Episteme (2025), 1–26
doi:10.1017/epi.2025.10070

episteme

ARTICLE

Knowledge from Falsehoods Reconsidered

Matteo Baggio

Università degli Studi di Torino, Torino, Italy
Email: matteo.baggio@unito.it

(Received 16 April 2025; revised 3 July 2025; accepted 10 July 2025)

Abstract

Recent epistemological debates have increasingly focused on the contentious counter-closure principle, which holds that, necessarily, if an agent *S* believes *q* solely on the basis of a competent inference from *p*, and *S* knows *q*, then *S* also knows *p*. This principle has drawn attention due to various challenges, particularly the issue of inferential knowledge derived from false premises. In this article, we pursue two objectives. First, we argue that the counter-closure principle is untenable but for reasons that depart from traditional critiques. Specifically, we will present a novel argument against the internalist approach that supports the cases of knowledge from falsehoods. Second, we show that the counter-closure principle's failure can be better addressed within an externalist framework by exploring novel theories of defeaters and the relationship between doxastic and propositional warrant.

Keywords: Counter-closure principles; inferential knowledge; propositional warrant; doxastic warrant; defeaters; epistemic internalism; epistemic externalism

1. Introduction

In recent years, the epistemological debate has increasingly focused on the contentious counter-closure principle (*CC*), which posits that:

> *CC* Necessarily, if an agent *S* believes *q* solely on the basis of competent inference from *p*, and *S* knows *q*, then *S* knows *p*.

At first glance, *CC* seems quite plausible. Indeed, many epistemologists regard it as an intuitive and largely uncontroversial principle of inferential knowledge. Furthermore, several persuasive philosophical arguments lend support to *CC*. For example, *CC* resonates with the intuitive idea that inference serves as a mechanism for *transmitting* epistemic goods, such as epistemic warrant or knowledge (cf. Nozick 1981). According to this view, if *S* believes *q* based solely on a competent inference from *p* and *p* does not instantiate epistemic goods, then there would be no epistemic goods to transmit to the conclusion *q*. Therefore, for *S* to know *q* solely through inference, *S* must also know *p*.

© The Author(s), 2025. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/epi.2025.10070 Published online by Cambridge University Press

Despite its appeal, *CC* faces several challenges. Luzzi (2019, 6–7) provides a general framework for constructing counterexamples to *CC*. Below is a comprehensive reconstruction of Luzzi's framework:

1. **Partial fulfillment of conditions**: *S* satisfies all but one condition – condition *X* – necessary for knowing a premise *p*. Consequently, *S* does not know *p*, even though their epistemic position regarding *p* is strong. If *S* had satisfied condition *X*, they would have known *p*.
2. **Non-overdetermination of belief**: *S*'s belief in the conclusion *q*, inferred from *p*, rests solely on the inference from *p* without additional reasons. If *S* had other reasons for believing *q*, *CC* would remain unchallenged, as *q* would have support independent of *p* lacking condition *X*.
3. **Excellent epistemic standing**: *S*'s position regarding *q* is strong; *q* meets all the same requirements as *p* but also satisfies condition *X*, unlike *p*.

Given this general framework, it's not surprising that various challenges to *CC* have emerged. For example, Warfield (2005) and Klein (2008) replace *X* with the truth condition, arguing that inferential knowledge can stem from a false yet warranted premise. Further, Luzzi (2010) substitutes *X* with safety and sensitivity conditions, proposing that knowledge can be inferred from a Gettierized premise. Finally, Murphy (2013) replaces *X* with the belief criterion, arguing that knowledge can arise from an unbelieved premise.[1] Together, these examples align with the profile outlined in points 1 through 3, suggesting that inferential knowledge can be derived from premises that do not qualify as knowledge.

This article addresses the first major challenge to *CC*: the case of inferential knowledge derived from falsehoods (*KFF*). We argue that *CC* is untenable but approach this conclusion differently than most *KFF* proponents. Advocates of *KFF* typically claim that *KFF* cases constitute genuine instances of inferential knowledge, thereby directly challenging *CC*. However, their arguments often rest on an internalist conception of epistemic warrant – a framework we contend is fundamentally flawed. In contrast, we argue that an externalist account of epistemic warrant offers a stronger and more compelling basis for rejecting *CC*.

The article is structured as follows. In §2, we introduce some *KFF* cases, and then we critically examine the primary argument against *KFF*, i.e., the proxy premise argument (§2.1). Next, we show how *KFF* supporters might leverage ideas within this argument to strengthen their position (§2.2). In §3, we aim to demonstrate that the *KFF* explanation discussed in §2.2 relies on an internalist approach (§3.1) and reveal the incompatibility between this internalist view and an externalist approach (§3.2). We conclude this section by arguing against the internalist *KFF* explanation (§3.3). In §4, we reexamine the externalist approach, with particular emphasis on the broad account presented by Graham and Lyons (§4.1). Ultimately, we argue that the *KFF* thesis can be more effectively reconciled with a revised interpretation of the externalist approach (§4.2).

2. *KFF* and the proxy premise argument

Let us consider the following scenarios:

[1]For a challenge to the *WFW* principle, which claims that every premise must be doxastically warranted for *S* to have inferential warrant for the conclusion, see Murphy (2015). Additionally, for a historical perspective on *CC*, see Saunders and Champawat (1964) and Hilpinen (1988).

https://doi.org/10.1017/epi.2025.10070 Published online by Cambridge University Press

Fancy watch

> I have a 7 p.m. 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 2:58 p.m.; therefore, I am not late for my 7 p.m. meeting." [. . .] I know my conclusion, but as it happens it is exactly 2:57 p.m., not 2:58 p.m. (Warfield 2005, 408).

Party hats

> Liz [carefully] counts 35 children at her son's birthday party and concludes that the 100 party hats she bought for the party are enough. However, there are 36 children at the party – one child ran to the bathroom after Liz started counting heads (Borges and Schnee 2023, 1; adapted from Warfield 2005, 407).

Ford car

> Nogot, Havit, and I are classmates. One day, I see Nogot driving a Ford, parking it in his garage, and claiming it's his car. Based on this, I conclude that Nogot owns a Ford, so someone in our class owns a Ford. Now, imagine that while it's true a classmate owns a Ford, it turns out not to be Nogot – it's Havit, about whom I have no beliefs regarding car ownership (adapted from Lehrer 1965).

These vignettes share a common structural feature: in each case, agent *S* draws a true conclusion from a false premise. Yet, there is also an important epistemic difference between these scenarios. In the first two, *S* appears to know the conclusion despite the false premise, while in the third, *S* does not know the conclusion, even though it is true. This epistemic discrepancy raises questions about a widely held assumption in epistemology, particularly in relation to Gettier cases such as the *Ford car* scenario. Traditionally, it is assumed that *S* cannot possess knowledge derived (solely) from a false belief. However, as the first two cases suggest, the relationship between false premises and knowledge is more nuanced than this traditional view allows.

While *KFF* cases, like *Fancy watch* and *Party hats*, seem intuitively convincing in terms of their epistemic outcomes, their structural similarity to classic Gettier cases requires further explanation from *KFF* proponents. Specifically, they must clarify why *KFF* cases should be considered instances of inferential knowledge rather than ignorance. In the following subsections, we will explore a widely accepted explanation offered by *KFF* advocates to support this distinction.

2.1. The proxy premise argument

KFF and *CC* proponents traditionally share the intuition that cases like *Fancy watch* and *Party Hats* are genuine examples of inferential knowledge. However, while supporters of *KFF* accept these cases as genuine instances of inferential knowledge, they do so at the cost of the *CC* principle. In contrast, *CC* advocates maintain that the principle's plausibility remains intact. To defend *CC*, these authors typically argue that the false premise does not play any epistemic role in such cases. Instead, only a known premise truly serves this function. This line of defense is known as the *proxy premise* argument.

Now, let's examine the proxy premise argument put forward by defenders of *CC*. As we will soon explore, this strategy may ultimately offer stronger support for *KFF* cases

https://doi.org/10.1017/epi.2025.10070 Published online by Cambridge University Press

than for *CC* itself. Before diving into that, however, it's important first to address the objection raised against *KFF* cases.

Traditionally, proponents of *CC* address the problem posed by *KFF* cases by asserting that a dispositional true belief – called the "proxy premise" – closely related to the false belief provides the sole epistemic basis for *KFF* inferences. According to this view, the proxy premise must *logically follow from*, or be *evidentially supported by*, the false premise and must also be *dispositionally known* by *S* (cf. Fitelson 2010; Ball and Blome-Tillmann 2014; Montminy 2015, 2023). For instance, in the *Fancy watch* case, a plausible proxy premise might be, "It is approximately 2:58 p.m."; in the *Party hats* case, it could be, "There are fewer than 100 children."

Now, to prevent the risk of epistemic overdetermination between the false belief and the proxy premise in *KFF* cases, *CC* advocates need to argue that the proxy premises are not just useful but also indispensable for establishing inferential knowledge. To do so, they claim that certain beliefs must be essential for *S* to know a conclusion inferentially. To ratify this requirement, *CC* advocates introduce the following condition:

> **Essentiality Condition**: A premise *p* is essential if, in its absence from *S*'s belief set, *S* would no longer know *q*.[2]

Within this framework, *CC* advocates avoid epistemic overdetermination by positing that the proxy premise – being both dispositionally known and epistemically robust – constitutes the sole essential ground for *S*'s inferential knowledge of *q*.

For instance, consider a *KFF* scenario. According to *CC* advocates, if *S* genuinely knows the conclusion *q*, *S* is dispositionally warranted in believing the proxy premise, such as "It is approximately 2:58 p.m." This belief, they argue, is evidentially supported by both the watch's reading and is entailed by the false premise that *S* explicitly believes, namely, "It is exactly 2:58 p.m." However, according to *CC* advocates, the false premise does not satisfy the essentiality condition. After all, if removed from *S*'s belief set, *S* would still know the conclusion because it remains grounded in the dispositionally warranted proxy premise. Indeed, the proxy premise remains essential to the inference: without it, *S* would lose the inferential knowledge of *q*, which *CC* advocates seeking to preserve. Moreover, the proxy premise is also dispositionally known by *S*, as accurate watches are typically slightly off rather than drastically inaccurate (cf. Montminy 2023, 66).

In light of this, *CC* advocates argue that the false premise becomes epistemically irrelevant in the inferential process. Focusing on the dispositionally known proxy premise, they argue that it alone supports *S*'s inferential knowledge, defending *CC* by arguing that the inference relies on true, known premises.

Without delving into the finer details of this argument, we wish to emphasize what we consider its central issue. According to advocates of *CC*, for *S* to know the conclusion of the alleged *KFF* inference, *S* must possess dispositional knowledge of a proxy premise (e.g., "It is approximately 2:58 p.m."). Crucially, however, *CC* proponents maintain that this proxy premise must not merely co-occur with the false belief but must either logically follow from it or be evidentially supported by it. In other words, the proxy premise derives its epistemic credibility precisely from its relationship to the false belief, "It is exactly 2:58 p.m." This dependency is not incidental but essential – without the false belief, the proxy premise loses its epistemic support. Thus, by their own epistemic standards, the false belief remains epistemically indispensable. As a result, their

[2]Montminy (2015) terms this the "Klein condition" to highlight its use by *KFF* proponents in explaining inferential knowledge.

https://doi.org/10.1017/epi.2025.10070 Published online by Cambridge University Press

purported solution reintroduces the very problem it aimed to resolve, since knowledge of the proxy premise itself becomes entangled in another instance of *KFF*.[3]

Moreover, as Luzzi (2019, 23) rightly observes, even if one were to argue that *S*'s original reasons provide independent dispositional warrant for the proxy premise (e.g., "My watch reads 2:58 p.m." supporting "It is approximately 2:58 p.m."), a pressing question remains: in an indistinguishable scenario where it is actually 2:58 p.m., why would *S*'s inferential knowledge be based on the correct (precise) premise, while when the premise is false, knowledge shifts to rely on an (approximate) proxy premise? This inconsistency further undermines the proxy premise argument. After all, it seems highly counterintuitive to suggest that the accuracy of the clock – the only changing variable in this new envisaged scenario – could somehow dictate the reasoning process of *S*.

In conclusion, the challenges posed by *CC* proponents highlight a critical issue in their reliance on the proxy premise argument. While they assert that this premise serves as the essential basis for inferential knowledge in *KFF* cases, they inadvertently open themselves to the risk of *KFF* at a meta-level if they lean too heavily on it. Alternatively, if they argue the proxy premise stands independently, they must confront a troubling thesis about the connection between real-world facts and inferential reasoning. Thus, these considerations reveal deeper complexities in understanding how we derive knowledge from premises, whether true or false.

2.2. Toward a (preliminary) solution

Despite objections from *CC* supporters, *KFF* advocates have found a valuable ally in the proxy premise idea. As a matter of fact, *KFF* proponents have skillfully turned some of the criticisms of the proxy premise argument into a strong defense of their position. Notably, both Klein (2008) and Feit & Cullison (2011) present a case for *KFF* that aligns with some ideas of the proxy premise argument. Specifically, they argue that, unlike in Gettier cases, knowledge can still be derived from a false premise in *KFF* cases because *a warrant* for the conclusion remains intact, even if the subject later discovers their belief in p is false.

To illustrate, consider the *Fancy watch* case. Suppose *S* learns the time is not exactly 2:58 p.m. Despite this, *KFF* proponents argue *S* can still know, "I am not late for the 7:00 p.m. meeting," based on a dispositionally warranted belief in a related but approximate premise, like "It is approximately 2:58 p.m." This suggests that even with contradictory information, the dispositional warrant for the approximate belief is strong enough to support the conclusion.

Now, contrast this with the *Ford car* case. Here, *S* believes "Someone in our class owns a Ford," based on the incorrect premise that "Nogot owns a Ford." However, in this case, *S*'s belief does not qualify as knowledge because the foundational grounds are false. If *S* learns "Nogot does not own a Ford," this removes any warrant for the belief that "Someone in the class owns a Ford." Crucially, *S*'s initial reasons cannot transfer to an alternative premise that would preserve their inferential knowledge after the false premise is compromised. Therefore, *S*'s inferential belief does not count as knowledge in this case.

To clarify these ideas further, we will now introduce some well-known technical terms that will be useful in our subsequent discussion. The next section will examine these terms more in-depth.

The main argument posited by proponents of *KFF* appears to be the following: the same initial reasons that provide a *doxastic warrant* for a false belief p also provide a

[3]Luzzi attributes to Arnold (2013) the first formulation of this counterargument.

https://doi.org/10.1017/epi.2025.10070 Published online by Cambridge University Press

propositional warrant for a related, approximate true proposition *a*. Now, since *a* logically follows from *p* – because any precise proposition p entails its approximation *a* – the two remain closely connected from an epistemological standpoint. As a result, even if *S* encounters a defeater for *p* (for example, learning that "The time is not exactly 2:58 p.m., but 2:57 p.m."), their inferential knowledge of *q* remains intact. This is because *p*, despite being false, serves as a premise in the inference to *q* and warrants *q* propositionally in much the same way *a* does. Thus, in a counterfactual scenario where *S* acquires a defeater for *p*, the knowledge of *q* is preserved. In this way, a false premise can still generate knowledge, provided its negation does not serve as a defeater. Therefore, even when the doxastic warrant for *p* is defeated, the propositional warrant for *a* continues to support the inference.[4] Below is a graphical representation of these ideas (see Figure 1).

Hence, *KFF* advocates argue that while secondary reasons (*a*) may serve as a fallback in counterfactual situations where primary reasons (*p*) fail, the primary reasons still support *S*'s inferential knowledge in the actual world. The distinction is crucial: counterfactual reasons do not carry the same epistemic weight as real-world reasoning. Therefore, *KFF* supporters stress the need to differentiate between counterfactual support (*a*) and the actual basis of *S*'s inferential knowledge (*p*).

The above considerations explain why *KFF* cases constitute knowledge, unlike Gettier cases. The solution relies on the role of defeaters and the distinction between doxastic and propositional warrant. These concepts distinguish cases where knowledge arises from false but *stable* premises (cf. Warfield 2005) from those where true belief is merely a product of *luck* (cf. Pritchard 2005). This raises an important question: how general is this solution, epistemologically speaking? Specifically, is it externalistically kosher? In the next section, we will argue that this is not the case.

3. *KFF* through the lens of epistemic internalism

This section serves two main purposes. First, it aims to demonstrate that while the traditional explanation for the plausibility of *KFF* cases might be convincing from an internalist perspective, it falls short when viewed from an externalist standpoint. Second, it will argue that, contrary to initial impressions, the internalist strategy faces far more significant challenges than the externalist one. We begin by outlining the reason-first approach that underpins the traditional explanation for *KFF* cases.

3.1. The reason-first approach

Proponents of *KFF* address the proxy premise argument by distinguishing between doxastic and propositional warrant. This distinction is traditionally introduced as follows:

- *S* has *propositional warrant* to believe proposition *p* iff they have sufficient epistemic reasons to do so, regardless of whether they actually hold that belief.
- *S* is *doxastically warranted* in believing *p* iff they have propositional warrant for *p*, believe *p*, and base their belief on these sufficient epistemic reasons (cf. Firth 1978).

[4]If *S* learns that ¬*p*, they can still rely on *a*, even though they lose *p*, because they possess propositional warrant for "If it is not exactly 2:58 p.m., then it is approximately 2:58 p.m." As we will explore shortly, this conclusion hinges on some problematic aspects of the reason-first account of epistemic warrant.

https://doi.org/10.1017/epi.2025.10070 Published online by Cambridge University Press

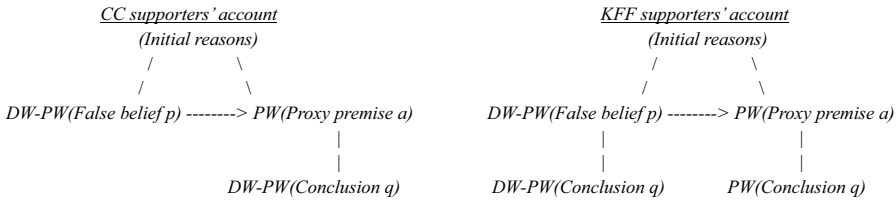


Figure 1. These graphs illustrate the epistemic support relations for *CC* (left) and *KFF* (right) supporters. Dotted horizontal arcs indicate entailment between propositions *p* and *a*, while vertical and diagonal arcs represent evidential or inferential dependencies. *DW* and *PW* refer to doxastic and propositional warrant, respectively.

Put simply, doxastic warrant is propositional warrant combined with its connection to *S*'s belief, specifically through the basing relation. Under this account, propositional warrant is conceptually fundamental: understanding the concept of doxastic warrant requires first grasping the concept of propositional warrant (cf. Melis 2018). This framework is often referred to as the *reason-first approach* (cf. Silva and Oliveira 2023).

The reason-first approach to epistemic warrant stems from an internalist view of non-factive positive epistemic support, which has long been dominant among epistemologists. Philosophers such as Ayer, Russell, Carnap, and C. I. Lewis, building on the work of modern thinkers, supported the *sense-data* theory of reasons. *In nuce*, this theory posited that empirical knowledge and its warrant depend entirely on internal mental states and procedural norms – meaning rules – that allegedly align an individual's sensory experiences with the world's facts. However, this theory declined due to the challenge of identifying these procedural norms.[5]

In response, epistemologists shifted towards a new dominant view: epistemic warrant is grounded exclusively in reasons. According to this view, *S* has propositional warrant for *p* just in case *S*'s psychological states support *p*. However, under this account, having adequate reasons for *p* does not automatically mean *S* is warranted in believing *p*. In other words, the rationality or credibility of *p* fundamentally depends solely on the *objective* relations between *S*'s reasons and *p*. The following quotations further illustrate this idea:

> The doxastic attitude that a person is justified in having is the one that fits the person's evidence. More precisely: EJ [Epistemic Justification]: Doxastic attitude *D* toward proposition *p* is epistemically justified for *S* at *t* iff having *D* toward *p* fits the evidence *S* has at *t*[6] (Feldman and Conee 1985, 15).

> [Epistemic Internalism]: An individual's having an epistemic warrant, and being epistemically warranted, in having a belief or in engaging in an epistemically relevant transition, supervenes on the non-factive kinds of psychological states, and kinds of relations among them, that are present in the individual's psychology (Burge 2020, 61).

These quotations emphasize that propositional warrant arises from the objective, or supervenient, relations rooted in an agent's internal psychological states. Notably, Conee and Feldman's quotation, drawing from Firth's (1978) work, underscores the alignment

[5]For a detailed overview, see Burge (2020).

[6]In this context, Feldman and Conee distinguish between two types of warrant. These authors call *propositional-type warrant* "justification" and *doxastic-type warrant* "well-founded belief."

https://doi.org/10.1017/epi.2025.10070 Published online by Cambridge University Press

between propositional warrant and the evidential support relations among propositions. Firth aptly captures this idea, stating:

> We might perhaps say, therefore, that this assessment of propositional warrant is a judgment about the evidential relationship between certain psychological states and proposition[s] [. . .] With appropriate qualifications we might want to call this a 'logical' relationship (Firth 1978, 218–9).

So, the reason-first approach focuses on determining when *S has a warrant* for a proposition rather than on when they are *warranted in believing it*. To be warranted in holding a belief, one must not only possess good reasons but also believe based on those reasons. Thus, propositional warrant remains independent of *S*'s cognitive limitations.

As a matter of fact, Firth adopts a broad interpretation of the term "logical," encompassing both deductive and non-deductive objective relationships. This suggests that, in the reason-first framework, propositional warrant depends on the objective connections between propositions and the reasons available to *S*, even if *S* is unaware of or unable to fully grasp them. Thus, on this view, the ultimate epistemic support for a belief lies in its logical relation to the agent's reasons, regardless of the agent's understanding (cf. Smithies 2015; Melis 2018; De Toffoli 2022). Kornblith further develops this idea, noting:

> Although the details here matter, and they vary significantly among different theorists, the big picture is quite straightforward: propositionally justified belief is explained in terms of the notion of a good argument, and goodness of argument is explained by the laws of logic, both deductive and non-deductive, together, perhaps, with a theory of probability (Kornblith 2022, 42).

Now, this focus on the objective structure of propositional warrant is particularly illuminating in the context of *KFF* cases, where *S*'s initial reasons may support both a false proposition and a true proxy premise.[7] For instance, consider the *Fancy watch* case: *S*'s initial reason – "My watch reads 2:58 p.m." – provides evidential support for both the false proposition "It is exactly 2:58 p.m." and the true proposition "It is approximately 2:58 p.m."[8] However, as responses to the proxy premise argument emphasize, *S* actually believes *p,* not *a*. This distinction is crucial: in the actual world, *S* has a doxastic warrant for *p* but not for *a* because *S* does not believe *a*. In the *Fancy watch* case, for instance, *S*'s doxastic warrant applies only to "It is exactly 2:58 p.m." because they base their belief on the specific reading of the watch, not the approximate time.

Therefore, *KFF* advocates address objections to the proxy premise argument by distinguishing between propositional and doxastic levels of warrant, a distinction that aligns with the reason-first approach. Specifically, they argue that while new reasons may defeat *S*'s propositional (and doxastic) warrant for "It is exactly 2:58 p.m.," the propositional warrant for "It is approximately 2:58 p.m." remains intact. Importantly, this propositional warrant, supported by *S*'s reasons evidentially and deductively, does not lead to a belief but provides a solid ground for a stable inference. As a result, *a* continues to *inertly* support the true conclusion, "I am not late for the 7 p.m. meeting."

[7]To be concise, we will refer to "the propositional warrant of the proposition expressed by sentence *X*" as "the propositional warrant of *p*."

[8]*S* has a propositional warrant for *p* on account of evidence *e*, as *e* strongly indicates that *p* is likely true. If *e* holds, it significantly increases the probability of *p*, giving *S* a compelling reason to accept *p* as true.

https://doi.org/10.1017/epi.2025.10070 Published online by Cambridge University Press

Thus, *S* can know this latter fact, even if the inference stems from a false belief about the exact time.

In contrast, while the reason-first approach can also be applied to explain Gettier cases, *KFF* proponents argue that their strategy does not fit these cases, further highlighting the plausibility of their rejoinder. In Gettier cases, the reasons available to *S* do not provide a propositional warrant for a proxy premise, making it difficult to infer conclusions like "Someone in my class owns a Ford," even if the conclusion is true and based on a false belief.

Therefore, if our considerations so far are accurate, it would be reasonable to conclude that the *KFF* proponents' solution to the proxy premise argument is at least *internalistically* adequate. Specifically, if one accepts a conception of warrant where doxastic warrant entails propositional warrant – with the latter being fundamental – they would agree with the *KFF* rejoinder. However, this raises the question: is this solution acceptable to externalists? The next subsection will explore the key differences between internalist and externalist approaches to *KFF*.

3.2. But what about externalism?

Traditionally, epistemic externalists, including Goldman (1979), Kornblith (1980), and more recent authors such as Graham and Lyons (2021), reject the reason-first approach. Consider, for example, the views of Goldman and Kornblith:

> [. . .] the bulk of this paper was addressed to [doxastic] justifiedness. This is the appropriate analysandum if one is interested in the connection between justifiedness and knowledge, since what is crucial to whether a person knows a proposition is whether he has an actual belief in the proposition that is justified (Goldman 1979, 21–2).

> The key to understanding the phenomenon of justified belief, as I see it, thus lies in taking the notion of doxastic justification to be the more fundamental notion (Kornblith 2022, 56).

Hence, for epistemic externalists, the relationship between propositional and doxastic warrant is understood in reverse compared to the reason-first view: doxastic warrant is conceptually fundamental, and propositional warrant is defined in terms of it. Here is Goldman:

> [S] is [propositionally] justified in believing p at t just in case a reliable belief-forming operation is available to him such that the application of that operation to his total cognitive state a t would result, more or less immediately, in his believing p and this belief would be [doxastically] justified (Goldman 1979, 23).

Within this framework, *S* has doxastic warrant for *p* if and only if the belief in *p* results from a reliable belief-forming process type.[9]

[9]For the purposes of this discussion, we will set aside the challenges to Goldman's theory raised by Plantinga (1993). As the following sections will show, the externalist account of warrant considered here includes a crucial requirement for competence (and proper functioning) (cf. Graham 2016; Burge 2020). Hence, by incorporating these requirements into the notion of epistemic warrant, we can effectively address Plantinga's counterexamples.

https://doi.org/10.1017/epi.2025.10070 Published online by Cambridge University Press

Now, the shift in focus from propositional to doxastic warrant carries significant implications for the externalist position. First, it shows that a warrant for a belief that p does not ensure a backup warrant for another proposition a. This is because the reliable processes or cognitive states needed to warrant an alternative belief may not be available to S at a given moment, thus failing to meet reliabilist standards. Second, it underscores that the ultimate grounds of rationality – aiming to believe the truth and avoid error – are now determined not by *the logical relations between internal states and propositions* (cf. Pollock and Cruz 1999), but by the agent's *contingent cognitive features* (cf. Graham 2016). Consequently, and more importantly, it appears that externalists may not benefit from the rejoinder proposed by *KFF* proponents against *CC* advocates.[10]

Consider *Fancy watch*: S checks their watch, feels confident in its accuracy, and, after careful observation, forms the belief, "It is exactly 2:58 p.m." Although this belief turns out to be false, it is still doxastically warranted. According to the reliabilist framework, epistemic warrant accommodates occasional errors (cf. Burge 2020). What matters is that the belief was formed through a generally reliable process, even if it led to a mistake in this particular instance.

However, a key question arises: at t_0, when S checks the watch and before forming the belief that p at t_1, does S have available a reliable process that could lead to an alternative epistemic path toward concluding q – specifically, one that concludes q by way of a propositional warrant for a proxy premise? According to Goldman's account, having a propositional warrant requires that, at time t, S must have a reliable process available that could provide a doxastic warrant for the relevant belief. Thus, the central issue is whether S has such a reliable process available at t_0 to form an approximate belief.

This issue introduces additional complexity. In general, belief-forming processes for approximate beliefs tend to be significantly less reliable than those that lead to exact beliefs (cf. Kornblinth 2017; Grundmann 2022). Several factors contribute to this difference.

First, belief-forming processes for exact beliefs typically rely on well-defined mechanisms and feedback loops that can be measured objectively, thereby ensuring the accuracy and truth of the belief (cf. Burge 2020). For instance, we can assess whether a person's belief about the time is correct by comparing their watch to an atomic clock. This kind of verification process provides clear and reliable feedback. By contrast, the formation of approximate beliefs often lacks such transparency, making their evaluation more difficult.[11]

Now, some might argue that approximate beliefs, being logically weaker, are more likely to be true in a broad or general sense. However, this supposed advantage is offset by the difficulty of assessing their accuracy. As a matter of fact, without clear methods of evaluation, their reliability becomes questionable (cf. Goldman 1979). For instance, consider someone who looks at an accurate watch and concludes, "The time is about 2:58 p.m." While this belief may be roughly correct for everyday purposes, its vagueness introduces a greater risk of error due to the inherent ambiguity of approximation.

This ambiguity becomes even more apparent when we try to assign a clear truth value to approximate beliefs. Unlike exact beliefs, which can be directly tested against objective standards, approximate beliefs often resist straightforward evaluation. For example, if an atomic clock shows 3:02 p.m., and someone believes "The time is approximately 2:58 p.m.," what truth value can we reasonably assign to that belief? The lack of precision

[10]Interestingly, Klein (2008, 29) acknowledges the distinction between reason-first and externalist approaches on these matters but does not further examine their potential implications for cases such as *KFF*.

[11]This does not mean that S must have a complete understanding of how their cognitive processes work. Rather, the call for transparency is directed at cognitive scientists, who are better positioned to competently evaluate the reliability of the relevant epistemic processes.

https://doi.org/10.1017/epi.2025.10070 Published online by Cambridge University Press

makes it difficult to provide a definitive answer, suggesting that approximation introduces a degree of interpretive openness that objective criteria alone may not resolve. In such cases, our judgments about truth may depend more on social conventions than on measurable facts.[12]

Second, belief-forming processes for approximate beliefs may be prone to systematic biases. Approximations can often underestimate or overestimate values, further reducing the reliability of these processes and weakening the epistemic support for our beliefs (cf. Tversky and Kahneman 1974). For example, if *S* tends to round up the time (such as rounding 2:57 to 2:58 p.m.), this rounding bias introduces minor inaccuracies that, over time, could lead to consistently flawed beliefs. In contrast, exact beliefs, like "It is exactly 2:58 p.m.," are less affected by such biases; *S*'s watch provides a specific time, and, unless it is consistently misread, there is no intermediary where rounding or similar biases could interfere.

Third, the cognitive complexity in forming approximate beliefs introduces a margin of error that is absent when forming exact beliefs. Epistemic externalism favors processes that minimize error and maximize truth. Thus, if an approximate process is more error-prone, it may not meet the reliabilist criteria for epistemic warrant (cf. Williamson 2000). For instance, if *S* tries to approximate the time while multitasking – glancing at the time while also focusing on the meeting – split attention could lead to errors. Estimations in such scenarios rely on less precise perceptual cues, resulting in less reliable beliefs. In contrast, forming the belief "It is exactly 2:58" by simply reading the watch enables an immediate and cognitively effortless belief formation with minimal mental interference. This simplicity reduces error and supports the reliabilist criteria for doxastic warrant.

Therefore, although approximate beliefs may be logically weaker, reliabilism does not grant them higher reliability simply because they are easier to align with the truth. The potential for error in approximation undermines their reliability by externalist standards, which require a high ratio of true to false beliefs, not just *close enough* beliefs. Given these challenges, it becomes clear that typical *a posteriori* processes for forming approximate beliefs are inherently unreliable. Their lack of transparency, susceptibility to biases, and cognitive complexity compromise their capacity to produce reliable beliefs consistently.

Furthermore, it is also noteworthy to highlight that even employing *a priori* processes would not resolve this issue. For instance, consider the claim that the proxy premise "It is approximately 2:58 p.m." could be warranted simply by understanding the statement "My watch reads 2:58 p.m." For this to be the case, there would need to be an inherent semantic link between the two, grounded in the meanings of terms like "approximately" and "reads" or in the sentences themselves. Yet such a link is not guaranteed. An epistemic agent might reliably understand "approximately" without understanding "reads" and vice versa. Moreover, a competent agent might reliably grasp the meanings of these sentences without recognizing any necessary semantic link between them (cf. Williamson 2007). Thus, since understanding one term or sentence does not inherently depend on the other, the warrant for this approximate belief cannot be derived solely from the meanings of the words or sentences.

Finally, before moving on, it is also important to clarify a potential misunderstanding regarding *S*'s ability to infer the proxy premise *a* from the false belief *p*. At time t_0, belief *p* has not yet formed in *S*'s cognition. This is a crucial point. Unlike internalists, who rely on perceptual or doxastic reasons that are logically connected to other propositions to explain *S*'s warranted states, *S* here operates within an externalist framework, where *S*'s

[12]Thanks to Federico Luzzi for a helpful discussion on this point.

https://doi.org/10.1017/epi.2025.10070 Published online by Cambridge University Press

only cognitive input is a perceptual (non-doxastic) state showing the exact time on their watch at t_0.

This clarification carries two key implications. First, unless this perceptual state engages in a reliable cognitive process available to *S* at t_0, it cannot independently provide epistemic support on its own. However, as discussed above, *S* lacks a reliable method to form *a* at t_0. Second, this latter point is reinforced by the fact that *S*, in the actual world, interprets the time displayed on the clock only in an exact way at t_0. Thus, from a strictly externalist standpoint, *S* also lacks any approximate cognitive states at t_0 that could lead to the formation of the proxy premise.[13]

In conclusion, any attempt to explain how *S* could derive a warrant for *a* from *p* (or their initial cognitive states) overlooks the limitations of *S*'s cognition at t_0. These limitations preclude the formation of the proxy premise and prevent *S* from acquiring propositional warrant for *a* at that time.[14]

Given these considerations, if externalists supporting *KFF*'s thesis cannot identify a reliable process available to *S* at t_0 to form an approximate belief based on initial cognitive states, we should conclude that *S* lacks propositional warrant for *a*. This suggests that externalists may not only miss the benefits of the proxy premise strategy proposed by reason-first advocates of *KFF* to counter *CC* but also struggle to distinguish *KFF* cases from Gettier cases. Notably, the *epistemic triangularization* proposed by the reason-first approach in *KFF* cases – linking initial reasons, doxastic warrant, and propositional warrant – appears incomplete within the externalist framework. Specifically, an epistemic link between the initial grounds and the approximate belief appears to be missing. Consequently, the lack of a stable path available to proponents of the reason-first approach further weakens the externalist position. Below is a graphical representation of these ideas (see Figure 2).

Despite these challenges to the externalist position, concluding that these considerations fully support the reason-first approach is premature. To understand this view's limitations in *KFF* cases, it is crucial to remember that the reason-first approach emphasizes the role of objective evidential relations between propositions and *S*'s available reasons in structuring (propositional) warrant. In the next subsection, we will present an argument against the reason-first view, highlighting the latent risks of this assumption.

3.3. Countering the standard explanation of KFF

To illustrate the limitations of the reason-first approach, let us consider the *Fancy watch* case, assuming the same initial conditions remain in place. To recap, in this scenario, *S* has a doxastic and propositional warrant for a false belief *p* and a propositional warrant

[13]This does not mean that *S* cannot perceptually focus on a more approximate time. Instead, our considerations suggest that, at t_0, *S* does not *intend* to form an approximate belief. This is likely because they prioritize punctuality and trust the accuracy of their watch. Therefore, this interpretation further clarifies why, despite the general unreliability of approximate belief-formation processes, *S* cannot "more or less immediately" form an approximate belief at t_0 in this context. Furthermore, as we will explore in the following sections, the transition to an approximate belief is more plausibly explained from an externalist perspective *via* an alternative pathway. Given that the transition from perceptual state to approximate belief is typically unreliable, it would be safer for the agent – especially one who highly values punctuality – to first form a belief about the exact time. From this exact belief, the agent could then more reliably infer the approximate time if needed.

[14]To remain consistent with the original case, we will assume that, at t_0, *S* has no access to testimonial sources for information about the relevant proxy premise.

https://doi.org/10.1017/epi.2025.10070 Published online by Cambridge University Press

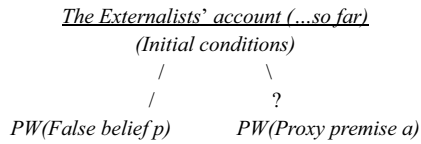


Figure 2. The graph illustrates the externalist account of *KFF* at t_0, where *S* has a *PW* for *p* but lacks one for *a* due to the absence of reliable processes available to *S* at t_0.

for a true proxy premise *a*. Notably, both warrants are supported by the same reason, "My watch reads 2:58 p.m."

Now, according to the reason-first approach, if *S* has a propositional warrant for *p*, *S* possesses a reason that, though false, can generate infinite propositional warrant tokens through logical links independently of *S*'s cognitive capacities. As a matter of fact, since propositional warrant is closed under logical consequence (assuming classical logic to adopt *KFF* proponents' assumptions), an infinite set of warranted propositions can be derived from an arbitrary belief ψ using *Disjunction Introduction*, i.e., $\psi \vdash \psi \vee \omega$.[15] For example, if *S* has a propositional warrant for ψ, *S* also has a propositional warrant for $\psi \vee \omega$, where ω represents any atomic or non-atomic proposition like "the wall is blue" or "if the wall is blue, then the wall is blue."

Interestingly, both friends and foes of *KFF* cases support this perspective. For instance, Klein (2008, 50–1), building on Hilpinnen's work, argues that for *KFF* cases to be genuine cases of inferential knowledge, the proxy premise *a* must be "sufficiently close" to the false premise *p*. While Klein suggests something less strict than an entailment relation might work to explain this epistemic closeness, his brief remarks on this point indicate he favors entailment. Thus, he seems to argue that *S* has propositional warrant for *a* because *a* is entailed by *p* (and because *S*'s initial reasons support it).

A similar view is expressed by Montminy, who defends the role of entailment in supporting *CC via* the proxy premise argument (cf. §2.1). In a recent work, however, Montminy (2023, 65) adds an important detail: while some true propositions may not be explicitly believed, an infinite set of "obvious" propositions remains dispositionally warranted and epistemically available to *S*. This highlights the wide scope of propositional warrant within the reason-first approach, showing that even a false premise can enable *S* to access a vast network of warranted propositions through entailment, such as $\psi \vee \omega$.

However, the reason-first approach faces challenges due to the arbitrary nature of the propositions that can be substituted into the disjunct ω. For example, consider the

[15]For the purposes of this discussion, we will proceed under the assumption that the traditional notion of propositional warrant is closed under (classical) entailment – a view that continues to receive considerable support in recent literature (cf. Ichikawa and Jarvis 2013; Smithies 2015; McCain and Moretti 2021). That said, it is also important to acknowledge that this position remains a matter of ongoing debate. Several philosophers, including Nozick and Dretske, have raised significant challenges to this perspective. More recently, scholars such as Turri (2010) and Boghossian (2014) have articulated further concerns. In considering such challenges, some authors – most notably Ichikawa and Jarvis (2013) and Coliva (2014) – have recently proposed distinguishing between the traditional notion of propositional warrant and what has been termed "ordinary propositional warrant" (cf. Melis 2018). On this account, ordinary propositional warrant represents a subset of the traditional notion that is aligned with the agent's cognitive and doxastic abilities. For a critical discussion of these proposals, see especially Melis (2018). Further efforts to clarify and refine the boundaries of the traditional concept of propositional warrant can be found also in the work of Pryor (2000) and Huemer (2001, 2007). For a critical examination of these and other internalist positions on epistemic warrant, see especially Burge (2020). Thanks to an anonymous referee for encouraging greater clarity on this issue.

https://doi.org/10.1017/epi.2025.10070 Published online by Cambridge University Press

following substitution: "All my approximate belief-forming processes are unreliable." If this instance of ω were derivable from the false premise p, it would present serious issues for the proxy premise.

Now, to see how this instance of ω could easily be derived, consider the relevant counterfactual situation where S, after forming the doxastically warranted belief "I am not late for my 7 p.m. meeting" based on the false belief p, learns that "It is not exactly 2:58 p.m." In this case, S gains a reason to reject the false belief, "It is exactly 2:58 p.m." However, since propositional warrant is closed under logical consequence, S would also acquire propositional warrant for any claim substituted into ω *via Disjunctive Syllogism*, i.e., $\psi \vee \omega$, $\neg\psi \vdash \omega$. Consequently, if ω represents the problematic proposition mentioned earlier, S would now have a reason that defeats the previously warranted proxy premise (cf. Pollock 1986).[16]

To illustrate this more clearly, consider ω as "All my approximate belief-forming processes are unreliable." In this case, $\psi \vee \omega$ can be read as "It is exactly 2:58 p.m., or all my approximate belief-forming processes are unreliable." Now, when we introduce the true information $\neg\psi$ ("It is not exactly 2:58 p.m."), we can apply *Disjunctive Syllogism* to derive ω. This leads to the result that S has propositional warrant for ω. However, this instance of ω challenges the credibility of the proxy premise's propositional warrant. Using Pollock's terminology, the derivation of ω introduces an *undercutting defeater* for S. Specifically, if R is a defeasible reason for S to believe p, R^* undercuts R just in case R^* shows S that their reasons for believing p are inadequate.[17] Thus, the derivation of ω undercuts the proxy premise's warrant, preventing S from confidently knowing, "I am not late for my 7 p.m. appointment" (cf. Klein 2008, 32).

A similar dynamic applies when considering *rebutting defeaters*. Again, following Pollock's framework, a rebutting defeater arises when if R is a defeasible reason for S to believe p, R^* is a rebutting defeater for R just in case R^* is a reason for S to believe $\neg p$. Let's break down this case to understand the issue. Previously, we saw how *Disjunctive Syllogism* worked in our case: given $\psi \vee \omega$ and $\neg\psi$, we were able to derive ω. Now, consider ω as a conditional statement where the antecedent is $\neg p$ and the consequent is any proposition that contradicts a. By applying *Modus Ponens* – i.e., ψ, $\psi \rightarrow \omega \vdash \omega$ – to $\neg p$ and the conditional statement that instantiates ω, it follows that S is propositionally warranted in believing a rebutting defeater for a upon learning $\neg p$. Therefore, deriving one's preferred ω undermines the propositional warrant of the proxy premise, preventing S again from confidently knowing, "I am not late for my 7 p.m. appointment."

But there is more! The problems with the reason-first approach become clearer when we consider Gettier cases. Suppose S has a propositional warrant for "Nogot owns a Ford." By applying *Disjunction Introduction*, S would also have a propositional warrant for "Nogot owns a Ford, or Havit owns a Ford." Yet, if we assume S later discovers that "Nogot does not own a Ford," we can use *Disjunctive Syllogism* again to derive "Havit owns a Ford" and thus argue that S acquires propositional warrant for this proxy premise.

[16]Here, we assume that a defeater, broadly speaking, is a reason that, when added to a subject's existing evidence for a target proposition, undermines the overall warrant needed for it to count as knowledge (cf. Pollock and Cruz 1999, 37). Further, we assume that defeaters are not constrained by factivity or the need for good evidence for S to believe $\neg p$, *pace* Klein (2008, 34) and de Almeida (2017, 308). For further discussion on this latter point, see Bernecker (2023).

[17]Graham and Lyons (2021) note that defeaters are rarely discussed in terms of reasons. Therefore, Pollock's definition could be refined to emphasize that warrant is the primary epistemic good being defeated.

https://doi.org/10.1017/epi.2025.10070 Published online by Cambridge University Press

However, this result contrasts with the conclusions typically drawn by the reason-first proponents in *KFF* cases. As a matter of fact, in the relevant scenarios where *S* discovers that "Nogot does not own a Ford," *S* might still retain a propositional warrant for a proxy premise supporting the knowledge that "Someone in the class owns a Ford."[18]

Therefore, these cases highlight a significant challenge for the reason-first framework: by imposing objective constraints on propositional warrant rather than cognitively fitting ones, this approach can become too permissive and yield counterintuitive results.

With this in mind, we will now turn to four potential objections that reason-first advocates might raise regarding *KFF* cases, and we will address each in turn.

3.3.1. Objection 1

A first objection targets our handling of the derivation of ω in the counterexamples discussed above. Specifically, it could be argued that when *S* learns $\neg\psi$, *S* loses propositional warrant for $\psi \vee \omega$, preventing the derivation of ω. More precisely, suppose $\psi \vee \omega$ can be derived from ψ, granting *S* propositional warrant for this disjunction. Now, if *S*'s warrant for $\psi \vee \omega$ is based *solely* on this derivation, then learning $\neg\psi$ would give *S* propositional warrant for $\neg\psi$. However, *S* cannot be warranted in both $\neg\psi$ and ψ. Hence, by a principle analogous to *CC* for epistemic warrant, i.e., *Warrant from Warrant* (*WFW*), *S* would lose propositional warrant for $\psi \vee \omega$. Here is *WFW*:

> *WFW* Necessarily, if *S*'s belief that *q* is competently inferred from a warranted premise *p* (without epistemic overdetermination), then if *S*'s belief that *q* is inferentially warranted, *p* is warranted (or contrapositively, if *S*'s belief that *p* is not warranted, then *q* is not warranted).

Despite this reply, several issues emerge, all rooted in *WFW*. The first, less significant point is that *WFW* has been questioned in the literature (cf. Murphy 2015). Although these critiques have not always gained widespread acceptance (cf. Luzzi 2019), it remains reasonable to conclude that *WFW* is at least a controversial principle within the realm of epistemic warrant.

Second, and more importantly, *WFW* pertains to doxastic, not propositional, warrant, as it addresses the epistemic warrant of beliefs formed through inferences that exclude overdetermination. Therefore, if *WFW* applies only to doxastic warrant, its relevance to our argument, which concerns propositional warrant, becomes unclear.

Finally, let us consider a plausible version of *WFW* for propositional warrant, which we will call *PWFPW*:

> *PWFPW* Necessarily, if *S can* competently infer *q* from a proposition p (without epistemic overdetermination), then if *S* has propositional warrant for *q* inferentially, then *p* is propositionally warranted for *S* (or contrapositively, if *p* is not propositionally warranted for *S*, then *q* is not propositionally warranted for *S*).

[18]For proponents of the reason-first approach, "Havit owns a Ford" can be independently warranted based on the reasons initially available to *S*. From "I see Nogot driving his Ford," any disjunction of the form $\psi \vee \omega$ can be derived. Therefore, when *S* acquires a defeater for *p*, *S* also gains a warrant for ω using *Disjunctive Syllogism*. Here, ω could be "Havit owns a Ford."

https://doi.org/10.1017/epi.2025.10070 Published online by Cambridge University Press

At first glance, *PWFPW* supports the case under consideration. As a matter of fact, it mirrors the belief-contraction mechanism raised in the earlier objection, now applied at the level of propositional warrant.

However, it is worth noting that, like the original *WFW*, *PWFPW* belongs to the broader class of counter-closure principles. While, to the best of our knowledge, *PWFPW* has not been directly criticized, it likely faces challenges similar to those of other counter-closure principles.

For instance, a key issue arises in *KFF* cases: *PWFPW* prevents *S* from recovering the proxy premise. To clarify, consider this: without the operation of belief-contraction imposed by *PWFPW*, losing the propositional warrant for a generic ψ would not affect the warrant for a disjunction like $\psi \vee \omega$ since $\psi \vee \omega$ logically follows from ψ. However, under *PWFPW's* belief-contraction mechanism, losing the warrant for ψ also results in losing the warrant for any disjunction $\psi \vee \omega$.

In the *KFF* cases, this leads to an unpalatable outcome: too much is sacrificed. As Klein (2008, 51) himself acknowledges, the proxy premise a is epistemically proximate to the false belief p because p always entails the true disjunction $p \vee a$. Hence, this proximity ensures that a remains epistemically available to *S* whenever they learn that $\neg p$. However, suppose *PWFPW* causes *S* to lose propositional warrant for $p \vee a$. In that case, *S* would also lose the propositional warrant for a.

Now, it could be argued that since *S*'s initial reasons independently support the proxy premise, its propositional warrant would not be entirely lost. However, it is crucial to recognize that the epistemic proximity between p and a is what ultimately salvages *KFF* cases. Without this proximity, *S* would struggle recovering a, as proponents of the reason-first approach emphasize. Therefore, if preserving closure under entailment for propositional warrant is essential for recovering the proxy premise through non-problematic cases of disjunction, then appealing to *PWFPW* fails to provide a satisfactory solution to the issues raised in the previous section concerning the problematic disjunctions of the form $\psi \vee \omega$.

In conclusion, while these considerations do not definitively disprove *PWFPW*, they strongly suggest that this principle falls short. An inference to the best explanation supports this conclusion: just as other counter-closure principles fail in other contexts, *PWFPW* might fare no better. Consequently, reason-first advocates of *KFF* must rely on something other than *PWFPW* to defend their account of *KFF*.[19,20]

3.3.2. Objection 2

A second objection to our argument might be framed as follows: some may contend that our *passe-partout* mechanism for deriving ω could provide *S* with a propositional warrant for every conceivable *defeater-eater* of chance (a defeater-eater being a defeater that reinstates the original warrant; cf. Klein 2008). In other words, by exploiting logical entailment, one could seemingly restore the warrant for the proxy premise by generating defeater-eaters that neutralize any proposed defeater, thus reinforcing the reason-first approach to *KFF* cases.

However, we believe that this objection actually highlights a deeper issue within the reason-first account of propositional warrant. Accepting this line of reasoning would push the reason-first framework toward an implausible idealization – one that assumes

[19]Moreover, note that even if independent reasons support a in standard fallback scenarios envisioned by reason-first proponents of *KFF*, *S* may still have indirect reasons (such as $\neg p$) to undermine a and no direct reasons to refute the rebutting defeater ($\neg a$).

[20]Thanks to Federico Luzzi and Lorenzo Rossi for their help in refining the point presented here.

https://doi.org/10.1017/epi.2025.10070 Published online by Cambridge University Press

epistemic agents possess an unlimited number of higher-order reasons for a vast array of claims, many of which exceed their cognitive capacities. However, this would not only undermine the practical basis of human knowledge but also distort what it means to possess warrant in genuine real-world epistemic contexts.

3.3.3. Objection 3

Another potential objection to our argument could arise from the claim that classical entailment may not adequately account for the closure of propositional warrant under logical consequence. While we recognize this as a legitimate concern, we also believe it places the burden on reason-first theorists to identify a more suitable logical framework. This is a significant challenge for those familiar with current debates in the epistemology of logic. As a matter of fact, anyone pursuing this kind of logical revision must propose an alternative framework that not only addresses the shortcomings of classical entailment in this context but also remains consistent with the fundamental commitments of the reason-first approach to propositional warrant.[21]

Alternatively, and perhaps more radically, one might argue that propositional warrant is not necessarily closed under logical consequence but rather under some weaker, non-deductive relation. However, the key question remains: which logical relation models propositional warrant in this case? Until such a framework is clearly articulated, the challenge of accurately modeling the objective evidential relations that underpin propositional warrant remains an open and pressing problem.

3.3.4. Objection 4

Finally, some may argue that our argument assumes that S learns that $\neg p$, but this assumption may not be necessary. After all, in standard *KFF* cases, S does not need to become aware of $\neg p$.

We acknowledge that this might be a reasonable objection. However, we also believe it overlooks our primary goal: to clarify the distinction between cases of inferential knowledge (*KFF*) and cases of inferential ignorance (*Gettier*), which we believe are one of the underlying reasons for denying *CC*. Hence, to bring this distinction into sharper focus, we find it both appropriate and natural to consider relevant counterfactual scenarios in which S learns that $\neg p$.

Moreover, we believe this latter objection can also be addressed from a different perspective. As Klein (2008, 44) himself points out, if $\psi \rightarrow \omega$ holds, then ψ is logically equivalent to ψ & ω. Now, if we consider a scenario where S has a propositional warrant for a conditional $\psi \rightarrow \omega$, derived *via* closure from p, and ω contains a true undercutting defeater, then S's knowledge of q is blocked, even without explicitly learning $\neg p$. For instance, suppose ω is, "My watch is accurate and fallible." By logical equivalence and *conjunction elimination*, i.e., ψ & $\omega \vdash \psi/\omega$, we can easily see that knowledge of q is blocked. In effect, although the watch is indeed accurate, this does not guarantee that it cannot be wrong at that particular moment.[22]

[21]If classical logic proves inadequate in this context, determining the most suitable non-classical framework becomes a significant challenge, particularly because the selection among competing logical systems is often underdetermined. In this context, intuitionistic, paraconsistent, relevant, and various substructural logics all present themselves as viable alternatives. However, the primary difficulty may not lie so much in deciding which logic to adopt but rather in developing clear criteria and robust methods for making such a selection. This issue underscores the complexity of the matter and reflects the ongoing, nuanced debates within the epistemology of logic (cf. Baggio 2025).

[22]Thanks to Federico Luzzi for a helpful discussion on this point.

https://doi.org/10.1017/epi.2025.10070 Published online by Cambridge University Press

~

Given these considerations, the remainder of this article will focus on reassessing the externalist view. Specifically, we will examine a new externalist framework that inverts the traditional relationship between propositional and doxastic warrant. Through this examination, we will argue that *KFF* cases constitute genuine instances of inferential knowledge and that they continue to pose significant challenges to *CC*.

4. *KFF* through the lens of epistemic externalism

This section explores an externalist account of propositional warrant and defeaters proposed by Graham and Lyons (2021), which presents a general externalist perspective on the relationship between these two key epistemic concepts. Due to its flexibility in framing externalist views, this account will play a crucial role in clarifying why it is the most effective framework for understanding the plausibility of *KFF* cases in challenging *CC*.[23]

4.1. Graham & Lyons's account

Graham and Lyons have recently put forward a general externalist account of propositional warrant and defeaters. To engage meaningfully with their view, it is important to first clarify a key aspect of their terminology. Specifically, they use the term "warrant" in a way that encompasses both propositional warrant and defeaters, effectively treating the possession of a defeater as equivalent to having propositional warrant. In contrast, throughout this paper, we have employed the term "warrant" in a more precise sense, meaning non-factive, positive epistemic support for a proposition or belief. Hence, to accurately reflect Graham and Lyons's usage while avoiding confusion, we will use the term "warrant/defeater" when referring to their conception of warrant as inclusive of defeaters. This terminological distinction is essential for a clear and accurate presentation of their account. Thus, with these clarifications in mind, we can now turn to Graham and Lyons's account:

> *S* has (*prima facie*) warrant for believing *p* at *t* iff a cognitive process that satisfies the general theoretical requirements for *prima facie* (doxastic) justification (a) is available to *S*, and (b) if used at *t*, taking as inputs only states that *S* is already in, does or would likely produce *p* as output (Graham and Lyons 2021, 61).

Graham and Lyons's account rests on two central ideas: (i) expanding Goldman's account of defeaters and propositional warrant and (ii) preserving Pollock's insights on defeaters in line with externalist tenets (cf. Graham and Lyons 2021, 57). In the remainder of this subsection, we will explore these two ideas in some detail, laying the groundwork for our subsequent analysis of *KFF* cases.

Regarding (i), Graham and Lyons's account closely parallels Goldman's account of propositional warrant. To recap, Goldman asserts that *S* has a propositional warrant for

[23]We chose to focus on the Graham and Lyons account for several reasons. First, it offers a broad, flexible perspective on externalism (more on this below). Second, other accounts, like Beddor (2015), align closely with this view, as noted by the authors we discuss. Third, while Melis's (2018) and De Toffoli's (2022) accounts are not strictly reason-first, they still incorporate internalist elements that can be easily reinterpreted in externalist terms. Given these considerations, and in light of space constraints, concentrating on this more general account allows us to develop the paper's dialectical points more effectively.

https://doi.org/10.1017/epi.2025.10070 Published online by Cambridge University Press

p at *t* just in case *S* has available a reliable process that would yield a doxastic warrant for *p* if applied to *S*'s cognitive states at *t*. Now, when Goldman's notion of defeaters is considered alongside his account of propositional warrant, Graham and Lyons's connection to it becomes clear. In Goldman's (1979) account, *S*'s warranted belief that *p* at *t* is defeated just in case *S* has another reliable process available at *t* that, if used, would have prevented *S* from believing *p*. Thus, for Goldman, Graham, and Lyons alike, a defeater for *p* is closely tied to having propositional warrant: if *S* were to apply an alternative reliable belief-forming process to their current cognitive states, *S* would arrive at a different judgment regarding *p* – such as disbelief, suspension of judgment, or a reduction in the degree of belief in the proposition.[24]

Despite having clear similarities to Goldman's view, Graham and Lyons introduce important adjustments in their account of warrant/defeaters. First, they argue that warrants/defeaters function as reasons in an *internalist* sense, but without requiring the agent to re-evaluate their beliefs consciously. That is, warrant/defeaters do not necessitate agent-level epistemic basing (cf. Graham and Lyons 2021, 58). Notably, their choice to characterize warrant/defeaters as "outputs" rather than beliefs reinforces this idea, suggesting that even non-doxastic states can act as warrant/defeaters for doxastically warranted beliefs.

For example, if *S* believes in *p* due to a reliable process but then perceives ¬*p* through another reliable process, that non-doxastic state acts as a warrant/defeater for *S*'s belief in *p*. Importantly, the existence of ¬*p* does not compel *S* to reassess their beliefs consciously; instead, *S* may immediately adopt the belief in ¬*p* or suspend judgment regarding *p*. Conversely, if *S* learns *p* but is misled by an illusion, such as the Müller-Lyer illusion, they may still maintain their belief in *p* despite having a non-doxastic defeater (cf. Graham and Lyons 2021, 62).

Second, Graham and Lyons's account deliberately leaves open the question of whether the outputs of a belief-forming process are always available to *S* (cf. Graham and Lyons 2021, 62). As a matter of fact, their account suggests that the available process "does or would likely produce *p* as an output," which implies that propositional warrant and defeaters may not always be immediately available to *S*. This indicates that Graham and Lyons view the availability of propositional warrants or defeaters as dependent on both the *cognitive complexity* involved in acquiring these epistemic resources and the *individual's cognitive capacities*.

For example, if forming a warrant/defeater for *p* (¬*p*) requires minimal cognitive effort – such as a few basic steps of reasoning – then ¬*p* could serve as a warrant/defeater for *p*, provided that *S* has the cognitive ability to form ¬*p* in their current state. However, if *S* has another available cognitive process that blocks ¬*p*, they may retain their belief in *p*. This aligns with Graham and Lyons's view, where *S* could use a *defeater-eater* to neutralize the original warrant/defeater. Conversely, if the process behind ¬*p* is reliable but beyond *S*'s cognitive capacities, ¬*p* would not function as a warrant/defeater for *p* (cf. Graham and Lyons 2021, 63).

Thus, while there is indeed some ambiguity in this account regarding how reliable cognitive abilities are identified or how modally stable they must be to ensure the availability of their outputs to *S*, this vagueness ultimately supports Graham and Lyons's account. As they note, their framework offers a flexible guide for externalists rather than a rigid blueprint for understanding propositional warrants and defeaters. Consequently, while the precise details of how reliable competencies are acquired can be filled in by different externalist accounts to address concerns about the scope and applicability of

[24]Graham and Lyons (2021, 60) argue that Pollock's traditional view is too rigid compared to Goldman's, who frames defeaters as leading to *non-belief* rather than a drastic shift from *p* to ¬*p*.

https://doi.org/10.1017/epi.2025.10070 Published online by Cambridge University Press

these reliable epistemic processes in various contexts, we can set aside such specifics for now and focus on the broader contours of this general approach.

Building on these considerations, Graham and Lyons propose the following analysis of warrant/defeaters:

- *S* has a rebutting defeater for her belief that *p* iff she has a warrant to believe ¬*p*.
- *S* has an undercutting defeater for her belief that *p* iff she has a warrant to believe that her warrants for believing *p* are inadequate.
- *S* has a defeater for her belief that *p* iff she has a rebutting or an undercutting defeater for her belief that *p* (Graham and Lyons 2021, 59).

So, regarding (ii) above, Graham and Lyons's analysis also aligns with Pollock's reason-first view of defeaters (cf. §3.3). By integrating the notion of warrant/defeater into Pollock's framework, Graham and Lyons offer a theory that bridges internalist and externalist perspectives, with defeaters becoming more dependent on one's cognitive capacities.

Given this overall framework, the critical question we must now address is whether this account enables proponents of *KFF* to counter *CC* effectively. In the next subsection, we will apply Graham and Lyons's account to resolve the tension encountered in the standard explanation of *KFF* cases.

4.2. Toward an(other) explanation of KFF cases

If we recall correctly, the main issue with the reason-first interpretation of *KFF* cases is that the requirement for propositional warrant to be closed under logical consequence needs to be revised. Specifically, reasons derived from the false belief *p* can, through simple logical inferences, generate a defeater that blocks the credibility of the proxy premise *a*. In contrast, the externalist faces a more troubling issue: the lack of propositional warrant for *a* at t_0. Thus, in order to coherently reconsider this latter perspective, the first question we must address is whether *S* has a reliable backup warrant for the proxy premise at subsequent times t_1 through t_n that enables them to sustain their inferential knowledge of *q* even when confronted with defeaters.

According to accounts from Graham, Lyons, and Goldman, it appears that *S* can possess a propositional warrant for a proxy premise at t_1. To illustrate, consider *Fancy watch* and assume that *S* has doxastic warrant for *p* at t_1. In this case, if we can reasonably assume that *S* can make inferences, *S* could form the belief that *a* at t_1 through a belief-dependent, conditionally reliable process (cf. Goldman 1979; see also Grundmann 2022). This process involves two key elements: (i) the slightly inaccurate but doxastically warranted premise *p*, which, although false, remains competently produced, and (ii) the conditional true belief *a*: "If it is not exactly 2:58 p.m., then it is approximately 2:58 p.m." available to *S*.

Now, to see how *S* could arrive at this latter belief, we can consider two approaches. The first, based on a Kantian-inspired concept of *containment*, suggests that *S* could infer that the exact time is contained within the approximate time. Alternatively, a more sophisticated reasoner might use the formal concept of a *set*, concluding that the set labeled "Approximately 2:58 p.m." includes the subset of "Exactly 2:58 p.m."[25]

[25]The reader should keep in mind that the approximate beliefs that *S* could form at t_n cannot be derived from *S*'s perceptual state at t_0 using these reasoning processes. This limitation arises from the non-propositional nature of perceptual states, which can cause errors when interpreting initial sensory data.

https://doi.org/10.1017/epi.2025.10070 Published online by Cambridge University Press

Consequently, since *S* is confident that their watch is accurate, they can reliably rely on the false belief that *p* to derive *a*. Notably, given *S*'s normal reasoning abilities and reliable methods, *S* has a propositional warrant for *a* at time t_1. Even if *p* is slightly inaccurate, *a* remains warranted because *S* can engage in a conditionally reliable process. The negligible inaccuracy in *p* does not compromise this reasoning, as *p* is sufficiently dependable to assure the reliability of *a*. Thus, if the conditional statement holds reliably and the reasoning process is executed competently, *a* can be propositionally warranted for *S* at t_1.

This conclusion is reinforced by recognizing that when *S* forms the false belief *p*, they already have propositional warrant for the conclusion *q*. For example, the inference from "It is exactly 2:58 p.m." to "I am not late for my 7 p.m. appointment" is straightforward and reliable. The inferential process (likely simple arithmetic or time reasoning) is highly reliable, as there is minimal room for error in calculating that over four hours remain before the appointment. Although the input belief "It is exactly 2:58 p.m." is off by one minute, this inaccuracy is negligible because the difference between 2:57 p.m. and 7 p.m. is so large. Thus, the reliability of the conclusion, "I am not late for my 7 p.m. appointment," remains unaffected by the small error in the input belief *p*.

So, the key difference between the externalist and reason-first approaches lies in how they structure epistemic support in *KFF* cases. According to the externalist perspective, at t_0, *S* begins with a perceptual state. At t_1, this state is processed by a reliable belief-forming process that provides doxastic warrant for *p*. Simultaneously, another reliable process secures propositional warrant for *a*.

Conversely, the reason-first framework holds that after *S* acquires initial reasons at t_0, they obtain doxastic warrant for *p* but lack it for *a* at t_1. However, *S* retains propositional warrant for both *p* and *a* at t_0. Below is a graphical representation to illustrate these ideas (see Figure 3).

The graph on the right highlights two key aspects of the externalist approach to countering *CC*. First, it shows that the epistemic status of the inferential belief that *q* depends inherently on the false belief that *p*. In contrast, the left graph shows that, under the reason-first approach, the belief that *q* is not inherently dependent on the belief that *p*. Thus, while the reason-first approach can be criticized for relying on an alternative path that may be epistemically available but cognitively inaccessible to *S* in supporting knowledge of *q*, this criticism does not apply to the externalist approach.

Second, and relatedly, the graph on the right highlights the essential role of doxastic warrant in acquiring knowledge from falsehoods within the externalist framework. Recall that, for externalists, propositional warrant requires satisfying conditions for doxastic warrant, which, in turn, requires holding a belief. This then implies that even when inferential knowledge is obtained through an alternative stable path, the propositional warrant for elements along that path depends on *S*'s doxastic abilities. Thus, externalism consistently argues that the epistemic availability of any alternative path is linked to *S*'s cognitive capacities, underscoring the importance of doxastic warrant in acquiring and maintaining knowledge.[26]

Now, a potential objection to *KFF*'s core intuition – that these cases involve genuine inferential knowledge – is the absence of *epistemic triangulation* within the externalist framework. Some might argue that this absence weakens the idea that such cases count as knowledge from falsehoods. However, we contend that this lack of triangulation

However, once *S* acquires *p* at time t_1, *S* can generalize more reliably from a clearly defined propositional attitude, ultimately leading to a more stable belief that *a*.

[26]For a related discussion of these *desiderata*, see Bernecker (2023) and Pritchard (2023).

https://doi.org/10.1017/epi.2025.10070 Published online by Cambridge University Press

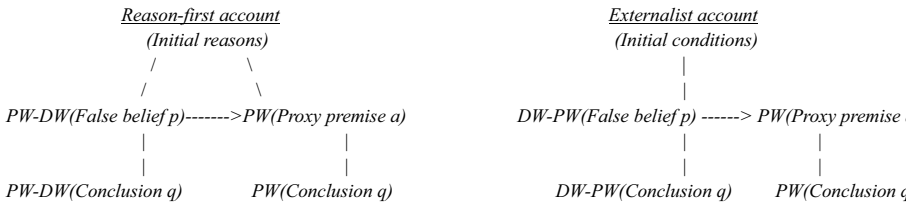


Figure 3. These graphs illustrate the epistemic dependency relations between the key epistemic elements of the internalist approach (on the left) and the externalist approach (on the right).

actually strengthens the view of knowledge as an achievement rooted in *S*'s cognitive abilities.

For example, we previously assumed that *S* could infer an approximate time from an exact one. However, this ability is not universally shared. A developing epistemic agent, like an infant, may lack the competence to make such inferences at certain stages of their development. For example, the child might not yet grasp concepts such as *set membership* or *containment*. In such cases, it seems intuitive to deny that the agent knows whether they are late, since they cannot reliably interpret the information provided by their belief about the watch. Without the necessary conceptual understanding, these agents clearly lack the epistemic competence required for such inferences. Therefore, by adopting a broad conception of relevant epistemic agents at play, the lack of triangulation in externalist accounts does not weaken the plausibility of *KFF* cases but allows for a more nuanced, context-sensitive understanding of inferential knowledge.

Having established the epistemic grounds of the inferential knowledge of *q*, we can now analyze defeaters. For an externalist, two key aspects of defeaters are crucial: the cognitive complexity involved in obtaining them and the nature of their semantic content. These factors are central to understanding how defeaters affect *KFF* cases within the externalist framework. Hence, we will begin by discussing the cognitive complexity component of defeaters.

In examining Graham and Lyons's account, we observed that the mere fact that a warrant/defeater logically follows from a proposition believed by *S* does not automatically make it epistemically available to *S*. For a warrant/defeater to be available, *S* must be able to infer it through a reliable cognitive process using the cognitive states at their disposal. Therefore, when considering the cognitive complexity of defeaters, the key issue is whether *S* can cognitively engage with the reasoning that leads to the undercutting and rebutting defeaters discussed in §3.3.

Unfortunately, the *KFF* cases we've analyzed in this paper provide little information about *S*'s cognitive background. Nevertheless, we can make three relevant observations about the processes involved in generating the defeaters discussed in §3.3.

First, both undercutting and rebutting defeaters in §3.3 depend on the effective application of deductive principles. Although these principles are generally straightforward and reliable, integrating them into a coherent argument presents a significant challenge for many epistemic agents. For instance, chaining *Disjunction Introduction* and *Disjunctive Syllogism* may stretch the cognitive capacities of some unskilled agents, hindering their ability to form the relevant defeaters.

Second, forming these defeaters requires the epistemic agent to engage with unfamiliar propositions that may not be immediately available to *S*. This suggests that the availability of these defeaters depends on *S*'s capacity to work with abstract or unfamiliar premises, a requirement that may not always be met.

https://doi.org/10.1017/epi.2025.10070 Published online by Cambridge University Press

Finally, the interplay of deductive reasoning with arbitrary premises complicates the process further. Such cognitive tasks may be beyond *S*'s reach and risk leading to unreliable outcomes. Hence, this complexity highlights the challenges *S* faces in generating and utilizing defeaters, illustrating the tension between reliable cognitive capacities and available epistemic states.

Based on these considerations, we can draw two general conclusions about the cognitive complexity involved in generating the defeaters discussed in §3.3. First, producing these defeaters requires greater cognitive effort than other inferences in *KFF* cases, as it involves combining at least two deductive principles. In contrast, forming the proxy premise *a* or inferring *q* typically involves only one inference.

Second, although *S* may be able to apply deductive processes reliably, the presence of arbitrary premises complicates the process of defeater formation within an externalist framework. This highlights the importance of *S* having available cognitive states that relate to their epistemic abilities and connects back to our earlier discussion regarding the semantic content of defeaters.[27] Therefore, let us now turn to the issues surrounding the semantic content of defeaters.

In an externalist framework, *S* would likely struggle to acquire the defeaters discussed in §3.3., as doing so would demand a level of conceptual sophistication that is typically beyond the cognitive reach of a non-expert agent. Specifically, *S* would need to trace the inferential chain to falsify *a* and understand the conceptual content of each step in that chain. This level of understanding is crucial. For example, grasping undercutting or rebutting defeaters requires advanced meta-epistemological skills, including a good understanding of abstract terms such as "reliability," "defeater," and "approximate beliefs."

Given this, while the *KFF* scenarios offer limited information and do not rule out the possibility that *S* may possess these skills, it seems safe to expect that an unskilled agent would lack the cognitive resources necessary to defeat the warrant for *a*. Additionally, even without considering these details, we have already observed that combining arbitrary premises with deductive reasoning typically leads to unreliable results. Therefore, unless the *KFF* cases are further clarified, it seems reasonable to conclude that these defeaters would not be epistemically available to *S*.

In conclusion, the abovementioned externalist analysis suggests that *S* can follow an undefeated, stable path to *q*, aligned with their doxastic abilities. Hence, this externalist account better predicts that *S* has the inferential skills to support the *KFF* advocates' thesis within a reasonable epistemic framework. Specifically, if *S* can derive the proxy premise *a* from *p*, we can clarify how *S* might know, "I am not late for my 7 p.m. appointment." By applying similar reasoning to that used for *a*, *S* can combine their belief in *a* with relevant propositions, such as "if it is approximately 2:58 p.m., then I am

[27]At this stage, an epistemic internalist might contend that the cognitive complexity component we have examined thus far simply reinforces the internalist's view that access or reflective awareness is essential for a mental state to count as a reason or defeater. However, we believe that the externalist notion of epistemic availability – grounded in reliable belief-forming processes that an agent can competently employ regardless of conscious access – offers a cognitively more plausible alternative. As a matter of fact, human beings often form beliefs through unconscious, automatic mechanisms that are both reliable and largely inaccessible to introspective awareness. Insisting on conscious access in such cases would not only exclude many ordinary beliefs from being warranted but would also place unreasonably high cognitive demands on agents when it comes to identifying defeaters. From this perspective, externalism seems better aligned with the realities of human cognition, securing epistemic warrant and defeaters through the availability of reliable processes and cognitive states, rather than relying on potentially fallible self-assessment. For a detailed discussion of these issues, see especially Bergman (2006). Many thanks to an anonymous referee for highlighting the importance of this point.

https://doi.org/10.1017/epi.2025.10070 Published online by Cambridge University Press

not late for my 7 p.m. appointment." This allows S to establish an undefeated propositional warrant for q, even if $\neg p$ is later discovered.

Crucially, this alternative inferential process remains reliable in counterfactual and real-world scenarios. As shown in the comparison between the externalist and internalist graphs, if S discovers that p is false, they can still infer both a and q. This consistency satisfies the externalist's criteria for propositional warrant, preserving S's epistemic position regarding q and meeting the conditions for doxastic warrant and knowledge in the actual world.

Lastly, the externalist framework also effectively addresses Gettier cases. In these circumstances, S lacks evidence regarding Havitt's car make, and their only method of forming a warranted belief about a is to guess which peer owns a Ford. Therefore, since S lacks reliable means to establish a stable alternative path, they cannot recover their knowledge of q.

5. Conclusions

Let's take stock. We have argued that the reason-first approach – typically favored by advocates of *KFF* cases – struggles to reliably distinguish between genuine and problematic instances of inferential knowledge. To address this challenge, we proposed adopting an externalist framework that provides a more refined account of propositional warrant and defeaters. Although this framework still requires further development, it effectively captures many of the key features present in *KFF* cases. Therefore, by reframing propositional warrant and defeaters through an externalist lens, we can deepen our understanding of how inferential knowledge can be secured, even when it is grounded in premises that are, unbeknownst to the agent, false.

Acknowledgements. I am deeply grateful to Silvia De Toffoli, Federico Luzzi, and Lorenzo Rossi for the many fruitful discussions we have shared on these topics. I also sincerely thank the audiences of the *Science and More* seminars at the University of Turin and the *Forschungsseminar* at the University of Innsbruck for their feedback and valuable contributions. Finally, I am indebted to an anonymous referee of this journal for their comments and constructive suggestions.

Funding. Funded by the European Union - Next Generation EU, Mission 4, Component 1, CUP I53D23006890001.

Conflict of interest. The author declares no conflict of interest.

References

Arnold A. (2013). 'Some Evidence is False.' *Australasian Journal of Philosophy* **91**(1), 165–72.

Baggio M. (2025). 'Logical Epistemology, Social Evidence, and the A Priori.' *Synthese* **205**, 1–28.

Ball B. and Blome-Tillmann M. (2014). 'Counter Closure and Knowledge despite Falsehood.' *Philosophical Quarterly* **64**(257), 552–68.

Beddor B. (2015). 'Process reliabilism's troubles with defeat.' *The Philosophical Quarterly* **65**, 145–59.

Bergman M. (2006). *Justification Without Awareness*. Oxford: Oxford University Press.

Bernecker S. (2023). 'Knowledge, Falsehood, and Defeat.' In R. Borges and I. Schnee (eds), *Illuminating Errors*, pp. 140–57. London: Routledge.

Boghossian P.A. (2014). 'What Is Inference?' *Philosophical Studies* **169**(1), 1–18.

Borges R. and Schnee I. (2023). *Illuminating Errors*. London: Routledge.

Burge T. (2020). 'Entitlement: The Basis for Empirical Epistemic Warrant.' In P. Graham and N.J.L.L. Pedersen (eds), *Epistemic Entitlement*, pp. 37–142. Oxford: Oxford University Press.

Coliva A. (2014). 'Moderatism, Transmission Failures, Closure, and Humean Scepticism.' In D. Dodd and E. Zardini (eds), *Scepticism and Perceptual Justification*, pp. 248–71. Oxford: Oxford University Press.

https://doi.org/10.1017/epi.2025.10070 Published online by Cambridge University Press

de Almeida C. (2017). 'Knowledge, Benign Falsehoods, and the Gettier Problem.' In R. Borges, C. de Almeida and P.D. Klein (eds), *Explaining Knowledge: New Essays on the Gettier Problem*, pp. 292–311. Oxford: Oxford University Press.
De Toffoli S. (2022). 'Intersubjective Propositional Justification.' In P. Silva and L.R.G. Oliveira (eds), *Propositional and Doxastic Justification*, pp. 241–62. London: Routledge.
Feit N. and Cullison A. (2011). 'When Does Falsehood Preclude Knowledge?' *Pacific Philosophical Quarterly* **92**(3), 283–304.
Feldman R. and Conee E. (1985). 'Evidentialism.' *Philosophical Studies* **48**(1), 15–34.
Firth R. (1978). 'Are Epistemic Concepts Reducible to Ethical Concepts?' In A.I. Goldman and I. Kim (eds), *Values and Morals*, pp. 215–29. Dordrecht: D. Reidel.
Fitelson B. (2010). 'Strengthening the Case for Knowledge from Falsehood.' *Analysis* **70**(4), 666–69.
Goldman A. (1979). 'What Is Justified Belief?' In G. Pappas (ed), *Justification and Knowledge*, pp. 1–25. Dordrecht: D. Reidel.
Graham P. (2016). 'Against Actual-World Reliabilism: Epistemically Correct Procedures and Reliably True Outcomes.' In M.A. Fernández Vargas (ed), *Performance Epistemology: Foundations and Applications*, pp. 83–105. Oxford: Oxford University Press.
Graham P. and Lyons J. (2021). 'The Structure of Defeat: Pollock's Evidentialism, Lackey's Framework, and Prospect for Reliabilism.' In J. Brown and M. Simion (eds), *Reasons, Justification, and Defeat*, pp. 39–68. Oxford: Oxford University Press.
Grundmann T. (2022). 'Dependent Reliability: Why and How Conditional Reliability Should Be Replaced by It.' *Philosophy and Phenomenological Research* **105**(1), 144–59.
Hilpinen R. (1988). 'Knowledge and Conditionals.' *Philosophical Perspectives* **2**, 157–82.
Huemer M. (2001). *Skepticism and the Veil of Perception*. Lanham, MD: Rowman & Littlefield.
Huemer M. (2007). 'Compassionate Phenomenal Conservatism.' *Philosophy and Phenomenological Research* **74**(1), 30–55.
Ichikawa J. and Jarvis B. (2013). *The Rules of Thought*. Oxford: Oxford University Press.
Klein P. (2008). 'Useful False Beliefs.' In Q. Smith (ed), *Epistemology: New Essays*, pp. 25–63. Oxford: Oxford University Press.
Kornblinth H. (2017). 'Doxastic Justification is Fundamental.' *Philosophical Topics* **45**(1), 63–80.
Kornblith H. (1980). 'Beyond Foundationalism and the Coherence Theory.' *Journal of Philosophy* **77**(10), 597–612.
Kornblith H. (2022). 'What Does Logic Have to Do with Justified Belief? Why Doxastic Justification Is Fundamental.' In P. Silva and L.R.G. Oliveira (eds), *Propositional and Doxastic Justification*, pp. 40–58. London: Routledge.
Lehrer K. (1965). 'Knowledge, Truth and Evidence.' *Analysis* **25**, 168–75.
Luzzi F. (2010). 'Counter-Closure.' *Australasian Journal of Philosophy* **88**(4), 673–83.
Luzzi F. (2019). *Knowledge from Non-Knowledge*. Cambridge: Cambridge University Press.
McCain K. and Moretti L. (2021). *Appearance and Explanation: Phenomenal Explanationism in Epistemology*. Oxford: Oxford University Press.
Melis G. (2018). 'The Intertwinement of Propositional and Doxastic Justification.' *Australasian Journal of Philosophy* **96**(2), 367–79.
Montminy M. (2015). 'Knowledge despite Falsehood.' *Canadian Journal of Philosophy* **44**(3–4), 463–75.
Montminy M. (2023). 'Harmless Falsehoods.' In R. Borges and I. Schnee (eds), *Illuminating Errors*, pp. 59–75. London: Routledge.
Murphy P. (2013). 'Another Blow to Knowledge from Knowledge.' *Logos and Episteme* **4**(3), 311–17.
Murphy P. (2015). 'Justified Belief from Unjustified Belief.' *Pacific Philosophical Quarterly* **98**, 602–17.
Nozick R. (1981). *Philosophical Explanations*. Oxford: Oxford University Press.
Plantinga A. (1993). *Warrant: The Current Debate*. Oxford: Oxford University Press.
Pollock J. (1986). *Contemporary Theories of Knowledge*. Lanham, MD: Rowman and Littlefield.
Pollock J. and Cruz J. (1999). *Contemporary Theories of Knowledge*. Lanham, MD: Rowman & Littlefield.
Pritchard D. (2005). *Epistemic Luck*. Oxford: Oxford University Press.
Pritchard D. (2023). 'Knowledge from Error and Anti-Risk Virtue Epistemology.' In R. Borges and I. Schnee (eds), *Illuminating Errors*, pp. 93–103. London: Routledge.
Pryor J. (2000). 'The Skeptic and the Dogmatist.' *Noûs* **34**(4), 517–49.
Saunders J. and Champawat N. (1964). 'Mr. Clark's Definition of 'Knowledge'.' *Analysis* **25**(1), 8–9.

https://doi.org/10.1017/epi.2025.10070 Published online by Cambridge University Press

Silva P. and Oliveira L. (2023). 'Propositional Justification and Doxastic Justification.' In M. Lasonen-Aarnio and C. Littlejohn (eds), *The Routledge Handbook of the Philosophy of Evidence*, pp. 395–407. London: Routledge.

Smithies D. (2015). 'Ideal Rationality and Logical Omniscience.' *Synthese* **192**(9), 2769–93.

Turri J. (2010). 'On the Relationship between Propositional and Doxastic Justification.' *Philosophy and Phenomenological Research* **80**(2), 312–26.

Tversky A. and Kahneman D. (1974). 'Judgment under Uncertainty: Heuristics and Biases.' *Science* **185**(4157), 1124–31.

Warfield T. (2005). 'Knowledge from Falsehood.' *Philosophical Perspectives* **19**(1), 405–16.

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

Williamson T. (2007). *The Philosophy of Philosophy*. Oxford: Blackwell.

Matteo Baggio is a postdoctoral research fellow at the University of Turin. He earned his PhD from the University School for Advanced Studies IUSS Pavia in Italy. Throughout his academic journey, he has conducted research at several international institutions, including the University of Bergen (2022), the Complutense University of Madrid (2022), and the University of Innsbruck (2025). His research primarily focuses on general epistemology, social epistemology, and the philosophy of logic.

Cite this article: Baggio M. (2025). "Knowledge from Falsehoods Reconsidered." *Episteme* 1–26. https://doi.org/10.1017/epi.2025.10070

https://doi.org/10.1017/epi.2025.10070 Published online by Cambridge University Press