing it from (d), and Smith knows (e), then Smith knows (d).
5. Therefore, it's not the case that Smith knows (e).

This argument is valid, and its premises seem, to many, unobjectionable; after all, premises 1–3 are stipulations of the case, and premise 4 is simply an instance of Counter-Closure. This argument, moreover, offers friends of NFL a bit of adjacent support. It may be, for instance, that the appeal of NFL derives, at least in part, from the truth of Counter-Closure (or *vice versa*).

Here, we hope to upset this happy theoretical alignment. We present a novel argument against Counter-Closure by appeal to knowledge from unknowable premises. If one can gain knowledge from unknowable premises—blindspots and Moorean abominations among them—then not only is Counter-Closure false, but also to the extent that NFL derives some plausibility from Counter-Closure, NFL looks much less appealing.

3. Knowledge from Blindspots

Let's start by considering a vignette:

Bamboozle: Juan is a convincing epistemologist, so convincing, in fact, that over the course of his lecture he gets his students to believe that (i) knowledge of the external world is impossible, but (ii) that he would only be able to convince them of (i) if he exists, and therefore, that there is an external world. One of these students, Rhys, reflects on this and comes to believe the following proposition: “There is an external world but I do not know it”. Suppose that Rhys then deduces from that “there is an external world”.

Against the background of some assumptions, the proposition *There is an external world but I do not know it* is unknowable. For suppose that Rhys knows it. Then it is true—that is to say, there is an external world but Rhys doesn't know it. Contradiction. Therefore, it is unknowable.⁴ Some would call such a proposition “abominable”,⁵ perhaps casting it to the flames of Moorean absurdity.⁶ As a shorthand, we'll just call these “blind spot propositions” or say that a certain proposition is “in one's blindspot”.⁷

But even though Rhys cannot know that there is an external world but he doesn't know it, he can justifiably believe it—after all, Juan has given him very convincing arguments for that proposition. Moreover, Rhys employs that justifiably believed blindspot proposition as a premise in his

⁴ Fitch (1963) and Church (2009). Of course, what makes them peculiar—however that peculiarity is described—is that they are unknowable *de se*. There is nothing, after all, odd about *someone else* believing that *Rhys doesn't know there is an external world and there is an external world*.

⁵ DeRose (1995).

⁶ Moore (1942).

⁷ The terminology is taken from Sorensen (1988).

competently deducing that there is an external world, in virtue of which he knows that there is an external world. If we are right, then the proposition on the basis of which Rhys knows that there is an external world is not only unknowable (to him), but it is also false.

Compare now Bamboozle with Warfield's example of knowledge from falsehood:

I have a 7pm meeting and extreme confidence in the accuracy of my fancy watch. Having lost track of the time and wanting to arrive on time for the meeting, I look carefully at my watch. I reason: "It is exactly 2:58pm; therefore I am not late for my 7pm meeting". Again I know my conclusion, but as it happens it's exactly 2:56pm, not 2:58pm. (Warfield 2005, p. 408)

Branden Fitelson (2010) has suggested that one could respond to cases like Warfield's by saying that the falsehood in question is not essential, in the sense that, had it been true, the subject would have still been in a position to know based on it. Fitelson then argues that there are cases of knowledge from falsehood which satisfy what we will call "Fitelson's counterfactual condition":

Fitelson's counterfactual condition: If the subject's belief p had not been false, then the example would not have constituted a case of inferential knowledge.

Warfield's case violates Fitelson's counterfactual condition because, had it been exactly 2:58 pm, the subject would still have inferentially known that he was not late for his 7 pm meeting. Fitelson presents his own case, which he argues satisfies his counterfactual condition:

I have a 7pm meeting and extreme confidence in the accuracy of both my fancy watch and the Campanile clock. Having lost track of the time and wanting to arrive on time for the meeting, I look out of my office window (from which the Campanile clock is almost always visible). As luck would have it (owing, say, to the fluke occurrence of a delivery truck passing by my window), the Campanile clock is obscured from view at that instant (which is exactly 2:56pm). So, instead, one minute later, I look carefully at my watch, which (because my watch happens to be running one minute slow) reads exactly 2:56pm. I reason: "It is exactly 2:56pm (p) therefore (q) I am not late for my 7pm meeting". Thus (supposing Warfield is right), I have inferential knowledge that q , based on a relevant premise p , which is a falsehood. Now for the twist. If my belief that p had been true, then (we can plausibly suppose) it would have been based on my reading (at exactly 2:56pm) of the Campanile clock, which would have read exactly 2:56. Unbeknownst to me, however, the Campanile clock has been (and would have been) stuck at 2:56 for some time.

The idea is that Fitelson's example satisfies his counterfactual condition because, had the belief that it is exactly 2:56 pm been true, it would have been based on his reading of the Campanile clock, but the Campanile clock is stuck at 2:56, and no knowledge about the time can be gained from reading a stuck clock.

Another related objection to alleged cases of knowledge from Falsehood comes from Coffman (2008), who suggests that, in all those cases, there is a nearby truth which is doing all the heavy lifting. More precisely, according to Coffman in any such case, there is a true proposition p' such that (i) the subject is (at least) disposed to believe p' and (ii) if the subject's inferential belief (that q) had been based on a belief in p' , the inferential belief would still have constituted knowledge. Fitelson

argues that his own example can be slightly modified so as to deal with this reply from Coffman: suppose that, if Fitelson had based his belief that he is not late for his 7:00 pm meeting, not on the proposition that it is exactly 2:56 pm, but rather on the proposition that it is *approximately* 2:56 pm, that would have been based on his reading of the Campanile clock (Fitelson thinks, say, that the Campanile clock is only approximately right, but his own fancy watch is always exactly right).

Fitelson's counterfactual condition can be strengthened. Consider instead the following *essentiality* condition:

Essentiality condition: Necessarily, if the subject's belief that p is true, then the subject does not know q .

Fitelson's case does not satisfy the essentiality condition, for it is obviously possible for Fitelson's belief that it is exactly 2:56 pm to be true and for him to have knowledge that he is not late for his 7:00 pm meeting. But our Bamboozle case does satisfy the essentiality condition. In any possible world in which it is true that there is an external world and Rhys does not know it, Rhys does not know that there is an external world. And given that our case satisfies the essentiality condition, it also obviously satisfies the counterfactual condition.

But is our case possible? We anticipate two sources of resistance: first, that in addition to being unknowable, blindspot propositions cannot even be justifiably believed. And second, that Rhys's knowledge is not derived from his competent deduction from a blindspot.

Suppose one thought that blindspot propositions cannot even be justifiably believed. One might point to something like the "felt inconsistency" of such propositions as grounds for skepticism about the possibility of justifiably believing something in one's blindspot. But this reaction is too strong. Start by noticing that a disjunction of blindspot propositions is not, in general, itself in one's blindspot. There is no felt inconsistency, abominableness, or Moorean absurdity in the air when one believes *de se* that *I am content to live in an ice hut and I doubt it OR I am not content to live in an ice hut and I doubt it*.⁸ Moreover, an author's statement in the preface of her manuscript is a long disjunction of blindspots (Sorensen, 1988), but we don't accuse authors of *unjustifiably* or *irrationally* acknowledging the likelihood that their work contains errors, despite thinking, of each particular claim, that *it* is not erroneous.

So, not only can a disjunction of blindspot propositions be justifiably believed, but intellectual humility might also even demand that one believes a disjunction of blindspot propositions. But now imagine that the author, after acknowledging her fallibility in the preface of her manuscript, proceeds to double-check her claims individually. One way to conceptualize what goes on in her double-checking is that the author engages in a long chain of disjunction elimination. She begins with the disjunction expressed in the preface, where each disjunct is itself a blindspot proposition:

$(P \text{ and not-}K(P)) \text{ OR } (P' \text{ and not-}K(P')) \text{ OR } (P'' \text{ and not-}K(P'')) \text{ OR } \dots$

⁸ The example is taken from the anthropologist Gontran de Poncins (1941).

She first eliminates the possibility expressed by *P and I don't know that P* by, say, double-checking her grounds for believing that *P* and deeming them adequate. But she is already in a position to justifiably reason as follows:

- 4) $(P \text{ and not-}K(P)) \text{ OR } (P' \text{ and not-}K(P')) \text{ OR } (P'' \text{ and not-}K(P'')) \text{ OR } \dots$
- 5) *It's not the case that (P and not-K(P))*
- 6) Thus: $(P' \text{ and not-}K(P')) \text{ OR } (P'' \text{ and not-}K(P'')) \text{ OR } \dots$

Again, her justification for premise 4 is her knowledge of her own fallibility, and her justification for premise 5 is her double-checking. At a certain point in this process of double-checking, perhaps as she nears the end of her manuscript, disjunction elimination will lead her to a blindspot. If disjunction elimination is good enough in the first $n - 1$ steps of this process to result in justified belief, why wouldn't it be good enough in the n th step to result in justified belief (even if it could not result in *knowledge*)? One answer to this question may be that by the point the author gets to the last disjunct, she can double-check the embedded claim. If the double-checking fails, then she will not believe the last blindspot, and if the double-checking justifies the last proposition—well, what then? Is she to believe that she is the first one ever to write a completely error-free book? Should she start planning for the global accolades that are sure to come (Christensen, 2004)? No, of course not: rather, she should think that she is still as fallible as the rest of us, and that the fallibility seeped into her double-checking procedures. So, if she double-checks the claim embedded in the last blindspot, she would either go back to believing a disjunction of blindspots or no blindspot at all. But what if she doesn't double-check the claim embedded in the last blindspot? Isn't she then justified in believing it? Maybe—or maybe she is not, because she is already (propositionally) justified in believing that there is nothing special about that last claim, and so that she should have the same attitude toward it as she does toward all the others.

Perhaps, then, the preface problem does not present a watertight argument for justified belief in blindspots. What about its cousin, the lottery problem? Suppose that Tim is convinced that he does not know, merely on the basis of statistical evidence, that his lottery ticket is a loser. Nevertheless, Tim is justified in believing that his lottery ticket is a loser. He puts two and two together and believes “My ticket is a loser but I don't know that it is”. Of course, he doesn't know this, but what is the argument that he doesn't justifiably believe it?⁹

In any case, regardless of whether either the preface or the lottery are examples of justifiably believed blindspots, we think that *Bamboozled* certainly is. Testimony is an extremely powerful source of justification, and in the right circumstances, it can justify subjects in believing blindspots. What of the worry that Rhys can himself figure out that the proposition in question is a blindspot? Doesn't that in itself count against the possibility of his justifiably believing it? Not really: Rhys already believes that Yul Brynner was bald, that Brian May is not bald, and that one hair doesn't make the difference between being bald and not being bald. Rhys realizes that this set of beliefs entails a contradiction. Nevertheless, he neither believes a contradiction nor does he give up his belief in the inconsistent triad. Many philosophers have tried to convince Rhys that he should give up

⁹ One argument against both our preface and lottery examples comes from extreme versions of knowledge-first epistemology which identify justification with knowledge. This is not the place to argue against such views, so the reader is invited to either agree with us regarding the implausibility of such positions or to take the results of this note to be some consequences of the rejection of those positions.

his belief in the tolerance principle that one hair doesn't make the difference between being bald and not being bald, but Rhys is more confident of that tolerance principle than he is of any philosophical theory against it. We think Rhys may very well be justified in this set of attitudes. If he is justified in believing each one of an inconsistent set of propositions, though, why would he not be justified in believing a blindspot?

One reaction to our claim that Rhys knows that there is an external world by way of competent deduction from a justifiably believed blindspot proposition was that such propositions cannot even be justifiably believed. This is too strong to be plausible, as the above considerations show. Another reaction is that, if Rhys knows that there is an external world, he knows it in some other way, not via competent deduction from a blindspot proposition.

In some respects, this second reaction is related to Coffman's idea that nearby truths are doing all the heavy lifting. We now turn to that objection.

4. The Justificatory Heavy Lifting

We claimed that Rhys can know that there is an external world by inferring it from the justifiably believed blindspot proposition *There is an external world but I do not know it*. But if one could locate a nearby proposition, one that Rhys was in a position to know, to do the justificatory heavy lifting, our claim would lose some of its appeal. So, what are some candidate nearby knowable propositions, and how does this "knowledge *despite* blindspots" strategy work?

Consider, as a surrogate, the nearby proposition *Juan said that there is an external world but I do not know it*. This proposition is not in Rhys's blindspot; there is, after all, no felt inconsistency, Moorean absurdity, or other abominableness to it. Compare: as Rhys looks in the fridge, he sees and thereby comes to know that there's no beer. Juan might, from the couch, insist that there's beer (maybe he remembers seeing some just yesterday). In some such situations, of course, Juan's testimony might undermine Rhys's knowledge that there's no beer, perhaps prompting him to further rummage around in pursuit of some. But in other situations, ones in which Rhys has already searched high and low, or ones in which the fridge is empty and so the presence or absence of beer would be obvious, Rhys is in a position to know the conjunction *Juan said there's beer but there's no beer*.

The issue is now to determine whether *Juan said that there is an external world but I do not know it* can really do the justificatory heavy lifting. The inference would have to go via Rhys' trust in Juan's testimony. But this way of recasting Rhys's predicament seems to introduce a new problem. Namely, that Juan, insofar as he is providing Rhys with evidence via testimony, is providing evidence for both the proposition that Rhys doesn't know that there is an external world *and* the proposition that there is an external world. If, instead, what Rhys has as evidence is merely that Juan says so, this should apply to each conjunct to which Juan testifies. Rhys would be left with the proposition *Juan said that I don't know there is an external world and Juan said there is an external world*. But from this, Rhys could only deduce that there is an external world by way of a background premise to the effect that Juan is a reliable informant, a premise which would interact with the first conjunct no less than to the second so as to support precisely that blindspot proposition this strategy was meant to circumvent.

We think this points to a general problem with the "knowledge *despite* blindspots" strategy, concerned as it is with looking for nearby knowable proposition to serve as explanatory surrogates

for the relevant blindspot proposition: as we already said, the *falsity* of the initial proposition is an essential part of what puts Rhys in a position to know that there is an external world. After all, if the initial proposition is true, then it's true that Rhys doesn't know there's an external world.

At this point, a friend of “No False Lemmas” or of “Counter-Closure” might be borderline exasperated. As if it weren't bad enough that we argued that one could justifiably believe a proposition in one's blindspot, we also argued that attempts to find nearby, knowable propositions to do the justificatory heavy lifting were doomed to fail. The problem, they might suggest, resides in the initial suggestion that Rhys knows there is an external world *via inference* from such a problematic proposition. Instead, more plausibly (the suggestion continues), Rhys knows there is an external world via whatever way he acquired justification for believing the blindspot in the first place. If that's right, the problem of knowledge from blindspots doesn't even get off the ground.

One way to make this sentiment more precise, we'll call the “disqualification strategy”. According to it, even though Juan's testimony provides evidential support for Rhys to believe the blindspot proposition, this evidence is *disqualified*, so Rhys's belief in the external world on its basis is irrational.

Evidence E in favor of a proposition P is *disqualified* by other evidence E' when E' is stronger—or perhaps somehow “more direct”—evidence in favor of P than is E.¹⁰ For instance, Alice tells Bob that there is beer in the fridge, and that this is sufficient evidence in Bob's situation to justify him in believing it. Suppose that Bob opens the fridge and sees a six-pack of Hazy IPAs. Seeing a six-pack of Hazy IPAs is also sufficient evidence to justify Bob in believing that there is beer in the fridge. Although both Alice's testimony and Bob's perceptual experience are each sufficient to justify Bob's belief, Bob's perceptual evidence *disqualifies* his testimonial evidence from serving as the (a?) basis for his belief. According to this line of thought, Bob ought to base his belief that there is beer in the fridge on his perceptual experience but not on Alice's testimony because the former is stronger or more directly relevant to the question at hand, namely whether there's beer.

Returning to the original example, Rhys has overwhelming evidence that there is an external world via his perceptual experiences; perhaps the evidence provided by perceptual experience in favor of the external world is far stronger than any evidence that could be provided by testimony or by logical inference. Consequently, Rhys' belief that there is an external world is improperly based; if Rhys ought to believe in an external world, he ought to believe in an external world *on the basis of his perceptual experiences*, not on the basis of testimony-plus-inference.

The problem with the disqualification strategy is that, even if in some sense Rhys would be better off by basing his beliefs differently, this does not show that Rhys, as he in fact is, is epistemically *irrational*; it only shows that Rhys is epistemically *suboptimal*. For instance, in his analysis of disqualification, Muñoz appeals to Harman's (1986) *Principle of Clutter Avoidance* in order to defend the claim that we ought not to base our beliefs on disqualified evidence. *Qua* thesis about epistemic optimality, the *Principle of Clutter Avoidance* is compelling. *Qua* thesis about epistemic rationality, the *Principle of Clutter Avoidance* is controversial. For instance, claiming that the

¹⁰ See Muñoz (2019). Disqualifiers “take a would-be justifier and make it irrelevant” (888). In this respect, disqualification is similar to defeat and evidential screening-off (on the latter, see Weatherston, 2019, ch. 11). The question of whether disqualification is, as Muñoz argues, irreducibly distinct from defeat and screening-off we set aside.

Principle of Clutter Avoidance places constraints on epistemic rationality would conflict with many forms of Evidentialism, according to which one's beliefs (at a time) are justified by one's total evidence (at that time), since it would require that agents ignore some of the evidence of which they could avail themselves.

An agent can violate the *Principle of Clutter Avoidance* while also having all of their beliefs supported by the evidence. For example, we can imagine that, for some reason, Douglas cannot get rid of their knowledge of all of the starting lineups of the Chicago Bears from 2000 to 2010. Given that Douglas is a finite creature, this leads to unfortunate situations where Douglas is unable to learn, or maintain knowledge in, more important propositions, in part, due to the fixity of their knowledge of the Bears' starting lineups at the beginning of the century. However, even if this is all true, this does not make Douglas epistemically irrational for believing that safety Mike Brown only played in six games for the Bears in the 2006 season. Indeed, he *knows* it despite failing to avoid clutter.

Similarly, we think for an agent who believes on the basis of disqualified evidence. While there may be something odd and perhaps suboptimal about the way that Rhys comes to believe that there is an external world, this does not undermine the claim that Rhys' ultimate belief in the external world is epistemically rational. Furthermore, it does not undermine the claim that Rhys' belief in the external world amounts to knowledge. Just like how Bob can know that there is beer in the fridge on the basis of Alice's testimony even while he's looking at beer, Rhys can know that there is an external world via inference from belief in a blindspot even while he's perceiving the external world. Thus, the disqualification strategy fails to secure the verdict that Rhys cannot know that there is an external world via inference from a proposition in his blindspot.

Still, one could think that it is not possible for Rhys to be justified in believing the blindspot proposition without *antecedent* justification for believing that there is an external world, and if Rhys has such antecedent justification, then Rhys' knowledge from blindspot is, if not disqualified, at the very least superfluous.

To address this worry, let us first note that we have so far been studiously avoiding the question of what kind of inference Rhys performs in going from his justified belief that there is an external world but he doesn't know it to his knowledge that there is an external world. Perhaps the obvious answer is that Rhys performs conjunction elimination, because the blindspot proposition is a conjunction of the propositions *There is an external world* and *Rhys doesn't know that there is an external world*. But it need not be this way. Suppose that, before presenting the skeptical arguments, Juan introduced a super-factive operator THAT. Importantly, THAT *p* is super-factive because it entails *p* even when embedded in contexts which would normally cancel those kinds of implications. It is thus unlike "it is true that", for *I believe that it is true that it is raining* does not entail that it is raining, but *I believe THAT it is raining* does entail that it is raining. What Rhys is convinced of, then, is the proposition *I do not know THAT there is an external world*. Although that proposition has an interesting structure in virtue of embedding a super-factive operator, it is not (we can suppose) a conjunction. Thus, it is not that Rhys first becomes justified in believing a conjunction one of whose conjuncts is that there is an external world, and then performs conjunction elimination to come to know that there is an external world. Perhaps it is not possible to know a conjunction without *thereby* knowing its conjuncts. But it is perfectly possible to learn that, say, Argentina won the 1986 Soccer World Cup on the basis of Juan saying "It has been reported THAT Argentina won the 1986 Soccer World Cup".

And that brings us to another important remark. We have so far concentrated on Rhys coming to know that there is an external world, but there is no need for the example to involve such a heavy-weight philosophical thesis. Instead, suppose that, on the basis of skeptical arguments, Juan convinces Rhys of the proposition *I don't know THAT Argentina won the 1986 Soccer World Cup*. That Argentina won the 1986 Soccer World Cup is not a proposition whose knowledge must antecede any other knowledge, and so that reason for supposing that knowledge from blindspots is impossible does not really get to the heart of the matter.

5. Defective Knowledge

The arguments in Sections 3 and 4 purported to cast doubt on two popular and mutually supportive theses: NFL and Counter-Closure. As a reminder, they say:

No False Lemmas: Necessarily, S's belief that p is knowledge only if it is not inferred from any falsehood.

Counter-Closure: Necessarily, if (i) S knows that p entails q and (ii) S comes to believe q solely on the basis of competently deducing it from p, and (iii) S knows q, then S knows p.

Our position is at odds with NFL and Counter-Closure because knowledge from blindspots is inferential knowledge from a proposition that is itself unknowable and false.

Zooming out from our particular examples, we hope to show that our position—that there is sometimes knowledge from blindspots—has wider appeal and motivation.

As might be apparent from the discussion in the last section, we think that an agent's being able to gain knowledge from a proposition in his blindspot is one among many ways of defective or non-ideal knowledge-acquisition, the others of which most of us are already familiar. Consider, for example, epistemic akrasia, a special sort of inner conflict wherein an agent believes both p and that her evidence fails to support p. Akratic agents are, like Rhys or Tim, liable to utter Moore-paradoxical-sounding sentences like "P, but I shouldn't believe that P"; whether or not these are truly abominable or have the same felt inconsistency as "P, but I don't know that P", there is surely a family resemblance between akratic agents and those who believe propositions in their blindspots.

While this is not totally uncontroversial, we think, along with many contemporary epistemologists, that some cases of epistemic akrasia are rationally permissible. Here is one case (of many) discussed recently by Hawthorne et al. (2021):

Unger Games: Following Unger, Stew believes that a belief in p is rational only if one has reasons for believing p and that something can be a reason for belief only if it is known. Moreover, Stew reasonably trusts an epistemologist—Peter Unger, in fact—who tells him that knowledge is unachievable. He thus believes that none of his beliefs are rational, thinking that the best that can be hoped for is some lesser status. But Unger has got all this wrong, and in fact many of Stew's beliefs are rational. (adapted from Hawthorne et al., 2021, p. 4)

Peter Unger convinces his impressionable student, Stew, that no one rationally believes anything because rational belief (*that there are dark clouds forming above Tucson*, say) requires a preponderance of epistemic reasons that no one in fact possesses, nor could possess. Stew might, after introspectively attending to his occurrent commitments, think to himself *there are dark clouds forming above Tucson*, but (*given Unger's lecture*) *I shouldn't believe it*. Suppose that Stew, in the midst of his *aporia*, infers that it will probably rain in Tucson later that day, so that he should ensure he has an umbrella handy. Perhaps his inference is a bit lucky in the sense that, had he spent more time reflecting on Unger's lecture and the misleading higher-order evidence it provides, he may have hesitated and remained agnostic about rain. Perhaps also his inference is defective because it flows from a state of inner conflict, global coherence being an ideal of epistemic rationality.

But, as one of us has argued, at least some cases like Stew's are ones of inadvertent epistemic virtue, wherein one forms the right belief for the right reasons while thinking otherwise.¹¹ In particular, Stew infers for reasons he thinks do not suffice to justify his inference. This, of course, is compatible with claiming that Stew cannot know the conjunction *there are dark clouds forming above Tucson*, but (*given Unger's lecture*) *I shouldn't believe it*. Even if, that is, Stew is in the unfortunate position of being unable to know what he believes, he can still generate knowledge from this position by way of manifesting certain epistemic virtues, responding to the right reasons being chief among them.

Stew's predicament in *Unger Games* is only superficially different from Michael's in *Churchlandia*:

Churchlandia: Following Churchland, Michael believes that there are no beliefs; completed cognitive science has no use for such folk psychological notions. Moreover, Michael reasonably trusts a philosopher—Paul Churchland, in fact—who tells him that beliefs are nothing but a fiction borne of ignorance and social utility. Michael thus believes that none of his (or anyone's) mental states are beliefs. But Churchland has got all this wrong, and in fact some of Michael's mental states are beliefs, the belief-seeming ones chief among them.

Michael might, after introspectively attending to his occurrent commitments, think to himself *there are dark clouds forming above Tucson*, but (*given Churchland's lecture*) *I don't believe it*. This is a proposition in Michael's blindspot. Suppose further Michael, in the midst of his *aporia*, infers that it will probably rain in Tucson later that day, so that he should ensure he has an umbrella handy.

Like Stew, Michael's inference might be thought lucky and defective: lucky because it could have easily been led astray by Churchland's misleading testimony, and defective because it falls short of an ideal of coherence. Nevertheless, Michael might display inadvertent virtue in inferring as he does; after all, he forms the right belief for the right reasons, while thinking otherwise. In particular, Michael infers from reasons he doubts exist (assuming that reasons for this or that inference are at least partly—if not wholly—constituted by one's beliefs).

The parallels between Stew and Michael, on the one hand, and Rhys and Tim, on the other, suggest that, at least sometimes, agents can gain inferential knowledge despite the defectiveness or non-ideality of the state from which they infer.

¹¹ Kearl (2022). See also Weatherson (2019).

6. Conclusion

We have argued for the possibility of knowledge from blindspots. This presents adherents of NFL and Counter-Closure with a distinctive challenge, since, not only is the blindspot in fact false, but also its falsehood is essential. In other purported cases of knowledge from falsehood, it has been argued that there are nearby true propositions with sufficient justificatory power to provide the agent with knowledge, rendering the falsehood inessential—it is a case of knowledge *despite* falsehood. This strategy of defending NFL cannot work against our cases, since cases of knowledge from blindspot satisfy not only Fitelson's counterfactual condition, but also our stronger essentiality condition: necessarily, knowledge is lost if the initial proposition is true.

Not only does knowledge from blindspots undermine NFL and Counter-Closure, but also it is interesting in its own right. It exemplifies one way in which we can gain inferential knowledge *defectively* or *non-ideally*. Viewed against this backdrop, perhaps what is most attractive about *NFL* and *Counter-Closure* is that they codify certain paradigmatic cases of inferential knowledge. But that still permits a great deal of wiggle room for discussing the varieties of non-paradigmatic inferential knowledge, perhaps understood in terms of theoretical distance from those ideals.

References

- Ball, B., & Blome-Tillman, M. (2014). Counter-closure and knowledge despite falsehood. *Philosophical Quarterly*, 64, 552–568.
- Christensen, D. (2004). *Putting logic in its place*. Oxford University Press.
- Church, A. (2009). *Referee reports on Fitch's 'A definition of value'* (pp. 13–20). Edited by J. Salerno (ed.).
- Coffman, E. J. (2008). Warrant without truth? *Synthese*, 162(2), 173–194.
- DeRose, K. (1995). Solving the skeptical problem. *The Philosophical Review*, 104(1), 1–52.
- Fitch, F. (1963). A logical analysis of some value concepts. *The Journal of Symbolic Logic*, 28, 135–142; reprinted in Salerno (ed.) 2009, 21–28.
- Fitelson, B. (2010). Strengthening the case for knowledge from falsehood. *Analysis*, 70(4), 666–669.
- Gettier, E. L. (1963). Is justified true belief knowledge? *Analysis*, 23, 121–3.
- Harman, G. (1973). *Thought*. Princeton University Press.
- Harman, G. (1986). *Change in view: Principles of reasoning*. MIT Press.
- Hawthorne, J., Isaacs, Y., & Lasonen-Aarnio, M. (2021). The rationality of epistemic akrasia. *Philosophical Perspectives*, 35, 206–228.
- Kearl, T. (2022). Evidentialism and the problem of basic competence. *Ergo*, 9(0). doi: <https://doi.org/10.3998/ergo.2271>

- Klein, P. (2008). Useful false beliefs. In Q. Smith (Ed.), *Epistemology: New essays* (pp. 25–62). Oxford University Press.
- Lee, K. (2021). Reconsidering the alleged cases of knowledge from falsehood. *Philosophical Investigations*, 44(2), 151–162.
- Luzzi, F. (2010). Counter-closure. *Australasian Journal of Philosophy*, 88(4), 673–683.
- Montminy, M. (2014). Knowledge despite falsehood. *Canadian Journal of Philosophy*, 44(3–4), 463–475.
- Moore, G. E. (1942). A reply to my critics. In P. A. Schilpp (Ed.), *The philosophy of G. E. Moore*. Northwestern University.
- Muñoz, D. (2019). Defeaters and disqualifiers. *Mind*, 128(511), 887–906.
- Poncins, G., (1941) [1988], *Kabloona* in collaboration with Lewis Galantieri. Carroll & Graff Publishers.
- Schnee, I. (2015). There is no knowledge from falsehood. *Episteme*, 12(1), 53–74.
- Sorensen, R. (1988). *Blindspots*. Clarendon Press.
- Warfield, T. A. (2005). Knowledge from falsehood. *Philosophical Perspectives*, 19, 405–416.
- Weatherson, B. (2019). *Normative externalism*. Oxford University Press.