

ARTICLE

Knowledge is Believing Something Because It's True

Tomas Bogardus^{1*}  and Will Perrin²

¹Pepperdine University, Malibu, CA, USA and ²Northern Illinois University, DeKalb, IL, USA

*Corresponding author. Email: tomas.bogardus@pepperdine.edu

(Received 18 November 2019; revised 4 March 2020; accepted 29 March 2020;
first published online 15 June 2020)

Abstract

Modalists think that knowledge requires forming your belief in a “modally stable” way: using a method that *wouldn't* easily go wrong (i.e. safety), or using a method that *wouldn't* have given you this belief had it been false (i.e. sensitivity). Recent Modalist projects from Justin Clarke-Doane and Dan Baras defend a principle they call “Modal Security,” roughly: if evidence undermines your belief, then it must give you a reason to doubt the safety or sensitivity of your belief. Another recent Modalist project from Carlotta Pavese and Bob Beddor defends “Modal Virtue Epistemology”: knowledge is a belief that is maximally modally robust across “normal” worlds. We'll offer new objections to these recent Modalist projects. We will then argue for a rival view, Explanationism: knowing something is believing it *because it's true*. We will show how Explanationism offers a better account of undermining defeaters than Modalism, and a better account of knowledge.

Keywords: Epistemology; knowledge; modal security; sensitivity; safety; explanationism; explanation; justification; Gettier cases

Introduction

The analysis of knowledge: what all the ages have longed for.¹ If you're an epistemologist, at least. Though the long and fruitless search has jaded many of us, we present here what we believe is the genuine article. To convince you, we'd like to discuss an emerging rivalry between two opposing views of knowledge, which we'll call “Modalism” and “Explanationism.” We'll raise objections to Modalism, and then argue for Explanationism. According to Modalism, knowledge requires that our beliefs track the truth across some appropriate set of possible worlds.² Modalists have focused on two modal conditions. *Sensitivity* concerns whether your beliefs would have tracked the truth had it been different.³

¹Our topic is propositional knowledge, i.e. knowledge-that. We'll leave this qualification implicit in what follows.

²For examples of Modalist projects, see Fred Dretske (1971), Robert Nozick (1981), Steven Luper-Foy (1984), Keith DeRose (1995), Mark Sainsbury (1997), Ernest Sosa (1999a, 1999b, 2007, 2009), Timothy Williamson (2002), Sherrilyn Roush (2005), Tim Black and Peter Murphy (2007), Duncan Pritchard (2007, 2009), Tim Black (2008), Jonathan Jenkins Ichikawa (2011), and Justin Clarke-Doane (2012, 2014, 2015, 2016).

³A more precise statement of the sensitivity condition would specify the *method* being used. For examples of sensitivity theorists, see Fred Dretske (1970), Robert Nozick (1981), and Kelly Becker and Tim Black (2012).

A rough test for sensitivity asks, “If your belief had been false, would you still have held it?” *Safety*, on the other hand, concerns whether our beliefs track the truth at ‘nearby’ (i.e. ‘similar’) worlds.⁴ A rough test for safety asks, “If you were to form your belief in this way, would you believe truly?” Modalism was a natural outgrowth of the revolutionary work done in modal logic during the last century. When Kripke, Lewis, and the gang give you a shiny new hammer, everything starts looking like a nail. Even knowledge.

According to Explanationism, by contrast, knowledge requires only that beliefs bear the right sort of explanatory relation to the truth. In slogan form: knowledge is believing something *because* it’s true.⁵ Less roughly, Explanationism says knowledge requires only that truth play a *crucial role* in the explanation of your belief.⁶ What’s a crucial role? We propose adapting Michael Strevens’ (2011) “kairitic test” for difference-making in scientific explanation, to the more general purposes of Explanationism. With regard to scientific explanation, Strevens’ proposal is this: Start with a deductive argument, with premises correctly representing some set of influences (potential explanans), and the conclusion correctly representing the explanandum. Make this argument as abstract as possible, while preserving the validity of the inference from premises to conclusion; strip away, as it were, unnecessary information in the premises. When further abstraction would compromise the validity of the inference, stop the process. What’s left in the premises are difference-makers, factors that play a “crucial role” in the scientific explanation. Now, explaining why a belief is held is importantly different from paradigm cases of scientific explanation. In order to adapt the kairitic test to Explanationism, we propose beginning with a set of potential explanans that explain the relevant beliefs being held, and then proceeding with the abstraction process until further abstraction would make the explanation fail. The remaining explanans are the difference-makers. On Explanationism, for the belief to count as knowledge, the truth of the relevant belief must be among these difference-makers in order for a belief to count as knowledge.

We believe that Explanationism’s time has come. The iron is hot; the field is white for the sickle. Below, we will raise objections to two recent Modalist projects. The first project comes from Justin Clarke-Doane and Dan Baras (Forthcoming), and centers on a principle they call “Modal Security.” The second project is a defense of “Modal Virtue Epistemology,” by Bob Beddor and Carlotta Pavese (Forthcoming). Our criticisms of these two projects point toward the shortcomings of Modalism generally. We’ll then set Modalism aside, and move to a motivation and defense of Explanationism. We’ll let the view stretch its legs in straightforward cases of perception, memory, introspection, and intuition. Then, we’ll put Explanationism through its paces, with harder cases of abduction, deduction, and induction, as well as a couple of types of apparent counterexamples that might occur to the reader. We’ll close by showing how Explanationism defuses Linda Zagzebski’s Paradox of the Analysis of Knowledge.

⁴For examples of safety theorists, see Mark Sainsbury (1997), Ernest Sosa (1999a), John Hawthorne (2004), and Duncan Pritchard (2005).

⁵Compared with many other analyses of knowledge on offer, this one is suspiciously simple. But we think Aristotle was right that we must not expect more precision than the subject matter admits. Also, we find it charming that Explanationism proposes to analyze as ancient and august a concept as *knowledge* in terms no less venerable: *explanation*, *belief*, and *truth*. On top of its philosophical merits, this is historically and psychologically plausible. For examples of broadly Explanationist projects, see Alan Goldman (1984, 1988), Steven Rieber (1998), Ram Neta (2002), Marc Alspecter-Kelly (2006), Carrie Jenkins (2006), Dan Korman and Dustin Locke (Forthcoming), and David Faraci (2019).

⁶As Alan Goldman put it, the truth of your belief “enters prominently” into the best explanation for its being held (1984: 101).

1. Modal Security

Justin Clarke-Doane, in his earlier work (2015; [Forthcoming](#): §4.6), as well as in more recent work with Dan Baras ([Forthcoming](#): 1), defends the following principle:

Modal Security: If evidence, E, undermines (rather than rebuts) my belief that P, then E gives me reason to doubt that my belief is sensitive or safe (using the method that I actually used to determine whether P).

Though this principle does not mention knowledge, we count it as part of a Modalist project because of the connection between undermining defeaters and knowledge. Baras and Clarke-Doane say that, in contrast to a rebutting defeater, an *undermining* defeater does not target the truth of a belief, and so “it must show that there is some important epistemic feature that our *belief that P* lacks. Safety and sensitivity, along with truth, have emerged from recent epistemology literature as among the most important epistemological features of beliefs.” So, if neither Safety nor Sensitivity is required for knowledge, presumably they’re not “important epistemic features,” the lack of which would allow for an undermining defeater, and doubts about which must be produced by any undermining defeater. And that would be trouble for Modalism generally, and for Modal Security in particular.

Baras and Clarke-Doane respond to several proposed counterexamples in the literature. For example, David Faraci (2019) proposes Eula’s Random Number Generator. Eula has good (but misleading) reason to believe her computer is running a program designed to produce prime numbers. Actually, it’s merely a random number generator. But, by chance it *happens* to produce only primes. And, as Baras and Clarke-Doane say, “these events are maximally deterministic. At every metaphysically – not just physically – possible world,” Eula gets the same results.⁷ So, for any number *n* that the computer spits out, if Eula believes that *n* is prime, her belief is formed safely. (And sensitively).⁸ Despite this safety (and sensitivity), the objection concludes, Eula gets a defeater for her belief when she learns she’s dealing with a random number generator, one that doesn’t produce these outputs *because* they’re prime. This is trouble for Modal Security, which says that defeaters must come by way of safety or sensitivity violations.

Baras and Clarke-Doane notice that the problem with this case is in the idea of a number generator that is both random, but which *must* – or even merely which *would* – deliver only primes. We put the problem this way: If the number generator is truly random, it can’t be the case that it must (or would) produce only primes. This is a problem, since we judge that Eula gets a defeater in this case only insofar as the number generator is genuinely random. But we judge that the number generator is a safe (and sensitive) method of forming beliefs about prime numbers only insofar as it’s not random, but must (or would) produce only primes. In other words, Faraci needs a case in which a subject gets a defeater for a safely and sensitively formed belief. But if

⁷In his own words, Faraci (2019: 22) says of this case, “At every possible world, Eula’s counterpart forms beliefs about which numbers are prime by consulting the Source’s counterpart, and at every possible world, the Source’s counterpart delivers the same answers as at the actual world. [There is no explanatory connection between this modal stability and the facts about which numbers are prime.]” But he adds, in a footnote, that all the argument really needs is that the number generator spit out only prime numbers in all the “nearby” worlds, so that the safety counterfactual comes out true: *if Eula were to form a belief about primes using this method, she would believe truly*. As long as the number generator doesn’t deliver these numbers *because* they’re prime, Faraci thinks, learning that it’s a random number generator should give her a defeater, even without targeting safety.

⁸If her belief were false – that is, if *per impossibile* this number weren’t prime – then the computer would not have produced it as an output, and so she would not have believed that it was prime.

Eula's number generator is truly random, then it's neither a safe nor a sensitive method. Yet if it's not truly random, and for some reason reliably produces primes, then Eula gets no defeater when she learns this.

We would like to present a new problem for Modal Security. Recall that Modal Security says that if evidence undermines your belief, it must give you reason to doubt that your belief is sensitive or safe. And this is because either sensitivity or safety is an "important epistemic feature," which we take to mean a necessary condition on knowledge. For if neither safety nor sensitivity is required for knowledge, it's hard to see how an undermining defeater could possibly succeed by targeting either of these conditions, and harder still to see why an undermining defeater *must* provide reason to doubt that these conditions obtain. If some condition is not necessary for knowledge – if, that is, if one can transgress this condition and still know full well that *p* is true – then why think evidence that this condition isn't met should defeat one's belief that *p*? And if neither sensitivity nor safety is required for knowledge, then that means some other condition is, in addition to truth and belief, since truth and belief are themselves insufficient for knowledge. And, presumably, an undermining defeater could succeed by targeting this other condition, without targeting sensitivity or safety. If these additional assumptions about the relationship between undermining defeat and knowledge are right, then a case of belief that was not formed sensitively or safely and yet is knowledge all the same would refute Modal Security.

We believe there are such cases, and here's a general recipe to whip one up: first, pick the most virtuous belief-forming method you can imagine, and have a subject form a belief via that method. Second, add a twist of fate: put the method in danger of malfunctioning, but let the danger remain purely counterfactual. Now, since things *could* have gone less well epistemically but *didn't*, it is quite tempting to admit that the virtuous method produced knowledge. And yet the belief was not formed safely, for there are many nearby possible worlds in which the method goes awry. And neither need the belief be formed sensitively, since, if the belief were false, it could easily happen that the method still produces that belief. Such a case nicely pries apart our concept of knowledge from our concepts of safety and sensitivity, by polluting very many of the "nearby" worlds with false beliefs, while maintaining a tight enough connection between belief and its truth to allow for knowledge.⁹ What we see in such cases is that, sometimes, our *actual performance* is enough for knowledge, even under threat of externalities.

For example, suppose that the world's most accurate clock hangs in Smith's office, and Smith knows this. This is an **Atomic Clock**: its accuracy is due to a clever radiation sensor, which keeps time by detecting the transition between two energy levels in cesium-133 atoms. This radiation sensor is very delicate, however, and could easily malfunction if a radioactive isotope were to decay in the vicinity. This morning, against the odds, someone did in fact leave a small amount of a radioactive isotope near the world's most accurate clock in Smith's office. This alien isotope has a relatively short half-life, but—quite improbably—it has not yet decayed at all. It is 8:20 am. The alien isotope will decay at any moment, but it is indeterminate when exactly it will decay. Whenever it does, it will disrupt the clock's sensor, and – for fascinating science-y reasons that needn't delay us here – it will freeze the clock on the reading "8:22."

The clock is running normally at 8:22 am when Smith enters her office. Smith takes a good hard look at the world's most accurate clock – which she knows is an extremely well-designed clock that has never been tampered with – and forms the true belief that

⁹According to Explanationism, this "tight enough" connection is the explanation relation. In Atomic Clock, it remains the case that Smith believes it's 8:22 because it is 8:22, even if, due to the threat of this isotope, the actual world is surrounded by an ocean of error.

it is 8:22 am. Now, does Smith know that it's 8:22 am? It strikes us as obvious that she does, and perhaps the same goes for you. Like Harry Frankfurt's (1969: 835–6) counterfactual intervener, the damage of this radioactive isotope remains purely counterfactual, and everything is functioning properly when she forms her belief. Yet, since the isotope could easily have decayed and frozen the clock at "8:22," and since Smith may easily have checked the clock a moment earlier or later, Smith might easily have believed it is 8:22 am without it being 8:22 am. Smith formed her belief in a way that could easily have delivered error. So, while the actual world remains untouched, the threat this isotope poses to the clock pollutes "nearby" modal space with error. This allows us to conceptually separate knowledge from safety.¹⁰

What's more, Smith's belief is not formed sensitively. Recall the test for sensitivity: *If it weren't 8:22, Smith wouldn't believe that it is (via this method)*. Let's run the test. How would things be for Smith and her clock if it weren't 8:22? Well, if it weren't 8:22 when Smith checked her clock, it would be slightly earlier or later. And the isotope has been overdue to decay for some time now. So, if Smith were to check the clock slightly earlier or later, she may well check a broken clock: the isotope may well have decayed, freezing the clock on "8:22" even when it's not 8:22. In that case, Smith would still believe it's 8:22 by checking her clock. So, if it weren't 8:22, Smith may well still believe that it is, using the same method she actually used. So, it's not the case that Smith *would not* believe it's 8:22 by checking this clock, if it weren't 8:22. So, Smith fails the sensitivity condition.

Therefore, neither safety nor sensitivity is a necessary condition on knowledge. So, if undermining defeaters must operate by targeting a necessary condition on knowledge, then undermining defeaters *can't* operate via safety or sensitivity, and a fortiori they *need not* so operate.¹¹ Therefore, Modal Security is false.¹² This ends our discussion

¹⁰For further development and defense of this case, please see Bogardus (2014), as well as Bogardus and Marxen (2014).

¹¹Atomic Clock targets the alleged safety and sensitivity conditions on knowledge. Yet, as we've noted, Modal Security is a principle that concerns only undermining *defeat*. If the additional assumptions connecting undermining defeat to knowledge discussed above give you pause for thought, you may be pleased to learn that Atomic Clock can be modified to target Modal Security directly. In Atomic Clock, there is no safety or sensitivity, but there is knowledge. So, if Smith learns all this, she doesn't get a defeater. But now suppose that the isotope does decay and the clock goes out. It freezes. Learning *that* does give Smith an undermining defeater for her belief. But, since the clock *remains* an unsafe (and insensitive) way to form a belief, the defeater doesn't give Smith a reason to think the clock is unsafe or insensitive. (She was already sure of that.) Or, if you think this still does give her reason to doubt safety or sensitivity, *even more*, we could add another decaying isotope that would flip the clock on, but this decay is really unlikely. We can adjust the probabilities so that the clock is *exactly as unsafe and insensitive as before*. There are exactly the same proportions of "nearby" worlds with a working clock and worlds with a broken clock. Since learning that the clock has frozen in this context would undermine Smith's belief without requiring any change in her confidence regarding safety or sensitivity, this is a problem for Modal Security. Notice also that this case makes trouble even for "Modal Security₃," from Clarke-Doane and Baras (Forthcoming: §8), which says: "If evidence, E, undermines our belief that P, then E gives us direct reason to doubt that our belief is sensitive or safe or E undermines the belief that <the belief that P is safe and sensitive>." For reasons already given, the first disjunct of the consequent of Modal Security₃ is not met. And, plausibly, neither is the second disjunct, for two reasons: First, in the case now under consideration, Smith doesn't even believe that <the belief that it's 8:22 is safe and sensitive>. In fact, since Smith had already been informed of the nearby isotope, Smith already knew that the belief that it's 8:22 is neither sensitive nor safe. Second, even if we read the second disjunct of the consequent of Modal Security₃ in such a way that Smith would need only receive reason to doubt (i.e. lower her credence) in the proposition that <the belief that it's 8:22 is safe or sensitive>, given the facts about the "offsetting" isotope we describe above, plausibly Smith doesn't receive any such reason.

¹²Going by its letter and not its spirit, Modal Security entails merely this: *either* safety or sensitivity is itself required for knowledge, *or* either safety or sensitivity entails some other feature F that is required for

of the first Modalist project we mean to address. We turn now to the second Modalist project: Modal Virtue Epistemology.

2. Modal Virtue Epistemology

In a recent paper, Bob Beddor and Carlotta Pavese ([Forthcoming](#)) endorse a version of the safety condition on knowledge, and they offer this analysis: “a belief amounts to knowledge if and only if it is maximally broadly skillful – i.e., it is maximally modally robust across worlds where conditions are at least as normal for believing in that way.” For a performance to be *maximally modally robust* is for it to succeed in *all* of the relevantly close worlds (Beddor and Pavese [Forthcoming](#): 10). This is why we don’t need a truth condition in their analysis: maximal modal robustness entails truth. As for “normal,” they provide this guide for understanding the notion: “identify the normal conditions for a task with those which we would consider to be fair for performing and assessing the task ... [O]ur intuitions about cases reveal a tacit grasp of these conditions.”

Beddor and Pavese offer a general strategy for responding to challenges to the safety condition, and here’s how it would go with respect to Atomic Clock. The task at hand is: forming beliefs about the time based on checking a clock. If the isotope had interfered with the clock, this interference would have rendered conditions significantly more abnormal for the task at hand than they are in the actual world, where the isotope is present but does not intervene. The decayed isotope puts poor Smith at an unfair disadvantage, as it were. According to Modal Virtue Epistemology, we should ignore situations like that, where the subject is disadvantaged by abnormal, unfair conditions. So, it turns out that Smith’s belief *is safe* (i.e. maximally broadly skillful), since *it is true in all worlds where conditions are at least as normal as they are in the actual world*. This how Beddor and Pavese mean to respond to proposed counterexamples to Modal Virtue Epistemology.

However, there is a serious problem for this Modalist strategy from Beddor and Pavese. On their view, knowledge requires safety, and their understanding of safety requires *maximal* modal robustness. It’s a high-octane brand of safety that they’re

knowledge (so that, if an undermining defeater operates by showing that your belief lacks F, it thereby also gives you a reason to think your belief lacks sensitivity or safety, *en passant*). Atomic Clock deals with the first possibility, showing that neither safety nor sensitivity is required for knowledge. As for the second possibility, neither believing sensitively nor believing safely by itself entails knowledge, because neither one entails truth. Nor do they entail truth even conjoined. And, as Linda Zagzebski (1994) taught us, if the non-true-belief part of a proposed analysis of knowledge doesn’t entail truth, then it should be possible to construct Gettier cases for the analysis. And one can do that for safety and sensitivity like so: pick the most virtuous, safest, most sensitive method you can think of. Like a clock, or a calculator. It’s a safe method, so there are no ‘nearby’ worlds in which you use it and it delivers error – but there are ‘distant’ worlds where this happens. It’s a sensitive method, so, there are no ‘nearby’ worlds in which a proposition is false and yet the method tells you that proposition is true – but there are ‘distant’ worlds where this happens. Now let a person employ this method to form a belief. Suppose this was one of the rare, ‘distant’ occasions when the method got it wrong. But, by sheer coincidence, events transpire that cause the belief to turn out true. A Gettier case. So, the belief will not be knowledge, even though it’s a true belief formed via a method that was both sensitive and safe. So, safety and sensitivity are not enough to entail whatever’s missing from the analysis of knowledge in order to avoid Gettier cases. This shows that there is some anti-Gettier condition that’s required for knowledge, yet not entailed by sensitivity or safety (or both). But then Baras and Clarke-Doane should agree that an undermining defeater might succeed without giving us reason to doubt safety or sensitivity, by instead giving us reason to doubt that this missing, anti-Gettier condition has been satisfied, where this condition is not entailed by sensitivity or safety (or both). And this means Modal Security is false.

using in this machine, requiring success in *all* relevantly close worlds, i.e. worlds where conditions are at least as normal for the task. But what if we're excellent yet not infallible at some task? Remembering what we ate for breakfast, for example. We're *really* good at this. But *sometimes* we get it wrong, even in normal, fair conditions.¹³ Modal Virtue Epistemology says that our belief was *not* maximally modally robust, and so *not* safe, and so *not* knowledge. On their view, knowledge requires infallibility in all scenarios at least as normal or fair as the actual scenario. So ordinary, fallible knowledge is a counterexample to Modal Virtue Epistemology. This is a problem.

But let's be charitable, even if we fail in all other virtues. Suppose Pavese and Beddor retreat to this – safety requires just *substantial* modal robustness, truth in *many* relevant nearby worlds. In that case, we get the second horn of the dilemma. Now the view is: a belief amounts to knowledge if and only if it is broadly skillful *to a high degree*. But, even still, there are counterexamples. Suppose that, very improbably and acting totally out of character, we give you some eye drops that make your vision blurry, and then pressure you to read an eye chart. Your vision is so blurry that either you get the wrong answer, or you get the right answer by chance. Neither one is knowledge. But what does this revised view say? Well, we'll have to check all the 'nearby' worlds at least as normal as the actual world, which will include all the many worlds where, acting in character, we don't administer the eye drops, and you have clear eyes. In many of these worlds, you get the right answer. So, the belief is safe, on this revised view. So, you have substantial modal robustness; the belief is substantially broadly skillful. So, it's knowledge, on this view. But this is certainly the wrong result.

We have good reason, then, to abandon Modal Virtue Epistemology, just as we did with Modal Security. The failures of these two prominent Modalist projects are suggestive of deep flaws within Modalism itself. We believe Modalism was an understandable application of exciting twentieth-century advances in modal logic to longstanding problems in epistemology. And it comes *close* to the truth. But it's possible, albeit difficult, to design counterexamples that tease apart knowledge from sensitivity and safety. So, we believe Modalism misses something crucial about the nature of knowledge: the connection between a believer and the truth can't be fully captured in modal terms, because it's an *explanatory* connection. Modalism, then, can have no long-term success as a research project. It's standing in its own grave. So, it's time to look elsewhere for the analysis of knowledge.

3. Explanationism

What is Explanationism? Alan Goldman (1984: 101) puts it this way: a belief counts as knowledge when appeal to the truth of the belief enters prominently into the best

¹³**Objection:** If we fail to remember what we had for breakfast, something abnormal must have been going on, requiring some special explanation. If so, the conditions weren't fair and normal after all, and so this isn't a counterexample. **Reply:** We doubt that first inference, but perhaps its attraction is due to the fact that knowledge by way of memory is somewhat *immediate*, and effortless. So consider instead a case that is more analogous to Beddor and Pavese's example of shooting hoops, namely a case in which knowledge comes by way of a process, a performance. Consider, for example, taking a standardized test, like the SAT or the GRE. We have a pretty firm grip on what constitutes fair and normal conditions for taking these tests. And, we think, even a fallible test-taker can nevertheless know at least some of the answers to these test questions in such conditions, just as even a successful, professional basketball player can miss some shots occasionally, even in fair and normal conditions. (And, in neither case does failure necessarily require some special explanation.) So, here, we see again that knowledge does not require success in all worlds where conditions are at least as normal for the task at hand. Thanks to Bob Beddor and also Sydney Keough for independently encouraging us to consider an objection like this.

explanation for its being held. And Carrie Jenkins (2006: 139) offers this definition: S knows that p if and only if p is a good explanation of why S believes that p, for someone not acquainted with the particular details of S's situation (an 'outsider').¹⁴ As for what makes something a *good* explanation, we recommend that Explanationism be open-minded, and commit itself only to this minimal understanding of explanation from Goldman (1984: 101): "Our intuitive notion of explaining in this context is that of rendering a fact or event intelligible by showing why it should have been expected to obtain or occur."¹⁵ As for what it is for truth to "enter prominently" or, as we'd say, to *figure crucially* into an explanation, we remind our reader of what we said above regarding Strevens' kairetic test for difference makers: the most abstract version of the explanation will feature all and only the difference-makers. If the truth of the relevant belief is among them, that truth figures crucially into the explanation, and the relevant belief counts as knowledge.

One motivation for Explanationism is the observation that debunking arguments across philosophy – against moral realism, religious belief, color realism, mathematical Platonism, dualist intuitions in the philosophy of mind, and so on – have a certain commonality.¹⁶ To undermine some belief, philosophers often begin something like this: "You just believe that because..." and then they continue by citing an explanation

¹⁴Jenkins (2006: 138) says, "I present this condition, however, not as an analysis of knowledge, but rather as a way of getting a handle on the concept and further the effort to understand what its role in our lives might be." We, however, advocate Explanation as a good old-fashioned *analysis* of knowledge. If you have Williamsonian scruples against the project of analyzing knowledge, please see Faraci (2019: 22–3), who helps dispel those worries.

¹⁵Goldman's talk of "rendering a fact or event intelligible" may occasion the worry that, since what counts as demystifying varies from subject to subject, on Explanationism, whether someone's belief that p counts as knowledge will vary from subject to subject. Some accounts of explanation – known as "pragmatic" or "contextual" accounts – may well have this implication. See, for example, Bas van Fraassen (1980) and Peter Achinstein (1983). But traditional, non-contextual accounts of explanation can give a different, objective reading of "rendering facts intelligible" or "demystifying phenomena." An explanation may succeed, render intelligible a phenomenon, and demystify it, even if someone disagrees or fails to appreciate this. Such an account may also assuage the worry that explanation itself is ultimately analyzed in terms of *knowledge*, perhaps by way of *understanding*. In which case, by analyzing knowledge in terms of explanation, Explanationism will go in a circle. But if explanation is a mind-independent relation, not to be analyzed in terms of mental states like understanding or knowledge, then there is hope for Explanationism to avoid this circle. Though we ourselves favor traditional, non-contextual accounts of explanation, we intend Explanationism to be consistent with each.

One may also worry that, if explanation itself is eventually cashed out in modal terms, then Explanationism is not a viable alternative to Modalism at all, and in fact reduces to a species of Modalism. One *might* think, for example, that if A explains B, then if A hadn't obtained, B would not have obtained either, and if A explains B, then if B were to obtain, then A would obtain as well. And the superficial similarity of such implications – with respect to Explanationism's claim that truth (A) must explain belief (B) – to the Sensitivity and Safety conditions might make one uncomfortable. However, we think both of those inferences above are invalid, subject to refutation by counterexample (e.g. cases familiar from the causation literature, featuring preemption, trumping, etc.). Also, if we're right that neither Sensitivity nor Safety are required for knowledge, and yet the condition laid down by Explanationism is required for knowledge, then this shows that the condition laid down by Explanationism does not entail either Sensitivity or Safety. While we're optimistic that explanation need not be analyzed in modal terms at all, even if that turned out to be the case, the fact that the condition laid down by Explanationism is distinct from both Safety and Sensitivity would still count as an important advance. (We thank our audience at the Institute of Philosophy in London for helpful discussion of this worry.)

¹⁶See Dan Korman (2019) for a nice overview of this style of argument across the entire landscape of philosophy.

that does not feature the truth of this belief. Our evident faith in the power of such considerations to undermine a belief suggests that we take knowledge to require that a belief be held *because it's true*, and not for some other reason independent of truth. Explanationism agrees.

We'll now further motivate Explanationism by taking the reader on a test drive through some common avenues of gaining knowledge, to see how Explanationism handles. Let's start with perception, using vision as a paradigm example. Consider an ordinary case, for example the belief that there's a computer in front of you, formed on the basis of visual perception under normal conditions. Your believing this is straightforwardly explained by the fact that there *is* a computer in front of you. It's true that there's more we could (and perhaps even should) add to the explanation – facts about the lighting, distance, functioning of your visual system, your visual experience, etc. – but the truth of your belief would remain a crucial part of this explanation. If we tried to end the explanation like so, "You believe there's a computer before you because it *looks* that way," we'd be ending the explanation on a cliffhanger, as it were. We should wonder whether things look that way because they are that way, or for some other reason.¹⁷ A satisfying, complete explanation in this case, then, will include the fact that there is a computer before you, which is the truth of the relevant belief. So, Explanationism tells us this is a case of knowledge, as it should.

Knowledge by way of memory can receive a similar treatment. Memory is, after all, something like delayed or displaced perception. When you remember that you had granola for breakfast, you hold that belief on the basis of a familiar kind of memorial experience. And, if this really is a case of knowledge, you're having this memory because you actually had the *experience* of eating granola for breakfast, and you had the experience of eating granola because you really *did* eat granola. This last link is required for a complete explanation. So, again, though there is much we might add to the explanation of your belief in this case to make it more informative, the truth of your belief will play a crucial role. And so Explanationism again gives the right result.¹⁸

We believe that Explanationism also handles introspection rather easily. Even if you don't think that introspection operates via representations, as perception does, probably we can agree that some sort of spectating occurs during *introspection*. And so a story can be told much like the story above about ordinary visual perception. When you believe, for example, that you're desiring coffee, or that you're having an experience as of your hands sweating from caffeine withdrawal, you believe this because of what you "see" when you attend to the contents of your own mind (or what you "*just see*," if you insist, rightly, that introspection is direct and non-representational). And you (just) see what you do because

¹⁷It would be like this: A man of modest means presents his wife with a letter he's just opened containing a \$50,000 check. "Where did this come from?" she excitedly and reasonably asks. "The mailbox," he truthfully replies. This is not a complete explanation. So too if we offer *how things look* to explain *how we believe things are*. Goldman (1988: 49–50) suggests that, in a case of believing there's a typewriter before one on the basis of perception, "the retinal image ... does not screen out the typewriter in the explanation for my belief," because "that the retinal image makes my belief more probable can itself be explained by the fact that the retinal image is much more probable given the typewriter."

¹⁸As for knowledge of events that happened in the past but which you don't remember, Alan Goldman (1984: 103) says something we consider very sensible: "On one variation a past fact explains subsequent perceptions and beliefs that, together with communication, explain a present belief about that fact. On a second variation a past fact leaves material traces (that it explains), which constitute present evidence that, together with beliefs about the explanatory relations themselves, explain present belief about the past fact. Knowledge about the future derives from similar chains as well. Presently observed facts can help to explain some belief about the future, and these facts can also explain via causal chains the coming to pass of the future events now believed forthcoming on the basis of this evidence."

it's there. You really do have the desire, and you really do have the experience as of your hands sweating. So, here too we find that truth plays a crucial role in the explanation of your belief. And Explanationism gets the right result again.

Now let's consider rational intuition. Here, things might look more challenging for Explanationism. Roger White (2010: 582–3) puts the concern like so: "It is hard to see how necessary truths, especially of a philosophical sort can explain contingent facts like why I believe as I do." But we believe Explanationism has an answer. In the ordinary cases of knowledge of necessary truths, rational intuition puts us in a position to "just see" the truth of the proposition. Emphasis on "just," since, in this case, no premises or inferences are needed. It's akin to the case in which we "just see" that there's a computer in front of us, non-metaphorically, using our eyes. (Though, in the case of perception, there is an appearance/reality distinction.) To the degree that it's plausible that, in the ordinary case, you believe there's a computer before you because there is a computer before you, it's also plausible that, for example, you believe that $2 + 2 = 4$ because, lo, this fact is there, before your mind's eye, as it were. Rational intuition has positioned you to appreciate this truth. The fact that $2 + 2 = 4$ is necessarily true is no barrier to its contingent presence before your mind playing a crucial role in the explanation of your belief.¹⁹

¹⁹An anonymous referee helpfully expressed this worry: the fact that $2+2=4$ is causally inert, and so the fact is literally not doing any work here. If there were no numbers, or if the number 2 bore the plus relation to itself and to 7, we still would have believed as we do. In response, we deny that if there were no numbers, or if $2+2=7$, we still would have believed that $2+2=4$. We believe that rational intuition puts us in direct contact with the fact that $2+2=4$. Rational intuition, on our view, does not operate via any sort of report, experience, representation, or appearance (e.g. what some call an "intellectual seeming," cf. Bealer 1998), and so does not require that $2+2=4$ have any causal power. On our view, there is no appearance/reality distinction when it comes to rational intuition: our grasp or apprehension of the relevant fact is unmediated, a direct acquaintance. So, there is no opportunity for hallucination or illusion when it comes to rational intuition. And, therefore, if it were false that $2+2=4$, we would have noticed that. Followers of physicalism may find it difficult to explain the direct apprehension of truth afforded to us by rational intuition. But surely rational intuition works in some way (otherwise, a nasty skepticism and self-defeat looms). And if it's difficult for physicalism to explain this, so much the worse for physicalism, we say. But even if you go in for quasi-experiential intellectual seemings, it's still plausible that, and if it were false that $2+2=4$, we would have noticed that, just as, in the ordinary case of seeing an apple immediately before you, if the apple weren't there, you wouldn't have believed that it was, because you would have noticed its absence. See also Nicholas Sturgeon's (1984: 65) response to a challenge to moral knowledge from Gilbert Harman. There, Sturgeon says, "if a particular assumption is completely irrelevant to the explanation of a certain fact, then the fact would have obtained, and we could have explained it just as well, even if the assumption had been false." Though see Joel Pust's (2000: 240–3) criticisms.

Of course, talk of rational intuition often raises skeptical concerns about its fallibility, especially in light of apparent disagreement, both of the intrapersonal variety (in the form of paradoxes), and the interpersonal variety (in the form of disagreeing intuitions among philosophers). This concern is perhaps more acute when rational intuition is said to be direct, unmediated, etc. How could such a power err? Readers wishing to explore this issue may consult Bealer (1998, 2004), BonJour (1998), and, more recently, Christensen (2009), Killoren (2010), Bengson (2015), Chudnoff (2016), and Mogensen (2017). We agree with Pust (2017), though, that "it can be difficult to determine if another's failure to have an intuition that p is epistemically significant, as they may have yet to really grasp or consider the precise proposition at issue." We think it a promising strategy worth exploring to diagnose alleged cases of disagreement over intuitions (as well as alleged cases of errors in intuition) as cases of proposition confusion in a Kripkean vein (cf. Kripke 1980: 143). Just as Kripke plausibly explained away the alleged intuitiveness of "It's possibly false that Hesperus = Phosphorus" as a *mistaken description* of the (true) intuition that "it's possibly false that the evening star = the morning star," one might similarly explain away alleged cases of disagreeing intuitions or intuitive error as actually involving a mistaken *description* of the relevant (and true) intuition. Exploring this strategy fully is beyond the scope of this paper.

In the next section, we'll turn to more difficult cases. These cases are so difficult that we call this section ...

4. Some common objections

We consider first the case of abductive knowledge. How could, for example, your belief that *all* emeralds are green be explained by the relevant fact, when you've seen only *a few* emeralds? Here, we believe that the field has already been plowed by previous philosophers. Carrie Jenkins (2006: 159), drawing upon Goldman's prior work (1984: 103), offers this explanation of abductive knowledge:

[I]n genuinely knowledge-producing cases of inference to the best explanation, where *p* explains *q*, *q* explains [that the agent believes that *q*], and [that the agent believes that *q*] explains [that the agent believes that *p*], and all these explanatory links are standard, then one can collapse the explanatory chain and simply say that *p* explains [that the agent believes that *p*], even when talking to an outsider. For example, if I believe that all emeralds are green on the grounds of an inference from known instances of green emeralds, it seems reasonable to say that the fact that all emeralds are green is what explains my belief that they are. For it is what explains the fact that all the emeralds I have seen so far have been green, and this is what explains my general belief.

You may have qualms about endorsing transitivity with regard to explanation (or grounding), but notice that this account of abductive knowledge doesn't require anything so strong as that. It's a limited, qualified claim that, *sometimes*, under the right conditions, explanatory chains collapse, and that this is one of those times with those right conditions. If Explanationism is true, then those "right conditions" will have to do with difference-making, and which factors are strictly required for successful explanation.

We move then to the truly difficult case of deductive knowledge, where previous work is not so definitive. The *prima facie* worry is this: when I come to know something by way of deduction – when I run an argument and draw a conclusion – don't the *premises* explain my belief in the conclusion, rather than the truth of the conclusion itself? Goldman (1984: 104–5) makes the following attempt to incorporate deductive knowledge into Explanationism:

I know that a mouse is three feet from a four-foot flagpole with an owl on top, and I deduce that the mouse is five feet from the owl. I then know this fact, although I have no explanation for the mouse's being five feet from the owl, although this fact does not explain my belief. Appeal to explanatory chains, however, allows us to assimilate this case to others within the scope of the account. The relevant theorem of geometry is the common element in the relevant chains here. First, the truth of that theorem explains why, *given that the mouse is three feet from a four-foot pole*, it is five feet from the top, although it does not explain that fact simpliciter. Someone who accepts the other two measurements but does not know why, given them, the mouse is five feet from the owl, has only to be explained the theorem of geometry. Second, the truth of that theorem helps to explain why mathematicians believe it to be true, which ultimately explains my belief in it. My belief in the theorem in turn helps to explain my belief about the distance of the mouse from the owl. Thus what helps to explain my belief in the fact in question also helps to explain that fact, given the appropriate background conditions, which here include the other distances. The case fits our account.

Goldman seems to use the following premises:

1. The Pythagorean theorem explains why, given that a mouse is three feet from a four-foot-tall flagpole with an owl on top, then the mouse is five feet from the owl.
2. The Pythagorean theorem explains why mathematicians believe it.
3. The fact that mathematicians believe the Pythagorean theorem explains why I believe it.
4. My belief in the Pythagorean theorem helps explain my belief that the mouse is five feet from the owl.

And Goldman wishes to derive this conclusion:

5. So, the Pythagorean theorem helps explain my belief that the mouse is five feet from the owl, and also explains why, given that the mouse is three feet from a four foot flagpole with an owl on top, the mouse is five feet from the owl.

If the explanatory chains collapse in the appropriate way, this conclusion does seem to follow: *the Pythagorean Theorem* helps explain my belief, and helps explain why my belief is true. But what's missing in all this – we think, and Goldman seems to acknowledge – is any account of how *the fact that the owl is five feet from the mouse* explains my belief that it is. Goldman shows, at most, how the Pythagorean Theorem helps explain my belief and also its truth (given certain conditions); but he doesn't show how the truth of my belief explains why I believe it. And we took that to be the primary desideratum of any Explanationist account of this case. Goldman (1988: 35) complicates his analysis to accommodate difficult cases, allowing that the analysis is satisfied “if that which explains a belief also ultimately explains or is explained by the fact to which the belief refers.” But, ideally, we'd like an Explanationism account of deductive knowledge that stays true to the original spirit of the analysis. So, we must look elsewhere.²⁰

Carrie Jenkins (2006: 160–1) attempts to improve on Goldman, and considers a case of Brian coming to know something by running the following argument:

6. If Neil comes to the party, he will get drunk.
7. Neil comes to the party.
8. So, he will get drunk

Brian knows (6) and (7), and believes (8) as a result of this argument. Does he know it, on Explanationism? Jenkins says “yes,” and here's her reasoning:

It's only because Neil will get drunk that the two premises from which Brian reasons are true together; that is to say, it is only because Neil will get drunk that Brian's argument is sound ... [A]fter noting that competent reasoners like Brian will (generally) only use sound arguments, we can explain the fact that Brian reasons as he does by citing the soundness of Brian's argument ... Finally, we note that, uncontroversially, Brian's reasoning as he did is what explains his belief that Neil will get drunk. So now we know that the fact that Neil will get drunk explains why Brian's argument is sound, which explains

²⁰An anonymous referee helpfully raised another concern, which you may well share: It is far from clear that the Pythagorean Theorem construed as a proposition about pure Euclidean space can explain anything in the non-platonic realm of matter. We take this to be another reason to leave Goldman's proposal behind, and look elsewhere for an Explanationist account of deductive knowledge.

Brian's reasoning as he did, which in turn explains why Brian believes that Neil will get drunk.

The first sentence is a little murky to us. Jenkins seems to be saying that (6) and (7) are true together, and the argument is valid, *because (8) is true*. But, obviously, an argument can be unsound and still have a true conclusion. (8) is no explanation of why "the two premises ... are true together." If anything, you might have thought, it's the other way around: (6) and (7) together explain the truth of (8).

We're also puzzled by Jenkins' claim that *the fact that this argument is sound* explains why Brian, a competent reasoner, used it. After all, there are lots of sound arguments that Brian didn't use. Imagine before you a sea of sound arguments that Brian might have used. He happened to use *this* one and not any other, and the fact that *it's sound* won't help us understand why, since they're all sound.

We'd like to improve on the efforts from Goldman and Jenkins, like so. Deductive arguments are a means by which we come to metaphorically "see" the truth of some proposition. Brian doesn't believe (8) because he believes (6), (7), and that (6) and (7) entail (8). Rather, those beliefs get him in a position to see the truth of (8), and *that* is why he believes (8). It's like an ordinary case of non-metaphorical vision. Vision is fallible, since it might be a hallucination or illusion. But, so long as your visual system is functioning well, it puts you in a position to see objects before you, and if one such is a chair, you believe there's a chair there *because* there's a chair there. The fact that there's a chair there is a crucial part of the explanation of why you believe there is. Similarly, when you see the premises are true and the inference valid, deduction helps you see the conclusion. And that's because, in a sound argument, the conclusion is already "there," so to speak, in the premises. When one appreciates a sound argument, one is not merely seeing the truth of the premises and the validity of the inference; one also sees the truth of the conclusion thereby. In that case, you believe the conclusion because it's true. This is fallible, when the component beliefs or premises might be false or otherwise unknown. But, when they're true and you know it, then deduction positions you to see the truth of the conclusion.

So, ultimately, Brian does believe (8) because it's true. His knowledge of the premises and validity helped him appreciate that (8) is true. It's like putting the pieces of a puzzle together, and then seeing the picture. In this analogy, the premises and inference are the pieces, and the conclusion is the picture. Now, when there's a completed puzzle of a pangolin in front of you, and you believe there's a picture of a pangolin there, a crucial part of the explanation of why you believe there's a picture of a pangolin there is because there's a picture of a pangolin there. In a similar way, when Brian knows those premises together and grasps the entailment, he's in position to see the truth of the conclusion. The conclusion is now "part of the picture," as it were, formed by the true premises and valid inference, and is thereby directly appreciable. And so it's correct that: Brian believes the conclusion because it's true. We may be misled into thinking otherwise by the way we standardly formalize arguments, with conclusions on a separate line from the premises. If you worry this may be happening to you, remind yourself that, with a valid argument, the conclusion is already "there" in the premises. When we represent valid Aristotelian syllogisms using Venn diagrams, for example, after we draw the premises, we needn't add anything further to represent the conclusion: it's already there. So, by appreciating the truth of the premises, we are in a position to appreciate the truth of the conclusion. And, if we do, it can be truly said that we believe the conclusion because it's true. We think this is a satisfactory Explanationist account of deductive knowledge, which can easily be adapted to other cases of deductive knowledge.

We'll turn now to inductive knowledge, especially of the future. On its face, such knowledge is a real puzzle for Explanationism: Absent some tricky backwards causation, how could your *current* belief that the sun will rise tomorrow be explained by the fact that the sun will rise *tomorrow*? Isn't it rather your current evidence that explains your current belief? Roger White (2010: 582–3) puts the concern like so: "I'm justified in thinking the sun will rise tomorrow. But assuming it does, this fact plays no role in explaining why I thought that it would." To solve this puzzle, we think we can combine the abductive and deductive solutions above. First, we'll use the explanatory-chain-collapse strategy from Goldman and Jenkins, like so:

9. You *believe* the sun rises every day because you've *seen* it rise many times in the past. (SEEN explains BELIEVE)
10. The *fact* that the sun rises every day explains why you've *seen* it rise many times in the past. (FACT explains SEEN)
11. So, the *fact* that the sun rises every day explains why you *believe* that the sun rises every day. (So, FACT explains BELIEVE)

If every link in this chain is required for a complete explanation, then, according to Explanationism, you know that the sun rises every day. Now, if you competently deduce that the sun will rise tomorrow from your knowledge that the sun rises every day, it will be the case that: the fact that the sun will rise tomorrow explains your belief that the sun will rise tomorrow. The known premise together with your grasp of the implication puts you in a position to appreciate the truth of the conclusion. In that case, the truth of the conclusion explains why you believe that it's true. And that's how inductive knowledge works on Explanationism, we suggest.²¹

File the following objection under "Miscellaneous":

The Case of the Epistemically Serendipitous Lesion. Suppose K suffers from a serious abnormality – a brain lesion, let's say. This lesion wreaks havoc with K's noetic structure, causing him to believe a variety of propositions, most of which are wildly false. It also causes him to believe, however, that he is suffering from a brain lesion. K has no evidence at all that he is abnormal in this way, thinking of his unusual beliefs as resulting from an engagingly original turn of mind ... surely K does not know that he is suffering from a brain lesion. He has no evidence of any kind – sensory, memory, introspective, whatever – that he has such a lesion;

²¹This same strategy may be employed to account for more mundane cases of knowledge of the future, for example your knowing that you will get a raise next year, when your boss tells you so. So long as your boss is a reliable testifier, i.e. so long as there's a law-like connection between her testimony and the truth, it's plausible that a universal generalization like "if my boss asserts p, then p is true" explains (the positive track record of truth-telling you've observed from your boss, which explains) why you believe this is so. So, on Explanationism, you know that this is so. And then, adding your knowledge of your boss's assertion about your raise, you deduce that you will get a raise next year, in such a way that the truth of this conclusion explains your belief in it, as described above. So, it's knowledge, according to Explanationism. One wrinkle: probably your boss is not so reliable that there's genuine *entailment* between her testimony and the truth. But, like much of our ordinary thought and reasoning, this indicative conditional may be non-monotonic, and this a case of "defeasible reasoning." That is, there is likely a qualification commonly left implicit, to the effect that: *unless something strange or anomalous happens*, if my boss asserts p, then p is true. If so, then in this case, what one really comes to know about the future, according to Explanationism, is the more modest claim that: *unless something strange or anomalous happens*, I will get a raise next year. But, on reflection, that sits well with us. We thank Bar Luzon for encouraging us to think about cases like this.

his holding this belief is, from a cognitive point of view, no more than a lucky (or unlucky) accident. (Plantinga 1993: 195)

The allegation here, adapting the case to target Explanationism, is that the fact that K has a brain lesion plays a crucial role in K's belief that he has a brain lesion, and so Explanationism must say K knows he has a brain lesion, which is absurd. But we believe that this case can be handled in a similar way to Faraci's case of Eula's random number generator, discussed above. Either the brain lesion is randomly producing beliefs, or it isn't. If it is randomly producing beliefs, then although K can't know that he has a brain lesion on this basis, Explanationism doesn't have to say that he does. And that's because, if the brain lesion is operating genuinely randomly, then its operation does not explain why K ends up with the belief that he has a brain lesion. We may have an explanation of why K believes *something or other*, but no explanation of the particular belief that he has a brain lesion. That's the nature of a truly random process: there is no demystifying explanation of its operation. The story is simply this: K had a brain lesion, and then – pop! – K ends up believing that he has a brain lesion. And there's nothing demystifying or explanatory about that pop.²²

On the other hand, suppose the brain lesion is *not* operating randomly when it tells K that he has a brain lesion. Rather, the brain lesion tells K he has a brain lesion because he *does* have a brain lesion. In a simple version of this case, when the brain lesion produces *only* the belief that there's a brain lesion, because it *is* a brain lesion, it looks like there is the right sort of explanatory connection for Explanationism to say that this is a case of knowledge. But – and now buckle up, Internalists – it also looks like that's the right result: this *is* knowledge. The brain lesion is a good way for K to learn that he has a brain lesion: the brain lesion “tells” K that he has a brain lesion *because it's true*. And what more could you ask from a brain lesion? This is merely a case of very unusual – though highly reliable – testimony.

But the case is a little more complicated as Plantinga describes it. As Plantinga has it, the brain lesion “wreaks havoc with K's noetic structure, causing him to believe a variety of propositions, most of which are wildly false.” In this more complicated version of the case, although the brain lesion tells K that he has a brain lesion not randomly but rather because it's true, the brain lesion also tells K many false things besides. And now it starts looking less like a case of knowledge. In fact, it starts looking more like Fake Barn Country (cf. Alvin Goldman 1976). In Fake Barn Country, one's eyes happen to alight on a real barn amid a sea of facades. In Chaotic Brain Lesion Country, one happens to receive a true belief from a brain lesion selecting from a sea of falsehood.

²²A similar response can be given to an objection involving an ordinary case of defeat. (Thanks to Bob Beddor for suggesting such cases.) **Objection:** Suppose Barry knows there's a red surface before him, on the basis of vision: he believes it because it seems that way, and it seems that way because it is that way. But now Bruno, a trusted friend and lighting expert, tells Barry there is unusual red lighting in the area. Even if Bruno is wrong about the lighting, surely it would be irrational for Barry to continue believing the surface before him is red, given Bruno's testimony. But it's possible that Barry does go on believing there's a red surface before him, because it looks that way, and it looks that way because it is that way. In that case, Explanationism must say this continues to be knowledge. But, intuitively, it isn't. So, Explanationism is false. **Reply:** We agree that, in this case, *before* Bruno comes along, Explanationism says, rightly, that Barry knows. But now, *after* he hears and understands Bruno's testimony, it seems something has gone awry with poor, irrational Barry. And whatever it is that's gone awry will enter into the explanation of why he now continues to believe there's a red surface, despite Bruno's testimony. Perhaps he's biased toward beliefs about red surfaces, or beliefs tend to “stick” in his mind, despite contrary evidence, or he's on drugs, or whatever. We suspect that, whatever the explanation turns out to be, the present truth of his belief will no longer figure crucially into it. And, if so, Explanationism would give the right result: Barry no longer knows there's a red surface before him, after hearing Bruno's testimony.

What can the Explanationist say about Fake-Barn-style cases? Perhaps this. In the ordinary case of seeing a barn, you believe there's a barn before you because it looks like there's a barn before you, and – we continue, for a complete explanation – it looks like there's a barn before you because there is a barn before you. Explanationism gives its imprimatur as it should: this is knowledge. But, as the fake barns proliferate, the second link in the explanatory chain becomes less plausible. Even if your eyes happen to fall upon a real barn in a forest of fakes, we might begin to think it false that *it looks like there's a barn before you because there is a barn before you*. As the barn facades proliferate, a rival explanation looms into view: that it looks like there's a barn before you because you're in a region full of structures that look like barns. In other words, it becomes plausible to say that you believe there's a barn before you because it looks like there's a barn before you, and it looks like there's a barn before you *because you're in Fake Barn Country*. (In Fake Barn Country, it's common to be appeared to barn-ly, and more common the more barn facades there are. In that case, we can – and perhaps we *should* – explain your belief by citing your presence among a multitude of things that look like barns. Given all those nearby structures, it seems *that particular* structure you happened to see – and the fact that it happened to be a real barn rather than a facade – plays no crucial role in the explanation of your belief.²³ Compare: While driving, a contractor's truck drops a large number of sharp objects – nails and screws – all over the road. The car following behind the truck gets a flat tire because it ran through this mess. Meditate for a moment on the suitability of this explanation. Now, while it may be true that one particular sharp object – a nail, let's say – punctured the tire, it's unnecessary to cite that particular object, or the fact that it was a nail rather than a screw, in order to explain the puncture, given all the other sharp objects nearby that nail, poised to puncture the tire in its place. All that figures crucially into the explanation of the punctured tire is the prevalence of these sharp objects. Perhaps the same goes with Fake Barn Country. What figures crucially into the explanation of your belief that there's a barn is the prevalence of objects that look like barns.) And the fact that there's a real barn before you plays no crucial role in *that* explanation. So, on Explanationism, as the barns proliferate, it becomes less and less clear that you have knowledge. At some point, perhaps, the knowledge disappears.²⁴

Something similar could be said about the more complicated version of Plantinga's serendipitous brain lesion. As the brain lesion produces more false-beliefs (which seem true from the inside), the following explanation of K's belief becomes more plausible: K believes that he has a brain lesion because it seems true to him, and it seems true to him because he's the victim of a brain lesion that produces many true-seeming beliefs. And, in this case, the truth of K's belief plays no crucial role in the explanation of his believing it.²⁵

²³Goldman (1988: 44) agrees: "In such cases ... the proper explanation for the belief appeals to the broader context of the perceiver's being in the vicinity of all these look-alike objects, any of which would produce the belief in question."

²⁴In recent years, it's become more common among epistemologists to express reservations about the claim that one doesn't know in Fake Barn Country. Perhaps this is because it's not clear whether and to what extent it's true that the alternative, knowledge-precluding explanation mentioned in this paragraph really does rise to prominence, and become the explanation of the subject's belief. We consider it a virtue of Explanationism that its verdicts on Fake Barn Country cases seem to be as indeterminate as our intuitions about whether there's knowledge in such cases.

²⁵A similar thought might apply in the case of statistical inference, on Explanationism. One might know that a lottery ticket *probably* will lose, on the basis of statistical evidence. But, since the fact that the ticket *will* lose doesn't figure into this evidence, one could not be said to know that the ticket will lose on this basis. Goldman (1988: 52–3) explains it thus: the future fact that my ticket will not win does not raise the antecedent probability of my holding this belief, "which is explained rather by my knowledge of the

Finally, let us close with a few words about Linda Zagzebski's (1994) Paradox of the Analysis of Knowledge, and how Explanationism answers it. Zagzebski argues that any view which allows "a small degree of independence" between the justification and truth components of knowledge will suffer from Gettier-style counterexamples. "As long as the truth is never assured by the conditions which make the state justified," she says, "there will be situations in which a false belief is justified" (Zagzebski 1994: 73). And, "since justification does not guarantee truth, it is possible for there to be a break in the connection between justification and truth, but for that connection to be regained by chance" (Zagzebski 1994: 67). Those are Gettier cases.²⁶ So the idea is that epistemologists are grappling with the following paradox:

- (A) Knowledge = TB + X
- (B) X does not entail TB. It's possible to have X without TB. False but justified beliefs, for example.
- (C) At least some cases of X without TB can be "restored" to TB + X, without there being knowledge.

If you think that (A) there's an analysis of knowledge to be had by adding a missing necessary condition to truth and belief, and also that (B) this missing condition does not entail truth, and yet also that (C) any such analysis will allow for Gettier cases, then you've contradicted yourself. That's the problem.

Explanationism's solution is to avoid Gettier cases by denying (B).²⁷ On Explanationism, S knows that p iff S believes that p because it's true. Clearly, the right biconditional entails truth. So, Explanationism is committed to denying (B). Zagzebski (1994: 72) says "few philosophers have supported this view," since it means one must "give up the independence between the justification condition and the truth condition." But, justification (in the Internalist sense, at least) doesn't figure into Explanationism's analysis. So, Explanationists are free to agree that there *can* be justified but false beliefs, in the sense of "justification" that Zagzebski seems to have in mind. The Explanationist could add that, perhaps, (Internalist) justification is just our best attempt to know *whether* the Explanationist condition is met. To prove that

number of tickets and my ability to infer inductively on that basis." This may help solve the puzzle of why we're reluctant to convict defendants on the basis of purely statistical evidence (cf. Littlejohn 2017), but why we're willing to convict on the basis of e.g. fallible eyewitness testimony (cf. Papineau *Forthcoming*). On Explanationism, eyewitness testimony can generate knowledge by providing an explanatory connection to the truth, even if we're uncertain *that* this connection holds. Statistical inference, on the other hand, provides at best an explanatory connection to – and, so knowledge of – only *probable* truth. And evidently our judgments about knowledge reveal that we much prefer a likely link to truth, rather than a link to likely truth.

²⁶See also Scott Sturgeon (1993) and Trenton Merricks (1995).

²⁷That is, on Explanationism, warrant entails truth. And this is plausible; cf. Andrew Moon (2012). This is an explanation of why, in principle, there can be no Gettier cases on Explanationism. But, perhaps to confirm this result, you might reflect on how Explanationism would handle a paradigm Gettier case. For example, Russell's Stopped Clock. Suppose your favorite and generally reliable clock has stopped this morning at 8:22 am. But, by the sheerest coincidence, you happen to look at it this evening at 8:22 pm. Since the clock reads "8:22," you form the true belief that it's 8:22 pm. And, given your history with this generally reliable clock, you have a high degree of (Internalist) justification. But, no knowledge. Why is that, on Explanationism? Well, you do believe that it's 8:22 because the clock reads "8:22." But, crucially, it's false that the clock reads "8:22" because it *is* 8:22. It reads "8:22" because it *was* 8:22, when the clock stopped. And, so, the truth of your belief does not figure crucially into the explanation of why you hold the belief. So, Explanationism gives the right result here. And similar treatments could be given of other Gettier cases.

it is. To prove *that* we know. But, knowing does not require knowing that one knows, or being able to prove that one knows. And that's because, in the final analysis, knowing is simply believing something because it's true.²⁸

References

- Achinstein P. (1983). *The Nature of Explanation*. New York, NY: Oxford University Press.
- Alspector-Kelly M. (2006). 'Knowledge Externalism.' *Pacific Philosophical Quarterly* 87, 289–300.
- Bealer G. (1998). 'Intuition and the Autonomy of Philosophy.' In M. DePaul and W. Ramsey (eds), *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*, 201–40. Lanham, MD: Rowman and Littlefield.
- Bealer G. (2004). 'The Origins of Modal Error.' *Dialectica* 58(1), 11–42.
- Becker K. and Black T. (2012). *The Sensitivity Principle in Epistemology*. Cambridge: Cambridge University Press.
- Beddor B and Pavese C. (Forthcoming). 'Modal Virtue Epistemology.' *Philosophy and Phenomenological Research*.
- Bengson J. (2015). 'Grasping the Third Realm.' In T. Szabo Gendler and J. Hawthorne (eds), *Oxford Studies in Epistemology*, Volume 5, pp. 1–38. New York, NY: Oxford University Press.
- Black T. (2008). 'Defending a Sensitive neo-Moorean Invariantism.' In V.F. Hendricks and D. Pritchard (eds), *New Waves in Epistemology*, pp. 8–27. Houndmills: Palgrave Macmillan.
- Black T. and Murphy P. (2007). 'In Defense of Sensitivity.' *Synthese* 154(1), 53–71.
- Bogardus T. (2014). 'Knowledge under Threat.' *Philosophy and Phenomenological Research* 88(2), 289–313.
- Bogardus T. and Marxen C. (2014). 'Yes, Safety is in Danger.' *Philosophia* 42(2), 321–34.
- BonJour L. (1998). *In Defense of Pure Reason*. Cambridge: Cambridge University Press.
- Christensen D. (2009). 'Disagreement as Evidence: The Epistemology of Controversy.' *Philosophy Compass* 4(5), 756–67.
- Chudnoff E. (2016). *Intuition*. New York, NY: Oxford University Press.
- Clarke-Doane J. (2012). 'Morality and Mathematics: The Evolutionary Challenge.' *Ethics* 122 (2), 313–40.
- Clarke-Doane J. (2014). 'Moral Epistemology: The Mathematics Analogy.' *Noûs* 48(2), 238–55.
- Clarke-Doane J. (2015). 'Justification and Explanation in Mathematics and Morality.' In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*, Volume 1. Oxford: Oxford University Press.
- Clarke-Doane J. (2016). 'What Is the Benacerraf Problem?' In F. Pataut (ed.), *New Perspectives on the Philosophy of Paul Benacerraf: Truth, Objects, Infinity*, pp. 17–43. New York, NY: Springer.
- Clarke-Doane J. (Forthcoming). 'Debunking Arguments: Mathematics, Logic, and Modal Security.' In R. Richards and M. Ruse (eds), *Cambridge Handbook of Evolutionary Ethics*. Cambridge: Cambridge University Press.
- Clarke-Doane J. and Baras D. (Forthcoming). 'Modal Security.' *Philosophy and Phenomenological Research*.
- DeRose K. (1995). 'Solving the Skeptical Problem.' *Philosophical Review* 104(1), 1–52.
- Dretske F. (1970). 'Epistemic Operators.' *Journal of Philosophy* 67, 1007–23.
- Dretske F. (1971). 'Conclusive Reasons.' *Australasian Journal of Philosophy* 49(1), 1–22.
- Faraci D. (2019). 'Groundwork for an Explanationist Account of Epistemic Coincidence.' *Philosopher's Imprint* 19(4), 1–26.
- Frankfurt H. (1969). 'Alternate Possibilities and Moral Responsibility.' *Journal of Philosophy* 66(23), 829–39.
- Goldman, A. (1976). 'Discrimination and Perceptual Knowledge.' *Journal of Philosophy* 73, 771–91.
- Goldman A. (1984). 'An Explanatory Analysis of Knowledge.' *American Philosophical Quarterly* 21, 101–8.
- Goldman A. (1988). *Empirical Knowledge*. Berkeley, CA: University of California Press.
- Hawthorne J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Ichikawa J.J. (2011). 'Quantifiers, Knowledge, and Counterfactuals.' *Philosophy and Phenomenological Research* 82, 287–313.
- Jenkins C. (2006). 'Knowledge and Explanation.' *Canadian Journal of Philosophy* 36(2), 137–64.

²⁸For helpful conversations and comments, thanks to Dan Baras, Bob Beddor, Julien Dutant, Gregory Gaboardi, Dan Greco, Alex Grzankowski, Sydney Keough, Dan Korman, Clayton Littlejohn, Bar Luzon, Joel Mann, Andrew Moon, Barry Smith, Jacob Zimbelman, and audiences at the Institute of Philosophy at the University of London, as well as at St. Norbert College.

- Killoren D.** (2010). 'Moral Intuitions, Reliability, and Disagreement.' *Journal of Ethics and Social Philosophy* 4(1), 1–35.
- Korman D.** (2019). 'Debunking Arguments.' *Philosophy Compass* 14(12), 1–17.
- Korman D. and Locke D.** (Forthcoming). 'Against Minimalist Responses to Moral Debunking Arguments.' In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics, Volume 15*. Oxford: Oxford University Press.
- Kripke S.** (1980). *Naming and Necessity*. Cambridge, MA: Harvard University Press.
- Littlejohn C.** (2017). 'Truth, Knowledge, and the Standard of Proof in Criminal Law.' *Synthese* 1–34. <https://doi.org/10.1007/s11229-017-1608-4>.
- Luper-Foy S.** (1984). 'The Epistemic Predicament: Knowledge, Nozickian Tracking, and Scepticism.' *Australasian Journal of Philosophy* 62(1), 26–49.
- Merricks T.** (1995). 'Warrant Entails Truth.' *Philosophy and Phenomenological Research* 55(4), 841–55.
- Mogensen A.** (2017). 'Disagreements in Moral Intuition as Defeaters.' *Philosophical Quarterly* 67(267), 282–302.
- Moon A.** (2012). 'Warrant Does Entail Truth.' *Synthese* 184(3), 287–97.
- Neta R.** (2002). 'S Knows that P.' *Noûs* 36, 663–81.
- Nozick R.** (1981). *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Papineau D.** (Forthcoming). 'The Disvalue of Knowledge.' *Synthese*.
- Plantinga A.** (1993). *Warrant: The Current Debate*. Oxford: Oxford University Press.
- Pritchard D.** (2005). *Epistemic Luck*. Oxford: Oxford University Press.
- Pritchard D.** (2007). 'Anti-Luck Epistemology.' *Synthese* 158(3), 277–98.
- Pritchard D.** (2009). 'Safety-Based Epistemology: Whither Now?' *Journal of Philosophical Research* 34, 33–45.
- Pust J.** (2000). 'Against Explanationist Skepticism regarding Philosophical Intuitions.' *Philosophical Studies* 106(3), 227–58.
- Pust J.** (2017). 'Intuition.' In E.N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2019 Edition). <https://plato.stanford.edu/archives/sum2019/entries/intuition/>.
- Rieber S.** (1998). 'Skepticism and Contrastive Explanation.' *Noûs* 32, 189–204.
- Roush S.** (2005). *Tracking Truth: Knowledge, Evidence and Science*. Oxford: Oxford University Press.
- Sainsbury R.M.** (1997). 'Easy Possibilities.' *Philosophy and Phenomenological Research* 57, 907–19.
- Sosa E.** (1999a). 'How to Defeat Opposition to Moore.' In J. Tomberlin (ed.), *Philosophical Perspectives 13: Epistemology*, pp. 141–54. Oxford: Blackwell.
- Sosa E.** (1999b). 'How Must Knowledge be Modally Related to What is Known?' *Philosophical Topics* 26 (1/2), 373–384.
- Sosa E.** (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, Volume I. Oxford: Oxford University Press.
- Sosa E.** (2009). *Reflective Knowledge: Apt Belief and Reflective Knowledge*, Volume II. Oxford: Oxford University Press.
- Strevens M.** (2011). *Depth*. Cambridge, MA: Harvard University Press.
- Sturgeon N.** (1984). 'Moral Explanations.' In D. Copp and D. Zimmerman (eds), *Morality, Reason and Truth*. Totowa, NJ: Rowman & Allenheld.
- Sturgeon S.** (1993). 'The Gettier Problem.' *Analysis* 53(3), 156–64.
- Van Fraassen B.** (1980). *The Scientific Image*. Oxford: Oxford University Press.
- White R.** (2010). 'You Just Believe that Because...' *Philosophical Perspectives* 24, 573–615.
- Williamson T.** (2002). *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Zagzebski L.** (1994). 'The Inescapability of Gettier Problems.' *Philosophical Quarterly* 44(174), 65–73.

Tomas Bogardus is an Associate Professor of Philosophy at Pepperdine University. His previous work on Modalism and the safety condition includes "Knowledge under Threat" (*Philosophy and Phenomenological Research*, 2014) and "Yes, Safety is in Danger" (co-authored with Chad Marxen, *Philosophia*, 2014).

Will Perrin is a graduate student at Northern Illinois University. His main philosophical interests are metaethics and epistemology.