

The Ignorance Dilemma and Awareness-First Epistemology

Paul Silva Jr. | *University of Cologne*

Draft, comments welcome | 25 March 2025

Abstract. There are cases in which an agent neither knows that p nor is ignorant of the fact that p . Every theory of knowledge, T , faces a dilemma in light of such cases: either T is too strong to explain the absence of factual ignorance in such cases, or T is too weak to explain the absence of knowledge in such cases. The solution is to embrace the first horn of the dilemma and to augment one’s theory of knowledge with an account of factual awareness that can explain why factual ignorance is absent in these cases. This paper develops a new, ignorance-based argument for the idea that factual awareness is a more general state of which knowledge is but one instance. It also provides further reasons to identify factual ignorance with the absence of factual awareness.

‘If I had to summarize [*Knowledge and Its Limits*] in two words, they would be: knowledge first. It takes the simple distinction between *knowledge and ignorance* as a starting point from which to explain other things, not as something itself to be explained. In that sense the book reverses the direction of explanation predominant in the history of epistemology.’

–Timothy Williamson (2000: v)
emphasis added

‘Knowledge can be uniquely characterized as the most general factive mental state; on this view, to know is to be in the *type* of mental state whose distinctive character is that it can only be held towards truths.’

–Jennifer Nagel (2023: 205)

Table of Contents:

1 Knowledge: The Contemporary Debate.....	2
2 Failing to Know Without Factual Ignorance.....	4
2.1 Paradigmatic Factual Ignorance.....	4
2.2 Known Anti-Expertise.....	5
2.3 Beliefless Remembering.....	7
2.4 Statistically Lucky Belief.....	8
2.5 Environmentally Lucky Belief.....	9
2.6 Safe Defeated Belief.....	10
3 The Ignorance Dilemma.....	11
4 Factive Pluralism.....	13
4.1 Factive Relations: Modal/Amodal, Stative/Non-Stative.....	13
4.2 From Non-Ignorance to Factual Awareness.....	15
5. Solving the Ignorance Dilemma.....	16
References.....	17

1 Knowledge: The Contemporary Debate

We don't know what knowledge is. There is, and has for decades been, so much persistent disagreement about the nature of knowledge among theory of knowledge researchers that any claim to know what knowledge is would have to be accompanied by a defense of the idea that persistent, widespread expert disagreement is not a defeater. I, for one, do not wish to defend *that* thesis. So I concede an uncomfortable truth for an epistemologist: we don't know what knowledge is.¹

But we should not for that reason conclude that we cannot easily identify many paradigmatic cases of knowledge. We also should not conclude that we cannot identify a set of conditions that are non-trivially jointly sufficient for knowledge. We can do both. You know that there is something above your head (if you're reading this indoors). You know that you are now reading an English sentence. You know that no square is a triangle. In these paradigmatic cases of knowledge a set of conditions obtain that are clearly jointly sufficient for knowledge.

Letting ' p ' be what is known in any case like these, the following conditions are jointly satisfied:

Truth: p is true.

Belief: You believe that p .

Justification: You believe that p for sufficient reasons.

Reliability: Your belief that p was formed in a suitably reliable way.

Safety: The belief you formed could not have easily been false.

Sensitivity: Were p false, you would not have believed that p .

These conditions demand further specification and most epistemologists would hold that, when suitably clarified, the conjunction of these conditions is sufficient for knowledge. At the same time, however, virtually all epistemologists would reject the idea that this set of conditions is jointly necessary for knowledge, wanting to strip away at least one of the conditions noted above.² Among these there are some who prefer an iconoclastic approach, injecting a bit of circularity into the conditions that remain or else defending the idea that knowledge is hardly better than true belief.³ These are the focal points of the contemporary debate about knowledge.

To give you a sense of the variety of views out there, look at Figure 1 which provides an unforgivably coarse-grained and incomplete summary:

¹ For those sympathetic to these ideas about defeat in philosophy see Goldberg (2013), Kornblith (2013), and Barnett (2019).

² For a certainty-first approach that is argued to validate all of the conditions above under some specification, see Climenhaga (2023). For an innovative challenge to the safety condition see Kelp's (2019: 52-53) safety dilemma. For responses see Mortini (2022) and Silva (2023: 161-165).

³ For knowledge-first circular approaches see Williamson (2000, 2009), Miracchi (2015), Kelp (2019; cp. 2021), and Simion (2019). For a virtue-theoretic rationality-centered approach see Wedgwood (2020).

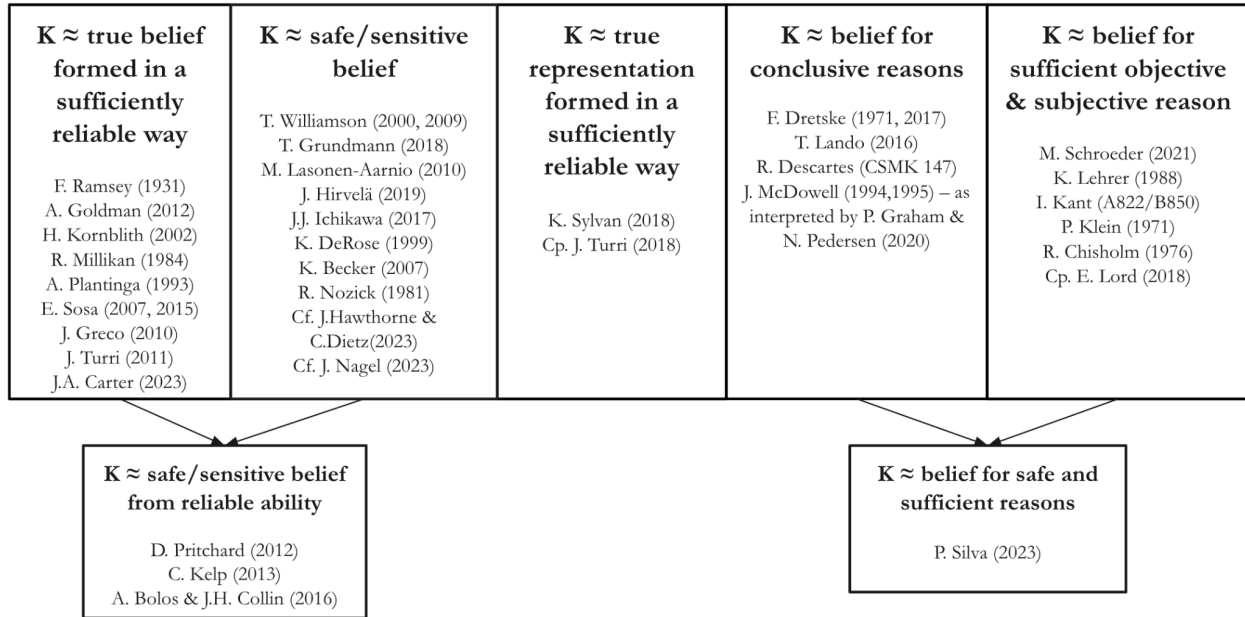


Figure 1. An Incomplete Mapping of Theories of Knowledge

Figure 1 fails to highlight many important distinctions among theories of knowledge, even within each category. So the symbol ‘K ≈’ should, for the most part, be read as ‘Knowledge is *very approximately* taken to be...’. For example, in the first cell we have broadly reliabilist views of knowledge all bundled together, and these views include process reliabilism, robust virtue reliabilism, and proper functionalism about knowledge. These are importantly different views. Similarly, no careful distinctions are drawn between safety and sensitivity theories nor the different ways of spelling out safety and sensitivity conditions. Innovative knowledge-first accounts are largely left out of the picture. Hetherington’s (2001) view of knowledge as true belief that enjoys at least a very small degree of justification is left out, as are explanationist approaches to knowledge.⁴ But all will need to confront the ignorance dilemma.

The ignorance dilemma is an explanatory dilemma. It involves a class of cases, $\alpha_n\text{-}\alpha_m$, about which we should make two judgements about an agent S and a fact p :

Knowledge and Ignorance Are Jointly Absent

¬Kp & ¬Ip: In each of $\alpha_n\text{-}\alpha_m$, S does not know that p and S is not ignorant of the fact that p .

Every theory of knowledge, T , will then be impacted by at least one of these two horns:

Horn#1. T is too strong, predicting that S does not know that p despite not being ignorant of the fact that p in $\alpha_n\text{-}\alpha_m$. Thus, T cannot explain the absence of ignorance in $\alpha_n\text{-}\alpha_m$.

Horn#2. T is too weak, predicting that S knows that p and is, therefore, not ignorant of the fact that p in cases $\alpha_n\text{-}\alpha_m$. Thus, T cannot explain why knowledge is absent in $\alpha_n\text{-}\alpha_m$.

⁴ See Bogardus and Perrin (forthcoming) and David Faraci (2019) for discussion and references.

If your theory of knowledge is afflicted by Horn#1, as nearly all are, be of good cheer. Horn#1 cannot and should not be avoided. To accommodate the non-ignorance data you need only to situate your theory of knowledge *within* a broader awareness-first framework along the lines suggested below. This will afford you the power to explain when and why factual ignorance is absent in α_n - α_m .

2 Failing to Know Without Factual Ignorance

This section gathers the data that drives the ignorance dilemma. It begins with a couple clear cases of factual ignorance that provide a contrast against which we are to understand the cases of non-ignorance that follow them.

Bear in mind that there are different objects of which one can be ignorant: one can be ignorant of *particulars and their properties* (e.g. the cat's existence and its color), one can be ignorant of *propositions* (e.g. the meaning of the sentence 'Cats cannot exist without quarks'), one can be ignorant of *facts/truths* (e.g. the fact that the nearest cat is vicious), one can be ignorant of *how to do things* (e.g. how to avoid the vicious cat), and one can be ignorant of *qualia* (e.g. what it is like to be attacked by the vicious cat). Our only concern in what follows is with **factual ignorance**, i.e. the kind of ignorance relation that takes a factive complement, which we refer to in natural language with the long expression 'ignorant of the fact that *p*' or the contracted expression 'ignorant that *p*'.⁵

2.1 Paradigmatic Factual Ignorance

Let's start with two paradigmatic cases of factual ignorance *and* the absence of knowledge. Here's the first case:

Caesar. Sam considers the claim: (C) that within 10 hours of crossing the Rubicon, Caesar wondered whether he would be a good father. Sam reasons that Caesar was a man who likely hoped for a biological legacy. So, perhaps, he wondered about that before crossing the Rubicon. But Sam also reasons that Caesar was likely pressed by other thoughts leading up to that momentous decision. So Caesar might well have failed to reflect on whether he would be a good father in that timeframe.

There is some fact of the matter here: (C) is either true or it's false. But Sam, like you and I, are completely in the dark. The fact, whatever it is, is something about which we are ignorant.

Here's the second case of ignorance. This one is provided by Littlejohn (2017: 26-27):

Experience Machine. Consider Nozick's (1981) experience machine. Agnes undergoes a series of experiences that dispose her to form false beliefs about her surroundings. It seems to her that she and everyone she cares about are flourishing. In the standard telling, her beliefs are all mistaken. This is not essential to the story. Agnes can be cut off from reality even if some of her beliefs happen to be true. Let us suppose that it [visually] seems to her that her brother has just crossed the stage at graduation and a smile stretches across Agnes' face because she believes he just graduated. What the lab technicians do not realize is that *precisely* as Agnes undergoes this experience her brother crosses the stage and accepts his diploma. While she believes correctly that her brother is graduating and is happy because she

⁵ See Le Morvan & Peels (2016), Peels (2023), Meylan (2024), and El Kassas (2025) for further discussion of the varieties of ignorance.

believes this, her reason for being happy is not that her brother is graduating. She cannot be rationally guided by such a fact, not when she's cut off from reality. (Littlejohn 2017: 26-27)

Littlejohn's use of the colloquial expression 'cut off from reality' is clearly being used to express the idea that Agnes is ignorant of some relevant fact(s) even though she appears to herself not to be. The fact that Agnes seems *to see her brother graduating* gives her an apparent reason to believe that her brother is graduating. But she does not see him do anything, and so she has a *merely* apparent reason to believe that he is graduating. And merely apparent reasons, as Littlejohn notes, do not put people in touch with facts in a way that excludes ignorance and enables one to treat the fact as a reason.⁶

So Agnes in Experience Machine has merely apparent reasons to believe some fact, while Sam in Caesar does not. In this way, the ignorance of Agnes and Sam is explained by the following condition:

Insufficient Non-Accidental Support

There is some salient fact p , but S fails to possess facts that non-accidentally support the truth of p to a sufficient degree.

In the case of Sam, there is an insufficient degree of support. In the case of Agnes there is a strong degree of support for the claim that her brother is graduating, but it is only coming from her non-factive visual experiences that are produced in an unreliable way. So the merely apparent reasons she has only accidentally support the claim that her brother is graduating.

These reflections suggests the following diagnostic principle for whether an agent is ignorant in a given case:

Typical Factual Ignorance

TFI: Typically, S is ignorant of the fact that p just when (and because) S fails to possess facts that non-accidentally support the truth of p to a sufficiently strong degree.

The idea this is meant to capture is that factual ignorance of p is *normally* indicated and explained by the factors just indicated.⁷ This principle clearly calls out for further clarification in relation to the non-accidentality relation and what degree of evidential support is sufficient to exclude ignorance. It also calls out for discussion of how and why the ideology of evidential support is used rather than reasons for belief. But I will not here attempt to defend any such accounts.⁸ My aim is just to contrast the paradigmatic cases of ignorance above with, what will be argued to be, clear cases of non-ignorance below. These limited aims can be achieved without a complete theory of what kinds of factors explain the presence/absence of factual ignorance, and these judgments can be upheld even if TFI turns out to be best thought of as a rough heuristic for identifying (non-)ignorance.⁹

⁶ For a recent development of the normative profile of factual ignorance see Silva and Siscoe (2024).

⁷ In addition to being supported by cases like those above, it is also supported by the idea that factual ignorance just is the absence of factual awareness (Silva & Siscoe 2024) and that factual awareness and its degrees are determined by one's possession of facts that non-accidentally support truths (Silva 2023; Silva & Siscoe 2025).

⁸ See Silva and Bernecker (2023) for discussion of how self-fulfilling and self-defeating beliefs can sometimes make it preferable to use the ideology of evidential support over reasons for belief. For earlier discussions of such cases see Foley (1991) and Conee (1987). For characterizations of evidential support and defeat that are distinct but relevantly related see Simion (2024).

⁹ TFI should not be resisted on the grounds that it is inconsistent with the phenomenon of basic knowledge. It is not. For discussion of when and how the fact that p can itself be one's reasons for believing that p see Silva (2022, 2023:Ch7).

2.2 Known Anti-Expertise

Take the omissive Moore-paradoxical schema ‘ p , and I do not believe p ’. It has often been observed that one cannot truly believe any present tense, first person omissive Moore-paradoxical statement of that form. The reason is that believing the first conjunct (i.e. believing p) is logically incompatible with the second conjunct (i.e. I do not believe p).

Omissive Moore-paradoxical propositions are not unusual. Indeed, they are frequently true of us and our evidence regularly supports them. This has often been overlooked. But *every* time we gain conclusive evidence for some true claim p that we know we do not yet believe, we also gain evidence to think: p , but I don’t believe it. For example, I know that I have no beliefs about the 51st digit in the decimal expansion of π . So at the instant, t , that I obtain conclusive evidence to think that that digit is n my total evidence strongly supports the claim that:

$q \& \neg \text{Bel}(q)$: n is the 51st digit in the decimal expansion of π , but I do not believe it.

Of course, my situation changes if, *after* t , I properly update my beliefs in response to my freshly acquired evidence. For after I update my beliefs I come to believe that: n is the 51st digit in the expansion of π . And I will, often enough, believe that I believe it. But *just prior* to that update – *at time* t – it is possible to have even recognizably conclusive evidence for the Moore-paradoxical claim that: n is the 51st digit in the expansion of π , but I don’t believe it.

Notice that the kind of ‘recognizably conclusive evidence’ in play here is evidence that could put *other people* in a position to know that: n is the 51st digit in the decimal expansion of π , but *I* do not believe it. So this is the kind of evidence which, if possessed by someone, *excludes ignorance*. And you could possess that evidence in the kind of case just described at the moment *before* your belief update (provided you’re a sufficiently self-aware thinker who knows that their belief updates take at least a little time). So it is possible for you to fail to be ignorant of the fact that $q \& \neg \text{Bel}(q)$ while not being in a position to know it. Since you cannot know p if you are not in a position to know p , this is a case where you can fail to be ignorant that $q \& \neg \text{Bel}(q)$, while also failing to know it (cf. Wedgwood 2014).

A particularly interesting omissive Moorean conjunction comes from John Buridan (c.1300-c.1358), who asked his readers to consider the sentence that says *I don’t believe this claim* (Buridan 1982). So take the claim $(B) = I \text{ do not believe } (B)$. (B) is logically equivalent to ‘ (B) , and I do not believe (B) ’. Accordingly, this is just a special instance of the omissive Moore paradoxical conjunction above where both conjuncts are identical. In this case we can make perfect sense of my refusal to *at any point in time* to form a belief in (B) . For I know that: if it’s believed, then it’s false; and if it’s not believed, then it’s true. Assuming modest introspective abilities, I know that I don’t believe (B) owing to its odd logical features relative to whether I believe it. In which case, *I know that in suspending and, hence, failing to believe it, it’s guaranteed to be true*. In which case, I have conclusive evidence for it and I’m able to recognize it as such. Now, in general, if I know I have conclusive evidence for p , then I’m not ignorant of the fact that p . But if I don’t believe p , then I don’t know it. So there’s a gap between the absence of ignorance and the absence of knowledge.¹⁰

Let us say a **known anti-expertise case** is any case, α_n , such that there is some thinker S for whom the following conditions obtain: (i) given what S knows, p is true iff S does not believe p , and (ii) S does not believe p , and (iii) S knows $e =$ (i) and (ii) obtain. With such cases we can make the following argument:

¹⁰ For discussion, defense, and relevant references for the belief requirement on knowledge and the possibility of known anti-expertise cases see Silva (2023: Ch2, Ch4).

1. There are possible known anti-expertise cases, i.e. possible case in which the following conditions obtain:
 - (i) given what S knows, p is true iff S does not believe p .
 - (ii) S does not believe p .
 - (iii) S knows $e =$ (i) and (ii) obtain.
 2. If S is a competent reasoner who knows a conjunction of the form $e \& e \leftrightarrow p$, then S is not ignorant that p .
 3. If S does not believe that p , then S doesn't know that p .
- Therefore,
4. There are possible situations in which S does not know that p , but S is also not ignorant that p . (fr. 1-3)

Premise 1 was defended above. Premise 2 is supported by the diagnostic principle, TFI, since knowing a conjunction of the form $e \& e \leftrightarrow p$ entails: that one has $e \& e \leftrightarrow p$ as evidence for p , that p is a fact, and that one has maximal evidential support for that fact. Premise 3 is a widely shared view in contemporary epistemology that will not be defended here.

2.3 Beliefless Remembering

Here is a case inspired by the work of Sven Bernecker (2009), who argues that factual memory does not depend on belief:

Beliefless Remembering. Jamie studies very hard for an exam on English history. Weeks ago she learned and came to know that (E) *Queen Elizabeth died in 1603*. And she reviewed the information to such an extent that it is stored in her long term memory so that she's able to very easily recall the information in the future. However, she later learns that 15% of the students had been studying from outdated textbooks, and that many – though not most – of the dates in the outdated textbooks are incorrect. She doesn't know whether she used the outdated textbook. So she suspends belief in (E). But she remains much more confident than not that (E) is true because it's much more likely than not that her textbook was reliable. But, unlike some of the others, her teacher ensured that she was given and studied from the accurate and authoritative textbook.

Here are the salient details of the case:

- D1. Jamie came to know that (E) on the basis of her study of a reliable, expertly written textbook, where her reception of that textbook was not an accident.
- D2. Jamie stored the information (E) so firmly in mind that she could easily recall it upon considering the question: when did Queen Elizabeth die?
- D3. Jamie later gives up her belief in (E), but remains much more confident than not that (E) is true, and this confidence is grounded in (i) her study of the reliable, expertly written textbook, and (ii) her knowledge that it's very unlikely that this date is wrong.

Those sympathetic to Bernecker's (2009: Ch3) theory of factual memory will be sympathetic to the idea that here we have a case where S remembers and can easily recall that p , while S fails to believe that p . It is near compelling to maintain the following:

If S came to know p at t and at $t+$ S continues to remember that p and can easily recall that p , then, even if S fails to believe that p , S is not ignorant that p .

Those in the grip of the ideas that remembering entails knowing and that factual ignorance just is failing to know, will likely resist this claim.¹¹ But we can still get what we need from the case above without relying on claims about factual memory. All we need is the following:

S is not ignorant of the fact that p if: (i) S came to know that p on basis b at some earlier time, (ii) S retains b and S knows that the objective chance that b is misleading is very low, and (iii) S is much more confident than not that p on basis b .

Whether or not one is willing to claim that S remembers that p when (i)-(iii) are met, one should at least feel the pull of the idea that S is not ignorant of the fact that p when (i)-(iii) are met. Indeed, the TFI further supports this intuitive judgment. For when (i)-(iii) are met an agent has reasons that non-accidentally support p to a very strong degree.

Now, as noted in the previous section: if S fails to believe that p , then it follows that S also fails to know that p . In which case, in Beliefless Remembering we have a case where ignorance and knowledge are *both* absent.

2.4 Statistically Lucky Belief

Take a standard lottery case:

Statistically Lucky Belief. Larry has a ticket in a fair lottery with very long odds, and he knows that there is a random process that will decide the winning ticket number. The winning lottery number has been drawn, although Larry has not heard the result yet. But reflecting on the random process that will determine the winning number, and the extremely high chance this gives to his ticket being a loser, Larry comes to believe (L) *that his ticket is a loser*. At the same time Larry knows that his ticket might be a winner. After all, he recognizes that his ticket has just as good a chance of winning as every other ticket. Besides Larry's competent assessment of the high objective chance of losing, Larry has no other evidence to think his ticket is a loser. As it turns out, Larry's belief that he owns a losing ticket is true – his ticket really is a losing ticket.

Here are the salient features of this kind of case:

- L1. Larry knows the process by which his ticket's status will be determined, and knows that while it is a random process, it's a process that ensures his ticket will very likely be a loser.
- L2. Larry's ticket is determined to be a loser by that process functioning in the way that Larry knows it is supposed to function.
- L3. Larry believes his ticket is a loser on the basis of his knowledge of how the winning ticket number is selected.

¹¹ See Bernecker (2009: Ch3) for arguments against this entailment, and see Silva (2023: Ch4) for discussion of an alternative entailment relation: remembering that p entails being aware that p , and that being aware that p does not entail believing or knowing that p .

Notice that this is completely unlike the cases of ignorance in Section 2.1. Larry's belief in (L) is not formed on the basis of evidence that only weakly supports (L) as in Caesar, nor is Larry's belief in (L) based on reasons that are only accidentally related to the truth of (L) as in Experience Machine. Quite the opposite: the facts that Larry has access to make it extremely likely that (L) is true and those facts are the very reasons that explain why (L) is true, i.e. that the lottery is designed to randomly select a winning ticket number from among a very large set of possible numbers, and it does just that.¹² For these reasons the TFI would suggest that here we have another case of non-ignorance.

More generally, this is simply not a case where it is at all intuitive to say that Larry is ignorant of the fact that his ticket is a loser. Indeed, from the first-person point of view, if Larry is a typical thinker, he would not be at all surprised to learn in a newspaper that some other ticket turned out to be the winner. He would simply think that (L) turned out to be true for the very reasons he thought it would.

Perhaps we can package the key aspects of this case into a claim:

Necessarily, if (i) S knows that p is objectively very likely to be true because of q , and (ii) p is true because of q , then S is not ignorant of the fact that p when S believes p on the basis of q .

It is noteworthy that epistemologists drawn towards the idea that knowledge lacks an anti-luck requirement have developed views of knowledge that tend to verify the idea that knowledge is present in lottery cases and cases of environmental luck.¹³ And if knowledge is present in lottery cases, then ignorance is absent. For these reasons, the claim of non-ignorance in lottery cases will be viewed favorably by many epistemologists.

In what follows I'm not going to directly defend the no-knowledge judgments in this case or the cases that follow in 2.5 and 2.6. Many, including myself, have done that elsewhere. However, I do have an indirect or 'debunking' argument to offer, but it can only be appreciated against the backdrop of the conclusions of this paper. I'll say a bit more about that below.

2.5 Environmentally Lucky Belief

The standard example of environmental luck cited by epistemologists involves fake barns that appear real from the roadside. But there are more realistic examples we could consider.

MF Doom was a rapper with a signature look: it involved wearing a metal mask modeled after the Marvel comic super villain Victor von Doom. Doom also had a signature act of deception. As Ta-Nehisi Coates (2009) relays it, Doom 'routinely sends out one of his comrades in the Doom costume and has him lip-sync the entire show. He sees this as a logical extension of the Doom [supervillain] idea.' When asked why he would deceive and frustrate his fans in this way, Doom responded: 'I don't leave any money on the table. But sometimes it's not enough money. When it's not enough money, I send the impostor.' (I guess that's the kind of thing a rap supervillain might be expected to do.) When further asked if this was wrong, Doom answered: 'I'm the fucking supervillain. I'm not your friend. You don't have to like me. You're paying for the experience of

¹² See Hicks & Wilson (2021) for more on how chance processes can figure in explanations.

¹³ The broadly reliabilist theories of knowledge noted in Figure 1, including Sylvan (2018) and Turri (2018), tend to have this implication.

dealing with a supervillain.’ (I guess that’s also the kind of thing a rap supervillain might be expected to say.)¹⁴

In any event, as Clayton Littlejohn has pointed out to me in conversation, Doom makes for a great example of environmental luck. Here is one way to tell such a story:

Environmentally Lucky Belief. Doom’s concert is about to begin. Clayton is headed there and is also a Doom expert. He knows how Doom sounds and moves when he performs so well that he’s never mistakenly identified an imposter as Doom. (To make things somewhat more precise, assume Clayton can tell Doom from his imposters 99.99% of the time and assume that Clayton knows this about himself.)

Clayton is late to the show and misses the first and second rap sets. These sets are performed by an imposter. But it’s an atypically skilled imposter. This imposter would fool even Clayton. So had Clayton made it to the first rap set, he would have mistakenly identified the imposter as Doom. And had he made it to the second rap set, he would have also mistakenly identified the imposter as Doom.

But Clayton only arrived during the final rap set. But at this point Doom had switched places with the imposter. Clayton walks in. He takes in the performance. He looks directly at Doom, listens carefully to Doom, and exercising his highly trained skill to discern Doom from imposters, he correctly identifies and comes to believe (D) *that Doom performed the final rap set.*

Is Clayton ignorant of the fact that Doom performed the final rap set? No, he’s not. After all, Doom in fact performed the final set and the reason Clayton believes it is because (i) Clayton is highly skilled at discerning Doom from imposters, and (ii) Clayton used that skill because he was seeing Doom perform the final set.¹⁵

It is condition (ii) that separates this case from that of Experience Machine in Section 2.1. For in that case Agne is, we can assume, skilled at identifying her brother in graduation-like circumstances. But she did not exercise that skill *because she was seeing her brother in graduation-like circumstances.* What she saw was *an illusory image* of her brother in graduation-like circumstances, and the generation of the image was in no way related to the fact that her brother was graduating. It was an image generated for the sole purpose of making her feel happy.

As with the previous section, it is noteworthy that epistemologists drawn towards the idea that knowledge lacks an anti-luck requirement have developed views of knowledge that verify the idea that knowledge is present in cases of environmental luck.¹⁶ And if knowledge is present in such cases, then ignorance is absent. For these reasons, the claim of non-ignorance in this case will be viewed favorably by many epistemologists.¹⁷

¹⁴ Dialogue reported by Byron Bowers on the Hannibal Buress Podcast, *Handsome Rambler*, with Talib Kweli and Byron Bowers (26 April 2017).

¹⁵ See McGlynn (2014:63), Piedrahita (2021), and Pritchard (2021: 238-240), who also argue that the failure to know due to environmental luck doesn’t leave one in the dark in such cases. Each of these authors, like Huemer (2001: 51-92) before them, indicate that there can be a gap between knowledge and ignorance.

¹⁶ See the broadly reliabilist views in Figure 1 as well as Sylvan’s (2018) and Turri’s (2018) views.

¹⁷ Indeed, many non-epistemologists have been shown in experimental studies to agree with the idea that knowledge is present in cases of environmental luck. See Colaço et al. (2014), Horvath & Wiegmann (2016), and Turri (2017). See Silva (2023:Ch8-9) for a diagnosis of the disconnect between the folk and the philosophers on this point.

2.6 Safe Defeated Belief

Take another case involving memory, once again inspired by Bernecker (2009):

Safe Defeated Belief. Last week Sam was studying history from a completely reliable textbook. In it he learned (and came to know) that (R) the Roman Colosseum was completed in CE 80. But today, Sam's classmates who are in competition with him to be valedictorian tell him that (R) is false because the Colosseum was not completed until CE 90. They present him with very plausible but misleading evidence to this effect. But Sam disregards their counter evidence for no particular reason. Rather, he continues to believe (R) on the basis of the textbook. (cf. Bernecker 2009: 78)

If knowledge requires having sufficient reason to believe, then being given credible counter-evidence is the kind of thing that defeats knowledge. We could strengthen the case by making it not fellow students, but tutors or lecturers who provide Sam with misleading counterevidence.

Now, Sam disregards the counter evidence and dogmatically sticks to his original belief in (R), and he continues to hold it for the reason he originally formed it, i.e. the fact that (R) was asserted in his textbook. But if he originally knew (R) because he believed it on the basis of his textbook's assertion, and he continues to believe (R) for that same reason, then the idea that he is ignorant of the fact that (R) seems strange. At the very least, Sam is in a position unlike the agents in Caesar and Experience Machine. And the TFI supports the idea that Sam is not ignorant. Again, the judgment of non-ignorance here will have a tendency to be affirmed by epistemologists who hold that safe/sensitive beliefs or beliefs held for conclusive (\approx safe or sensitive) reasons are beliefs that constitute knowledge.¹⁸ And since knowledge excludes ignorance, these epistemologists are under pressure to agree with the present idea that this is not a case of ignorance.

3 The Ignorance Dilemma

At this point the ignorance dilemma comes into view. Exactly how the dilemma plays-out for each reader will depend on their view about the nature of knowledge. I'll take the liberty of anchoring discussion of the dilemma against the backdrop of the idea that knowledge is *safe belief for sufficient reasons*. On this view knowledge requires at least the following: truth, belief, safety, and the absence of defeaters (where defeaters are reasons that prevent one from having sufficient reason to believe p even if one's reasons for believing p ensure one's belief is safe). This is a view of knowledge that, when suitably spelled-out and embedded in an awareness-first framework, might have the potential to explain the data that has anchored the intuitions driving epistemological theorizing in the last sixty years and also might have the potential to explain why certain debates in the theory of knowledge have been intractable (Silva 2023).¹⁹

¹⁸ See Figure 1 for references. See also Lasonen-Aarnio (2010) for discussion and defense of knowledge in such cases.

¹⁹ This view of knowledge is akin to Descartes conception of knowledge as 'conviction based on a reason so strong that it can never be shaken by any stronger reason' (1988: CSMK 147). It is also akin to Fred Dretske's (2017: 349; cf. 1971) conception of knowledge as belief for conclusive reasons, where 'your reasons for believing p are so strong that, given simply these reasons, you can't be wrong...'. It is also akin to Mark Schroeder's (2021: 224-242) Kantian conception of knowledge as belief for sufficient objective and subjective reasons, where one's reasons are objectively sufficient just in case the facts one has are not defeated by *any* unknown facts. Objective sufficiency is the sort of thing that, according to Schroeder, ensures not only that one's belief is justified by the facts, but also ensures that one could not have easily been wrong in cases of environmental luck (Schroeder 2021: 229-230). Schroeder is less clear about the implications of his view for lottery cases.

Any view of knowledge that requires belief, safety, and the absence of defeaters is a theory of knowledge on which *none* of the cases from 2.2-2.6 is a case of knowledge. Accordingly, any such view of knowledge will be one that validates the \neg Know column in Table 1:

	Knows	\neg Know		\neg Ignorant	Ignorant
α_1 : Known Anti-Expertise		✓		✓	
α_2 : Beliefless Memory		✓		✓	
α_3 : Statistically Lucky Belief		✓		✓	
α_4 : Environmentally Lucky Belief		✓		✓	
α_5 : Safe Defeated Belief		✓		✓	

Table 1. Mapping the Joint Absence of Knowledge and Ignorance

You might not agree with *all* of these case judgements. You might prefer a more permissive theory of knowledge. You might prefer a more demanding theory of ignorance. But so long as you agree that *in at least one* of these cases there is the absence of both knowledge and ignorance, then you will be in a position to appreciate the following dilemma.

Let α_n be any situation in which there is some fact p such that a thinker neither knows nor is ignorant that p . For any theory of knowledge T one of the following holds:

Horn#1. T is too strong, predicting that S does not know p despite not being ignorant of the fact that p in α_n . Thus, T cannot explain the absence of ignorance in α_n .

Horn#2. T is too weak, predicting that S knows that p and is, therefore, not ignorant of the fact that p in α_n . Thus, T cannot explain why knowledge is absent in α_n .

To see the dilemma in action take three competing theories of knowledge. Dretske's (1971, 2017) view that *knowledge is belief for conclusive reasons*. This implies that there is no knowledge in α_1 - α_4 because conclusive reasons are, according to Dretske, sensitive reasons for belief. And in α_1 - α_4 the target thinker either fails to believe p (α_1, α_2) or she fails to believe p for sensitive reasons (α_3, α_4). Provided the target thinker is not ignorant in at least one of those cases, this creates an explanatory gap. In virtue of what is that thinker not ignorant of the fact that p ? This is something Dretske's theory of knowledge cannot explain. Since one can believe p for a conclusive reason even if one has misleading defeaters, this predicts the presence of knowledge in α_5 . At this point the other horn of the dilemma emerges: while this theory rightly predicts non-ignorance in α_5 (because it implies knowledge in α_5 and knowledge excludes ignorance), it is too weak to explain why knowledge is absent in α_5 . (Again, I'm *assuming* what others have argued for elsewhere: that the no-knowledge judgement in α_5 is correct.)

For another example take Pritchard (2012) and Kelp's (2013) *anti-luck virtue epistemologies* on which knowledge is, very roughly, safe belief due to a reliable cognitive ability. This implies that there is no knowledge in α_1 - α_4 because the target thinker either fails to believe p (α_1, α_2) or she fails to safely believe p . Provided the target thinker is not ignorant in at least one of those cases, this creates the explanatory gap noted in Horn#1. And the other horn of the dilemma likewise emerges: while

this theory rightly predicts non-ignorance in α_5 (because it implies knowledge and knowledge excludes ignorance), it is too weak to explain why knowledge is absent in α_5 . Even if the theory came with the ability to handle defeaters, it would still face Horn#1.

Lastly, take Sylvan’s (2018: 199; 213-214) *minimalist theory of knowledge* on which knowledge is, roughly, true representation (e.g. a true belief, or a true seeming state) that is formed in a way that manifests an agent’s reliable cognitive ability to truly represent the world (cp. Turri 2018). Because this theory does not require belief or any anti-luck (safety, sensitivity) condition, α_1 - α_5 cannot be ruled-out as cases of knowledge for those reasons. Moreover, because α_1 - α_5 are cases in which the target thinker represents the relevant fact from an exercise of a reliable cognitive ability, it follows that in *all of the cases* there is knowledge. This allows the theory to entirely avoid Horn#1. For knowledge excludes ignorance. That said, the minimalist theory of knowledge faces Horn#2: it’s such a weak theory of knowledge that it cannot explain why knowledge is absent in any of α_1 - α_5 .

So long as knowledge and ignorance are jointly absent in at least one of α_1 - α_5 , no theory of knowledge will avoid both horns of the dilemma. The solution, I suggest, is to take on a pluralist theory of factive relations that can draw a line between knowledge and a logically weaker amodal factive relation that can explain the absence of ignorance in α_1 - α_5 .

4 Factive Pluralism

Factive pluralism, of the sort I’m concerned with, is the view that there are amodal stative factive relations distinct from knowledge and that also fail to entail knowledge. Factive pluralism is, thus, inconsistent with the idea that knowledge is the only amodal factive stative relation and that every amodal factive stative relation entails knowledge. In this way, factive pluralism is inconsistent with a broadly Williamsonian (2000) approach to factive relations.

4.1 Factive Relations: Modal/Amodal, Stative/Non-Stative

To get a sense of the landscape of factives, take the following factive expressions in Table 2:

Modal Stative Factives	Amodal Stative Factives	Non-Stative Factives
‘sees that’ ‘hears that’ ‘perceives that’ ‘remembers that’ ‘is informed that’ ‘intuits that’ ‘demonstrated/proved that’	‘knows that’ ‘aware [/of the fact] that’ ‘realizes that’ ‘noticed that’ ‘discovered that’ ‘understands that’ ‘learned that’ ‘recognizes that’	‘ignorant [/of the fact] that’ ‘unaware [/of the fact] that’ ‘oblivious [/to the fact] that’ ‘overlooks [/the fact] that’

Table 2. A Partial Mapping of Factive Expressions

Modal stative factive expressions are ‘modal’ in the sense that they implicate a cognitive modality (at some level of specificity) that at least partially *explains how* one came to stand in the implicated stative factive relation at issue (e.g. by seeing, perceiving, remembering, being informed, intuiting,

proving).²⁰ Amodal stative factives, by contrast, are silent about which cognitive modalities were in play. It is in this sense that knowing that p , being aware that p , and realizing that p , etc. are amodal: they fail to imply whether one cognitive modality or another brought it about that one knows that p , is aware that p , realized that p , etc. Importantly, the factive pluralism I have in mind does not imply that for every amodal stative factive expression there is a unique factive stative relation to which it refers. It only requires that this class relations not be limited to knowledge and knowledge-entailing relations.

Non-stative factives are *factives*: you cannot be ignorant of the fact that p unless p is true, you cannot be unaware of the fact that p unless p is true, and so on. But they are *non-statives* in that they entail that there is a stative factive relation that one *fails* to stand in toward the fact that p . For example, necessarily, S is unaware of the fact that p only if that S is not aware of the fact that p . All non-stative factives have this quality. For this reason, when it comes to *denials* of the non-stative factives, the logical implication is that *there exists* some factive stative relation that S stands in. For example, if S is *not unaware* that p , then it semantically follows that S is *aware* that p .

Now there is more to say about the total class of factives than there is space to say here. What matters is that we've seen a range of cases, α_1 - α_5 , in which S is *not ignorant of some fact* and this implies that there exists a stative factive relation that S stands in toward the relevant fact. But we cannot identify that stative factive relation with knowledge. For in cases α_1 - α_5 the target thinker does not have knowledge. So if the stative factive relation is not knowledge, then it must be some other stative factive relation. But what, then, could it be?

Silva and Siscoe (2024) have argued that, contrary to the standard view, factual ignorance is not to be identified with the failure to know. It is, rather, to be identified with the failure to be factually aware.²¹ I will not rely on that identity claim to move forward. I will rely on the modest position that:

The Absence of Ignorance Entails Awareness

$\neg Ip \rightarrow Ap$: Necessarily, if S is not ignorant of the fact that p , then S is aware of the fact that p .

This is something that can be embraced by even those who identify factual ignorance with the failure to know since knowledge is itself awareness-entailing. This conditional is intuitive and its denial is starkly counterintuitive in concrete cases of non-ignorance:

#Buridan is not ignorant of the fact that he does not believe (B), but he is unaware of the fact that he does not believe (B).

#Jamie is not ignorant of the fact that Elizabeth died in 1603, but she is unaware of the fact that Elizabeth died in 1603.

These claims are self-contradictory in the way that claims about married bachelors are self-contradictory. Some might seek to leverage a connection between awareness and consciousness to undermine $\neg Ip \rightarrow Ap$. This conditional cannot, for example, be true if awareness is essentially a conscious state. The reason is that non-ignorance is not essentially a conscious state. Consider the way in which knowledge excludes ignorance, but fails to entail consciousness. Asleep and

²⁰ See French (2013) for discussion of the multiple meanings that can be expressed with 'seeing'.

²¹ For implicit or explicit endorsement of views in the vicinity, see Huemer (2001: 51-92), McGlynn (2014:63), Piedrahita (2021), and Pritchard (2021: 238-240).

unconscious at night one's knowledge states may persist, and if one's knowledge that p persists then one's non-ignorance of the fact that p likewise persists.

But this way of seeking to undermine $\neg Ip \rightarrow Ap$ rests on the mistaken idea that being conscious that p is identical to being aware that p . It is not. Rather, to be conscious that p is better identified with *being consciously aware that p* . And not all awareness states need to be conscious states. That's the whole point of having the modifier 'consciously'. This insight is borne out in cases and Chalmers (1996: 221–22) illustrates this, saying: 'I may be (nonoccurrently) aware that there is a bicycle downstairs, without there being an associated bicycle experience. This sort of awareness without experience is most pronounced with propositional [=factual] awareness...' Factual memory serves as an apt example. You learned, memorized, and have not since forgotten your multiplication tables. So you are aware of the fact that the product of 6 and 4 is 24. But you were aware of that fact *even before* this very example made you consciously aware of it by directing your attention towards it. So being aware that p should be distinguished both from being conscious that p as well as attending to the fact that p . Additionally, cognitive scientists and philosophers study unconscious perception (Prinz 2010; Block 2016). Since the concept of perception is a concept of one way of being aware – i.e. perceptual awareness – the study of unconscious perception is the study of unconscious awareness. While the object of their studies is typically perceptual awareness of particulars, this still provides evidence that the general concept of awareness is to be disassociated from the concept of consciousness.²²

With $\neg Ip \rightarrow Ap$, Table 3 follows from Table 1:

	Knows	\neg Know	Aware	\neg Ignorant	Ignorant
α_1 : Known Anti-Expertise		✓	✓	✓	
α_2 : Beliefless Memory		✓	✓	✓	
α_3 : Statistically Lucky Belief		✓	✓	✓	
α_4 : Environmentally Lucky Belief		✓	✓	✓	
α_5 : Safe Defeated Belief		✓	✓	✓	

Table 3. Mapping Factual Awareness

4.2 From Non-Ignorance to Factual Awareness

What is factual awareness? There are different ways of approaching this question. I'll sketch two ways that have been developed in earlier work, and then turn towards a new way that stems from this paper's reflections on factual ignorance.

The Determinate Approach to factual awareness is a bottom-up approach that asks you to begin with determinate states of awareness such as: perceiving that p , remembering that p , intuiting that p , being informed by S that p , and having proved that p . The background assumption here is that each of these is a determinate way of being aware that p . Since each of these entails being aware that p it

²² Indeed, the academic search engine Google Scholar gives more than 1.4 million hits for 'unconscious perception'. All of the top hits I looked through take it as a serious question whether and in what circumstances there can be unconscious perceptions. This presupposes that the concept of unconscious perception is not self-inconsistent.

follows that factual awareness cannot be more demanding than these states. Each of these determinate states constitutively involves non-accidental representation of the fact that p . For stripping away either the non-accidentality component or the representation component would be inconsistent with each of the determinate states in question. Such reasons support the idea that awareness that p just is some kind of non-accidental true representation of the fact that p . Add to this arguments for the idea that one can see that p , remember that p , etc. without belief, safety, or undefeated justification, and it will follow that factual awareness is a state that is logically weaker than knowledge.²³

The Generalization Approach is a top-down approach. It asks you to begin with *knowledge* conceived of as a high-grade state that requires truth, belief, a safety condition, and an undefeated justification condition. The first move is to generalize away from belief to representation more generally, and hence to move away from the need for undefeated justification for belief. The second move is to generalize away from the modally strong demands of safety, and to the logically weaker demands of reliability. When we do this, we reach a state of reliably (and hence non-accidentally) formed true representation. Since knowledge exists and since knowledge is an instance of this kind of state, it follows that this kind of state exists.²⁴

At this point, it's difficult to make the final step to the claim that factual awareness just is this generalized state of which knowledge is but one instance. The Determinable Approach could thus be leveraged here to make that argument. Or one could just claim that the Generalization Argument gives us the resources to do a bit of conceptual engineering and that epistemologists should just use the expression 'aware that' to refer to *reliably (and hence non-accidentally) formed true representational states*.

However, it is here that our discussion of α_1 - α_5 and the absence of ignorance forges a new link for us to exploit in thinking about the nature of factual awareness. We start with $\neg Ip \rightarrow Ap$ and the judgement that both knowledge and ignorance are absent in α_1 - α_5 . For if ignorance is absent in α_1 - α_5 , then it follows that they are cases of factual awareness. Now we can look at what those cases have in common to help isolate the factors that explain why they are cases of non-ignorance/awareness.

When we look at α_1 - α_5 what we see is a set of cases in which an agent represents a fact p from a reliable ability to do so. In α_1 what we have is a case in which S inferentially represents a fact p by knowing some other claims that S knows entail that p . In α_2 what we have is a case in which S remembers a fact from a reliable ability to thus remember. Remembering that p is a way of representing the world as being such that p . In α_3 - α_5 what we have are a set of cases in which p is true and S believes (and thus represents) that p due to her reliable ability to form true beliefs.

What we see, then, across these cases is awareness of a target fact (truth) that appears to be constituted by an agent's reliable ability to represent facts. Process reliabilists and proper functionalists are free to argue that the ideology of abilities is not necessary to understand what's going on in these cases. That is fine. The space of possible awareness-first epistemologies that can do the necessary theoretical work to address the ignorance dilemma is likely not even limited to these views.

5. Solving the Ignorance Dilemma

We are now in a position to solve the ignorance dilemma. Any theory of knowledge that requires belief, safety, or a no-defeater condition will be one that predicts that knowledge is absent in at least

²³ See Silva (2023: Ch4) for more on this approach, which emphasizes *visually perceiving* that p .

²⁴ See Huemer (2001: 51-92) and Silva (2023: Ch5) for development of this generalization strategy.

one of α_1 - α_5 . Thus, any such theory of knowledge will face Horn#1. For example, if knowledge is *safe belief for sufficient reasons* it follows that knowledge is absent in all of α_1 - α_5 . So this is a theory of knowledge that faces Horn#1, but not Horn#2.

The first part of the solution for theories of knowledge facing Horn#1 is to on-board a modest version of factive pluralism on which factual awareness is, roughly, *reliably formed true representation*. Reasons for adopting such a view have been noted above. The second part of the solution is to on-board the view that factual awareness is sufficient for excluding factual ignorance:

Awareness Excludes Ignorance

$A_p \rightarrow \neg I_p$: Necessarily, if S is aware of the fact that p , then S is not ignorant of the fact that p .

For only with this additional thesis can one leverage the presence of awareness (and the conditions that explain its presence) in α_1 - α_5 to in turn *explain* the absence of ignorance in those cases. Bear in mind, that we used $\neg I_p \rightarrow A_p$ to *identify* which factive stative relation was present in α_1 - α_5 , but it's that relation which in turn is needed to *explain* why ignorance is absent in α_1 - α_5 .

Together $A_p \rightarrow \neg I_p$ and $\neg I_p \rightarrow A_p$ have the making for an elegant theory of factual ignorance, i.e. that factual ignorance is just the absence of factual awareness. This view has been defended on independent grounds elsewhere (Silva & Siscoe 2024) and it offers us just what is needed to provide a fully developed response to the ignorance dilemma.²⁵ We expect too much of a theory of knowledge if we expect it to explain all facts about the absence of ignorance. For that we need a theory of factual awareness, and the reason that a theory of knowledge comes up short here is that knowledge is *but one* (particularly demanding) instance of factual awareness. But there can be factual awareness in the absence of knowledge, and that is just what we find in α_1 - α_5 .

You'll recall that I did not directly defend the no-knowledge judgment in the cases from 2.4-2.6, i.e. in standard lottery cases, cases of environmental luck, and cases involving defeated but safe beliefs. It is at this point that you'll be in a position to appreciate an indirect or 'debunking' argument that has the potential to explain why there has been debate over whether these are cases of knowledge. The argument is this: those epistemologists who have affirmed knowledge in those cases have, in part, been guided by genuine insights about the absence of ignorance in such cases. Add to those insights the implicit assumption that the absence of ignorance entails knowledge, and you'll find these same epistemologists having a tendency to affirm knowledge. But we need not take that last step. Insights about the absence of ignorance do not always carry in their wake facts about the presence of knowledge. For the implicit assumption is, if not mistaken, at least cast into serious doubt by the evidence given above.

To return to the quotes with which we began from Williamson (2000) and Nagel (2023). The arguments of this paper suggest that we can effectively push theory of knowledge research forward by anchoring our work in the contrast between *ignorance* and *an amodal factive mental state*. But the arguments above also suggest that we would be wise to patiently consider the question whether that amodal factive state is knowledge or factual awareness.

One last thing. I began by noting epistemology's present failure to deliver on one of its central aims: it has failed to put us in a position to know what knowledge is. Shall we conclude, therefore, that we are ignorant of what knowledge is? We need not. For it remains possible to be aware of the fact that knowledge has a certain nature even if we fail to know it. As epistemologists who have devoted an inordinate amount of time to thinking about the nature of human knowledge, we should find that to be a comforting possibility.

²⁵ It also offers us what we need to explain degrees of factual ignorance. See Silva & Siscoe (2025).

Acknowledgments: Special thanks to audiences at Goethe University, University of Delaware, and VU Amsterdam.

References

- Barnett, Zach (2019). Philosophy Without Belief. *Mind* 128 (509):109-138.
- Becker, Kelly (2007). *Epistemology Modalized*. New York: Routledge.
- Bernecker, Sven (2009). *Memory: A Philosophical Study*. Oxford University Press.
- Block, Ned. (2016). The *anna parenina* principle and skepticism about unconscious perception. *Philosophy and Phenomenological Research* 93(2): 452–459.
- Bogardus, Tomas & Perrin, Will (forthcoming). A Defense of Explanationism against Recent Objections. *Episteme*: 1-12.
- Bolos, Anthony and James Henry Collin (2018). A sensitive virtue epistemology. *Synthese* 195(3): 1321–35.
- Buridan, Jean (1982). John Buridan on self-reference: chapter eight of Buridan's *Sophismata*, with a translation, an introduction, and a philosophical commentary. Cambridge University Press. Edited by G. E. Hughes.
- Carter, J. Adam (2023). *Stratified Virtue Epistemology: A Defense*. Cambridge University Press.
- Chalmers, David (1996). *The Conscious Mind*. Oxford University Press.
- Chisholm, Roderick M. (1966). *Theory of knowledge*. Englewood Cliffs: Prentice-Hall.
- Climenhaga, Nevin (2023). How Infallibilists Can Have It All. *The Monist* 106 (4):363-380.
- Coates, Ta-Nehisi (2009). The Mask of Doom. *The New Yorker*. September 21, 2009, issue.
- Colaço, David; Buckwalter, Wesley ; Stich, Stephen & Machery, Edouard (2014). Epistemic Intuitions in Fake-Barn Thought Experiments. *Episteme* 11 (2):199-212.
- Conee, Earl (1987). Evident, but Rationally unacceptable. *Australasian Journal of Philosophy* 65: 316–326.
- El Kassar, Nadja (2025). *How Should We Rationally Deal with Ignorance?: A Philosophical Study*. Routledge.
- DeRose, Keith (1995). Solving the skeptical problem. *Philosophical Review* 104 (1):1-52.
- Descartes, Rene (1988). *The Philosophical Writings Of Descartes*, 3 vols, translated by John Cottingham, Robert Stoothoff, and Dugald Murdoch (Volume 3 including Anthony Kenny), Cambridge: Cambridge University Press.
- Dretske, Fred I. (1971). Conclusive reasons. *Australasian Journal of Philosophy* 49 (1):1-22.
- Dretske, Fred (2017). ‘Golden Gettier’. In Rodrigo Borges, Claudio de Almeida, and Peter D. Klein, eds, *Explaining Knowledge: New Essays on the Gettier Problem*. Oxford: Oxford University Press, 345–54.
- Faraci, David (2019). Groundwork for an Explanationist Account of Epistemic Coincidence. *Philosopher’s Imprint* 19(4), 1–26.
- French, Craig (2013). Perceptual experience and seeing that p. *Synthese* 190(10): 1735–51.
- Foley, Richard (1991). Evidence and Reasons for Belief. *Analysis* 51 (2):98-102.
- Goldberg, Sanford (2013). Defending Philosophy in the Face of Systematic Disagreement. In Diego E. Machuca, *Disagreement and skepticism*. New York: Routledge: 277-294.
- Goldman, Alvin I. (2012). *Reliabilism and Contemporary Epistemology: Essays*. New York: Oxford University Press.
- Graham, Peter J. & Pedersen, Nikolaj J. L. L. (2020). Dretske & McDowell on perceptual knowledge, conclusive reasons, and epistemological disjunctivism. *Philosophical Issues* 30 (1):148-166.
- Greco, John (2010). *Achieving knowledge: a virtue-theoretic account of epistemic normativity*. Cambridge University Press.
- Grundmann, Thomas (2018). Saving safety from counterexamples. *Synthese* 197 (12):5161-5185.
- Hawthorne, J. & C. Dietz (2023). The Safety Conception of Knowledge. In *Externalism About Knowledge*, (ed) Luis Oliveira. Oxford University Press.
- Hetherington, Stephen (2001). *Good knowledge, bad knowledge: on two dogmas of epistemology*. Oxford University Press.
- Hicks, Michael Townsen & Wilson, Alastair (2021). How chance explains. *Noûs* 57 (2):290-315.

- Hirvelä, Jaakko (2019). Global safety: how to deal with necessary truths. *Synthese* 196 (3):1167-1186.
- Horvath, Joachim & Wiegmann, Alex (2016). Intuitive expertise and intuitions about knowledge. *Philosophical Studies* 173 (10):2701-2726.
- Ichikawa, Jonathan Jenkins (2017). *Contextualising Knowledge: Epistemology and Semantics*. Oxford University Press.
- Kant, Immanuel (1996). *Critique of Pure Reason*. Werner Pluhar. Indianapolis: Hackett.
- Kelp, Christoph (2013). Knowledge: The Safe-Apt View. *Australasian Journal of Philosophy* 91 (2):265-278.
- Kelp, Christoph (2019). *Good Thinking*. London: Routledge.
- Kelp, Christoph (2021). *Inquiry, Knowledge, and Understanding*. Oxford: Oxford University Press.
- Klein, Peter D. (1971). A proposed definition of propositional knowledge. *Journal of Philosophy* 68 (16):471-482.
- Kornblith, Hilary (2002). *Knowledge and its place in nature*. New York: Oxford University Press.
- Kornblith, Hilary (2007). Is Philosophical Knowledge Possible? In Christoph Jäger & Winfried Löffler, *Epistemology: Contexts, Values, Disagreement. Papers of the 34th International Ludwig Wittgenstein-Symposium in Kirchberg, 2011. The Austrian Ludwig Wittgenstein Society*. pp. 285-304.
- Lando, Tamar (2016). Conclusive Reasons and Epistemic Luck. *Australasian Journal of Philosophy* 94 (2):378-395.
- Lasonen-Aarnio, Maria (2010). Unreasonable Knowledge. *Philosophical Perspectives* 24 (1):1-21.
- Lehrer, Keith (1988). Metaknowledge: Undefeated justification. *Synthese* 74 (3):329 - 347.
- Le Morvan, Pierre, and Rik Peels (2016). The Nature of Ignorance: Two Views. In *The Epistemic Dimensions of Ignorance*, edited by Martijn Blaauw and Peels, Rik, 12–32. Cambridge: Cambridge University Press.
- Littlejohn, Clayton (2017). How and why knowledge is first. In J. Adam Carter, Emma C. Gordon, Benjamin Jarvis, eds, *Knowledge First*. Oxford University Press: 19–45.
- Lord, Errol (2018). Epistemic Reasons, Evidence, and Defeaters. In Daniel Star (ed.), *The Oxford Handbook of Reasons*:Oxford University Press.
- McDowell, John (1995). Knowledge and the internal. *Philosophy and Phenomenological Research* 55 (4):877-93.
- McDowell, John Henry (1994). *Mind and world: with a new introduction*. Cambridge: Harvard University Press.
- McGlynn, Aidan (2014). *Knowledge First?* Palgrave MacMillian.
- Meylan, Anne (2024). In Defence of the Normative Account of Ignorance. *Erkenntnis* 89:207-221.
- Millikan, Ruth Garrett (1984). Naturalist Reflections on Knowledge. *Pacific Philosophical Quarterly* 65 (4):315-334.
- Miracchi, Lisa (2015). Competence to know. *Philosophical Studies* 172(1): 29–56.
- Mortini, Dario (2022). A new solution to the safety dilemma. *Synthese* 200(137): 1–17.
- Nagel, Jennifer (2023). Seeking safety in knowledge. *Proceedings and Addresses of the American Philosophical Association* 97:186-214.
- Nozick, Robert (1981). *Philosophical explanations*. Cambridge: Harvard University Press.
- Peels, Rik. *Ignorance: a philosophical study*. Oxford University Press.
- Piedrahita, Oscar A. (2021). Lucky Ignorance, Modality and Lack of Knowledge. *Pacific Philosophical Quarterly* (3).
- Plantinga, Alvin (1993). Warrant: The Current Debate. *Warrant and Proper Function*. Oxford University Press.
- Prinz, J. J. (2010). When is perception conscious? In B. Nanay (Ed.), *Perceiving the world*, Oxford University Press: 310–332.
- Pritchard, Duncan (2012). Anti-Luck Virtue Epistemology. *Journal of Philosophy* 109 (3):247-279.
- Pritchard, Duncan. (2021). Ignorance and Normativity. *Philosophical Topics* 49(2): 225-243.
- Ramsey, Frank (1931). Truth and probability. In Braithwaite R. B. (ed.), *The Foundations of Mathematics and Other Logical Essays*. Routledge Kegan & Paul. pp. 156–98
- Schroeder, Mark (2021). *Reasons First*. Oxford University Press.
- Silva Jr., Paul (2022). Basic knowledge and the normativity of knowledge: the awareness-first solution. *Philosophy and Phenomenological Research* 104(3): 564–86.

- Silva Jr., Paul (2023). *Awareness and the Substructure of Knowledge*. Oxford University Press.
- Silva, Paul & Bernecker, Sven (2023). Evidence, reasons, and knowledge in the reasons-first program. *Philosophical Studies* 181 (2):617-625.
- Silva Jr., Paul & Siscoe, Robert Weston (2024). Ignorance and awareness. *Noûs* 58 (1):225-243.
- Silva Jr., Paul & Siscoe, Robert Weston (2025). Awareness by degree. *Philosophy and Phenomenological Research* 110 (1):172-200.
- Simion, Mona (2019). Knowledge-first functionalism. *Philosophical Issues* 29(1): 254–67.
- Simion, Mona (2024). *Resistance to Evidence*. Cambridge University Press.
- Sosa, Ernest (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I*. Oxford, GB: Oxford University Press.
- Sosa, Ernest (2015). *Judgment and Agency*. Oxford University Press.
- Sylvan, Kurt (2018). Knowledge as a Non-Normative Relation. *Philosophy and Phenomenological Research* 97 (1):190-222.
- Turri, John (2011). Manifest Failure: The Gettier Problem Solved. *Philosophers' Imprint* 11.
- Turri, John (2017). Knowledge attributions in iterated fake barn cases. *Analysis* 77 (1):104-115.
- Turri, John (2018). Virtue epistemology and abilism on knowledge. In Heather D. Battaly, *The Routledge Handbook of Virtue Epistemology*. Routledge: 209-316.
- Wedgwood, Ralph (2014). Rationality as a Virtue. *Analytic Philosophy* 55:319–338.
- Wedgwood, Ralph (2020). The internalist virtue theory of knowledge. *Synthese* 197 (12):5357–5378.
- Williamson, Timothy (2000). *Knowledge and Its Limits*. Oxford University Press.
- Williamson, Timothy (2009). *Probability and Danger*. Amherst Lecture in Philosophy.