Out in the Open: Public Evidence and the Limits of Experience

Ted Poston*

Word Count: 9268

December 19, 2025

Abstract

Public evidence plays a central role in the justification of scientific theories—but does its importance extend beyond science, for instance, to political or religious belief? To address this question, we first need a clear account of public evidence. This paper develops such an account, characterizing public evidence as non-experiential evidence that meets the non-factive epistemic conditions for common knowledge. I argue that public evidence not only underwrites the justification of scientific theories but also constrains how experience can justify belief more generally. The result is a novel account of public evidence with broad applicability across domains of inquiry.

Epistemology has taken an experiential turn. While reason aspires to universality, experience is inherently local. This shift has fueled defenses of intuition,[1] appearances,[2] and religious experience[3] as legitimate sources of justification. This leads to a fragmented map of reasonable landscapes, where what is reasonable for a person who grew up in Babylon is different from one raised in Byzantium and different from one in Bethlehem. To take one instance, the turn from attempting to justify religious belief in terms of common evidence to

*I am grateful to Rodrigo Borges, Juan Comesaña, Marilie Coetsee, Kevin Dorst, Branden Fitelson, John Keller, Lorraine Keller, Tom Kelly, Jon Matheson, Kevin McCain, Tim Perrine, Daniel Rubio, Ernie Sosa and Mike Veber for their valuable feedback on earlier drafts. I also thank the audiences at the University of Wisconsin, Vanderbilt University, St. Joseph's University, the MIT Epistemology Reading Group, and the Southeastern Epistemology Conference for their insightful comments and suggestions.

[1]Bealer (1998).

[2]Huemer (2001).

[3]Plantinga (1967, 1983, 2000); Alston (1991).

conceiving of religious experience as the primary means by which religious belief is justified leads straightforwardly to the problem of religious diversity. If it is religious experience that justifies religious belief in contrast to reasoning about the public evidence, then there is little hope for a non-skeptical resolution for the justification of religious worldviews.[4]

It was not always this way. The early modern philosophers–specifically Descartes and Locke–had a conception of reason that was universal. Philosophy, on this view, aimed to articulate a standpoint accessible from anywhere where reason, working on a foundation of common facts, supports a common civic, political, and religious rationality. Rationality then would issue the same prescriptions in Babylon, Byzantium, and Bethlehem. We've learned that that universal ambition cannot be achieved because of the dubious theses associated with the Enlightenment model of rationality.[5] To mention three: the foundations of knowledge are theory neutral, the superstructure is entirely supported by deduction and enumerative induction, and the sharp distinction between facts and values. In short, Enlightenment rationality rested on a picture of inquiry that sharply separated theory from observation and fact from value, and that modeled justification as proceeding by relatively simple routes from observation to theory.

Yet something of the Enlightenment's universal ambition endures in the thought that science exemplifies a model of epistemic success. Part of the explanation for this success is that scientific justification is secured by public evidence and eschews appeals to incommunicable insight. While the relationship between observation and theory in science is neither sharp nor straightforward, scientific practice remains a paradigm of justified belief precisely because it demands that justification be grounded in publicly accessible evidence.[6]

It's surprising, therefore, that there is little discussion of the nature of public evidence.[7] The task of this paper is to defend a conception of public evidence that applies more generally in epistemology, not just in the context of scientific justification. I argue that we should understand public evidence as non-experiential evidence that meets the non-factive epistemic conditions for common knowledge. The first three sections develop the concepts and arguments needed to formulate and defend that thesis. The final section argues that public evidence is an epistemic constraint on the way experience provides justification.

[4]Witness the great pumpkin objection to Plantinga's reformed epistemology. See Plantinga (1983).

[5]See, for example, Kuhn (1962, 1977).

[6]For a recent development of how the method of science generates significant knowledge of the world see Strevens (2020). Strevens's view is that the epistemic success of science lies in the methodological rule—what he terms "The Iron Rule" to resolve scientific debates through, and only through, empirical tests Strevens (2020, 93).

[7]The recent state of the art volume *The Routledge Handbook of the Philosophy of Evidence* has a short two page discussion of the "publicity of evidence" in the context of a discussion about internalist and externalist views of evidence Fratantonio (2024, 58-60). Bob Beddor and Finnur Dellsén have a recent paper where they discuss the public evidence norm in science (see Beddor and Dellsén (2025)).

1 Two faces of public evidence

Public evidence plays a central role in science. Timothy Williamson observes that "science depends on public evidence."[8] John Earman, in a note complaining about the irrelevance of externalist theories of knowledge to discussions of scientific justification, remarks that "because science is a community enterprise the only forms of justification that are scientifically relevant are those which are statable and open to public scrutiny."[9] Earlier in the twentieth century, Hempel insists that "all statements of empirical science be capable of test by reference to evidence which is public, i.e., evidence which can be secured by different observers and does not depend essentially on the observer."[10] Apart from the logical positivists' discussions over the nature of observation sentences in connection with the objectivity of science, there have not been substantive discussions on the nature of public evidence informed by recent developments in epistemology.[11]

1.1 Two criteria of adequacy: externality & limitation

Public evidence is often thought to play two seemingly incompatible roles. The first role is that public evidence is out there in the world, available to everyone to learn about. It is not a mental state like an appearance, a seeming, or an intuition. The evidence that a giant meteor caused the Cretaceous–Paleogene extinction sixty-six million years ago is *not* that appears to me to be so.[12] Rather the evidence includes a global iridium-rich layer, the Chicxulub crater in Mexico, shocked quartz and tektites, worldwide tsunami and fallout deposits, and a sudden mass extinction pattern in the fossil record. The evidence for this is public; it is out there in the world in sharp contrast to the way in which I may come to know whether I am feeling tired. Thus, we have the following criterion of adequacy.

Externality: Public evidence is external.

Externality is schematic, but it captures three related ideas. First, the public evidence consists of facts about the world. Second, the public evidence is accessible to all interested parties. Third, the public evidence is not a private "internal" mental state. An account of public evidence shouldn't prejudge the truth of these claims, but an adequate account should

[8] Williamson (2000, 185).

[9] Earman (1993, 37).

[10] Hempel (1952, 22).

[11] See Thomas Kelly's discussion of the role of public evidence in the positivists in section 4 of Kelly (2016). See also Kelly (2025) for some remarks on the contrast between public and private evidence.

[12] If someone had such an appearance, you might well take that as evidence that there is something defective about their cognition.

explain the prominence of these features in discussions of public evidence. These three ideas triangulate the phenomenon of interest. Public evidence is the kind of thing that can be successfully demonstrated such that multiple individuals can attend to it. The contrast is with the way only I can attend to the character of my headache.

Yet just as public evidence is external, public evidence limits what is reasonable to believe. If Grassy Knoller has a pet theory about the Kennedy assassination, it is incumbent on him to account for the publicly available evidence. If he fails to engage with the public evidence—such as the Zapruder film, the findings of the Warren Commission, or ballistic analyses—then Grassy's theory lacks credibility. The point is not that public evidence eliminates disagreement, but that it sets the boundaries within which reasonable disagreement can occur. Public evidence acts as epistemic guardrails, preventing reasonable belief from veering too far afield from what is known.

Even so, there is a common way of thinking about public evidence that taken at face value does not limit what is reasonable to believe. Facts about the fossil record, the distribution of elements in soil and rock samples, the possibility of meteor strikes, and so on were not even conceived in ancient Greece. These unconceived facts did not constrain what it was reasonable for Anaximander or Aristotle to believe about the history of the earth. As Williamson observes "not all truths are evidence [because] [s]ome sort of epistemic accessibility is required."[13] An adequate conception of public evidence must satisfy the following thesis:

Limitation: Public evidence limits what it is reasonable to believe.

Externality and **Limitation** track two features of evidence that are often held to be in tension with each other. One way to make the tension into an explicit inconsistency would be to hold that **Externality** implies that public evidence is not the *content* of a mental state but that **Limitation** requires that public evidence is the content of a mental state. This inconsistency must be avoided by any adequate account of public evidence.

This dual way of thinking about evidence tracks with ways evidentialists have talked about evidence. Richard Feldman and Earl Conee distinguish between *scientific evidence* and *justifying evidence*.[14] Scientific evidence comprises states in the world—fingerprints, smoke, koplik spots—that reliably indicate other states—the culprit, fire, measles. This evidence is publicly available and can serve as a neutral arbiter in scientific debates. They characterize scientific evidence as "E is scientific evidence for P provided that E is publicly available and E reliably indicates the truth of P."[15]

Scientific evidence, as they conceive of it, does not limit what it is reasonable to believe.

[13] Williamson (2007, 209).

[14] See Conee and Feldman (2008, 84-85).

[15] Conee and Feldman (2008, 84).

A doctor unaware that koplik spots reliably indicate measles is not justified in believing that a person has measles when she sees these small lesions in a patient's mouth. In this case, a reason to believe that the person has measles requires knowing that koplik spots reliably indicate measles. Feldman and Conee characterize justifying evidence directly in terms of having a reason to believe: "justifying evidence is by itself a reason for belief, something one could in principle cite as a justifying basis for belief."[16]

Thomas Kelly draws a similar contrast, distinguishing between evidence in terms of "physical objects . . . one might place in a plastic bag, dig up from the ground, send to a laboratory, or discover among the belongings of an individual of historical interest" and evidence in terms of "the kind of thing which can make a difference to what one is justified in believing or what it is reasonable for one to believe."[17] With both Feldman & Conee and Kelly, we see a dual face to evidence, one tied to **Externality** seen in scientific evidence and another tied to **Limitation** seen in what directly bears on reasonable belief. In the following I develop two conceptions of evidence that fit with **Externality** and **Limitation**.

1.2 O-evidence & M-evidence

To develop an adequate theory of public evidence, I begin by distinguishing between two conceptions of evidence: one that treats evidence as consisting of facts, and the other that identifies evidence with a kind of mental state. Let's start with a specific case.

1.2.1 Fingerprints

Talk about 'evidence' is ambiguous between something that is out there in the world and something that directly bears on reasonable belief. Consider the following case. You are a detective investigating the theft of Rembrandt's *The Abduction of Europa* from the Getty Museum. You discover a telling fingerprint near the scene of the crime. The fingerprint–the three-dimensional physical object with a unique pattern—is evidence. This object has significance for discovering the culprit because it is a reliable sign of a person likely to be the culprit.

The significance of the physical object tracks the other way of thinking about evidence. The connection between the fingerprint and the likely culprit is currently unknown. The forensic technician observes that the fingerprint has a unique radial loop near its center. She indicates that this rare feature is likely to identify the culprit. You've now gained some additional evidence by coming to learn about a property of the fingerprint. But the evidence does not yet constrain any beliefs you might have about the identity of the culprit. The investigation continues, and the discovery is made that Moriarty's fingerprint exhibits this distinctive

[16] Conee and Feldman (2008, 85).
[17] Kelly (2016).

pattern. You have now gained evidence that constrains what it is reasonable to believe about likely culprits.

It's natural to say something like this: you've gained evidence about the culprit by studying your evidence. But while that's natural, it is also puzzling. If 'evidence' were used in the same way, it'd be like trying to get rich by inspecting your wallet. But money, like physical evidence, doesn't multiply just because you're staring at it. We need to make sense of how the object is evidence and also provides evidence. To do this, let's distinguish two kinds of evidence.

1.2.2 Objective evidence

To make sense of how one gains evidence without changing what evidence you have, we need to distinguish between two ways we often talk about evidence. Facts are often taken as evidence. It is the fact that the fingerprint has a unique central radial loop that provides a key bit of evidence. It is facts about the fossil record, the structure of DNA, or the shape of the continents that provide crucial explananda. To be sure, objects are often talked about as evidence–fingerprints, fossils, koplik spots, etc., but objects have a myriad of properties and only some of the properties of objects are relevant to a claim. It is a specific fact about an object that is evidence. Let us call this kind of evidence objective evidence (*o-evidence*). **Externality** is connected with thinking of evidence as o-evidence, though I'll argue that that's not the only way to understand it.[18]

1.2.3 Mental state evidence

A different sense of evidence is mental-state evidence (*m-evidence*). Whereas o-evidence is a fact consisting in an object and its properties, m-evidence is the content of a mental state. In the case described, you gain m-evidence by coming to learn that the fingerprint has an uncommon radial loop. The proposition ⟨the fingerprint has a distinctive radial loop⟩ is m-evidence by being the content of a mental state with assertive force. Similarly, once you learn 'Moriarty's fingerprint has that unique pattern' you gain evidence that Moriarty is the culprit.[19]

The same point applies in scientific contexts. One gains m-evidence by learning various facts about the fossil record. It remains an open question whether m-evidence extends beyond knowledge (or justified belief) of objective evidence. If, for example, intuitions count as m-evidence, then m-evidence would not be exhausted by knowledge of o-evidence. In the case

[18] In the debate over the ontology of evidence, objective evidence fits with factualism about evidence. My view is that reflection on public evidence offers a more nuanced view about the ontology of evidence.

[19] M-evidence doesn't require an explicit connecting fact. E.g., if you know p and know p implies q, that is enough to have justification for believing q. You do not need to explicitly have modus ponens represented as the content of some mental state. At some basic level, we have epistemic abilities that are different from the contents of mental states.

of public evidence, however, the m-evidence aims to be grounded in knowledge (or justified belief) of the relevant o-evidence. This reflects the idea that it is facts about the world that ultimately ground public evidence.[20]

This distinction clarifies the role of evidence in constraining rational belief. When we speak of o-evidence, we refer to objects and their properties that stand in causal, nomological, logical, or mathematical relations to other states of the world. Until those relations are taken up in cognition, however, o-evidence does not constrain what it is reasonable to believe. M-evidence, by contrast, satisfies **Limitation**: it is only when a proposition is known or justifiably believed that it can shape what it is reasonable to believe.

1.2.4 Phenomenalistic & public conceptions of evidence

Understanding m-evidence in terms of that which justifies belief makes sense of epistemologists sometimes peculiar ways of talking about evidence. Descartes holds that the ultimate basis of knowledge is clear and distinct perceptions.[21] Russell maintained that sense-data are the ultimate basis for empirical knowledge.[22] The logical positivists argued that observation sentences are the basis for scientific theory.[23] Quine held that evidence is restricted to the stimulation of one's sensory receptors.[24] Michael Huemer holds that seemings—perceptual appearances and intellectual appearances—are evidence.[25] In each case, the ultimate evidence is gained by being in some kind of mental state.

This way of thinking about evidence sits uneasily with scientific practice. It would be odd for Darwin to cite sense-data, the stimulation of sensory receptors, or seemings as evidence for the theory of evolution. Instead, he appeals to publicly accessible facts: the fossil record, the geographical distribution of species, and comparative anatomical structures. These are paradigmatic examples of public evidence. In contrast, epistemologists have tended to focus on evidence that comes from mental states that are essentially private. As Thomas Kelly observes '[t]his picture of evidence stands in no small measure of tension with the idea that a central function of evidence is to serve as a neutral arbiter among competing views."[26] It is the publicity of Darwin's evidence, the fact that it is available for scrutiny, debate, and shared evaluation, that enables the open and intersubjective character of scientific reasoning, in contrast to the first-personal and privileged nature of private experiential evidence.

[20] Michael Strevens articulates a related idea in his discussion of the Iron Rule, according to which scientific arguments are settled by empirical tests rather than by appeals to intuition. See Strevens (2020).

[21] See Descartes (1641/1996).

[22] See Russell (1914/1993).

[23] See Neurath (1932).

[24] "The stimulation of his sensory receptors is all the evidence anybody has had to go on, ultimately, in arriving at his picture of the world." Quine (1969, 75).

[25] Huemer (2001).

[26] Kelly (2016) See §4.

Publicity, as it contrasts with privacy, is a key feature of scientific evidence. Kelly states: "[I]t is a central methodological norm of science to eschew as inadmissible (e.g.) any alleged episodes of incommunicable insight in considering whether to accept or reject a claim."[27] Timothy Williamson shares a similar thought, "If one's evidence were restricted to the contents of one's own mind, it could not play the role that it actually does in science. The evidence for the proposition that the sun is larger than the earth is not just my present experiences or degrees of belief."[28] Williamson's central point in this context is that evidence is not something a person can infallibly identify. He argues that internalist conceptions of evidence are often driven by the assumption that individuals have privileged access to their evidence. But if evidence must always be recognizable from the inside, this yields a conception of evidence that diverges sharply from the one operative in scientific practice. On my account public evidence is m-evidence that is grounded in o-evidence. While the m-evidence is internal, what grounds it is not.

The distinction between o-evidence and m-evidence alleviates the tension between a phenomenalistic conception of evidence and a public conception of evidence. M-evidence, in virtue of being the content of mental state, makes a difference to what it is reasonable for a person to believe. But o-evidence, in virtue of being a state of the world that stands in a reliable connection to other states, does not constrain reasonable belief. A clear distinction between o-evidence and m-evidence is crucial for developing an account of public evidence. The aim is to find a general account of public evidence that satisfies both **Externality** and **Limitation**.

2 Limitation & M-evidence

2.1 General conditions

From our above discussion, we saw that m-evidence is what makes a difference to what is reasonable to believe. In order for evidence to satisfy **Limitation**, it must be the content of a mental state. In this section I argue for a particular conception of m-evidence.

> M-evidence is the content of a mental state with a creditable epistemic standing.[29]

Williamson observes that there are two general conditions for an item e to be evidence for a claim h. "First, e should speak in favour of h. Second e should have some kind of creditable

[27] Kelly (2016) See §4.

[28] Williamson (2000, 193).

[29] My view of m-evidence aligns with Juan Comesaña experientialism. See Chapter 6 'Experientialism' in Comesana (2020).

standing."[30] For e to speak in favor of h is, in the main, for e to boost the rational credence of h.[31] Where cr is a probability function that captures rational credence, e is evidence for h when $cr(h \mid e) > cr(h)$.

Limitation requires that an item e is evidence only if it is the content of some mental state. We thus have three general conditions for an item e to be evidence for h.

1. e should speak in favor of h.

2. e should have some kind of creditable standing.

3. e is the content of some mental state where e is hosted as true.

In the third condition, e's being hosted as true captures the idea that e has what Jim Pryor calls "phenomenal force" and what Laurence Bonjour refers to as "assertive content."[32] This mode of hosting the content explains why the content has a world-to-mind direction of fit as opposed to a mind-to-world fit that contentful states such as desires have.

2.2 Is m-evidence propositional?

This section argues that a general account of m-evidence should allow that assertive mental states may have non-propositional content that provides reasons for belief. This possibility fits well with a clear distinction between public and private evidence.

Williamson offers several lines of argument that evidence is propositional.[33]

Explanation First, he observes that we often choose between hypotheses on the basis of their ability to explain some proposition.[34] The hypothesis that 'my son drank all the milk' explains why 'there is no more milk in the fridge.' The nature of explanatory reasoning requires that the evidence used in an explanation is propositional. To be sure, one can say things like 'Explain the fossil record!' that seems to have the form of asking to explain an object. But the demand asks for the impossible unless the context makes it clear which claims about the fossil record stand in need of explanation.[35]

[30] Williamson (2000, 186).

[31] I say 'in the main' because evidence may speak in favor of h without changing its degree of rational credence. This can happen when, e.g., one has excellent testimonial evidence for a claim and then gains even more testimonial evidence. The additional evidence need not change one's rational credence in the claim even though it is evidence for the claim. See Christensen (1999); Joyce (2005); McCain and Poston (2014) for discussion.

[32] See Pryor (2000, 547) and BonJour (1985, 78).

[33] See Neta (2008, 95–101) for a criticism of Williamson's three lines of argument.

[34] See Williamson (2000, 194-5).

[35] Williamson (2000, 195).

Probabilistic reasoning Second, Williamson notes that probabilistic confirmation requires that evidence is propositional.[36] One can determine the probability of a hypothesis given some evidence by determining the probability of the evidence given the hypothesis and the prior probability of the hypothesis. But it is propositions that have a probability. There is a certain probability that *a red card was chosen from the deck*. And there's a probability that *a queen of hearts was chosen given that a red card was chosen*. In each case, probability attaches to a proposition.

Hypothesis elimination Third, Williamson notes that a certain function of evidence is that it can be inconsistent with other hypotheses.[37] Popper made this feature the centerpiece of his account of the testing of scientific theories. A test of a scientific theory aims to falsify the theory. Popper explicitly ties this to statements writing, "Every test of a theory, whether resulting in its corroboration or falsification, must stop at some basic statement or other which we decide to accept."[38] So testing of a theory aims for evidence that is propositional.

Experience Lastly, Williamson considers whether perceptual experiences provides evidence that is not propositional.[39] The content of perceptual experience is rich, much richer than the vocabulary we have at our disposal to express its content. Consider looking at a spectrum of blue color samples. The colors represented in your visual experience are much more fine-grained than a normal adult's color vocabulary. But a person can still discriminate between the shades.

Does this imply that experiences provide non-propositional m-evidence? No. Williamson observes that while experiences do not consist of propositions, experiences confer the status of evidence on propositions.[40] He explains,

> [T]he evidence for an hypothesis h consists of propositions $e_1, \ldots, e_n$, which count as evidence for one only because one is undergoing a perceptual experience ϵ.[41]

The challenge is to specify the propositions that the experience confers evidence on. The color spectrum example suggests that one can appeal to demonstrative concepts to take up in cognition the fine-grained detail in the experience. Then the proposition 'that [shade] is thus [shade]' becomes an item of m-evidence for you. Yet it is compatible with Williamson's

[36] Williamson (2000, 195-6).
[37] Williamson (2000, 196-7).
[38] Popper (1959, 86).
[39] Williamson (2000, 197-200).
[40] Williamson (2000, 197).
[41] Williamson (2000, 197).

thought that this proposition is evidence for you because the non-propositional experience is itself evidence for you. My point here is not to settle this issue but to rather observe that a general account of m-evidence should be consistent with the possibility that the non-propositional experience itself constrains reasonable belief and thus counts as m-evidence.[42]

Experiential evidence and public evidence Williamson's argument that evidence is propositional is useful for thinking about a potential account of public evidence. His first three reasons for thinking that evidence is propositional—explanatory reasoning, probabilistic reasoning, and hypothesis elimination—are all kinds of reasoning used in scientific debates over the upshot of the public evidence. Scientific facts are used in explanation, for updating one's probabilities, and are inconsistent with various hypotheses. More generally, the public facts can be used for reasoning in these ways.

It is likewise striking that these facts are different from the way experience provides evidence that may be different in kind from the propositional evidence. In these cases, the representational content of experience is too rich, too fine-grained to be captured by descriptive concepts. Rather the content must be incorporated into thought by demonstration. This is different from demonstrating a fingerprint or a fossil. In the former case, the demonstration is to an experienced quality; in the later case it is to an object out in the world. I'll argue below in section 3 that there is a key difference between non-experiential and experiential evidence. Here I note that Williamson's tacit contrast between *reasoning* and *experience* points in a similar direction.

2.3 Is propositional m-evidence knowledge?

Williamson's argument for $e = k$ proceeds as follows.[43]

1. All evidence is propositional.
2. All propositional evidence is knowledge.
3. All knowledge is evidence.

   Hence,

4. All and only knowledge is evidence.

I've argued in §2.2 that a general account of m-evidence should allow for the possibility that 1 is false. A primary virtue of $e = k$ is that it brings together **Externality** and **Limitation**.

[42] See McCain (2014) for a defense that non-propositional experience can provide evidence.
[43] Williamson (2000, 193).

If one's evidence is one's knowledge, then evidence is both factive and the content of a mental state. Moreover, given that knowledge can be shared, $e = k$ holds promise for an account of public evidence.

However, despite aligning with the two key features of public evidence I have emphasized, $e = k$ is ultimately inadequate. Knowledge has a structure where some items of knowledge are more basic and evidentially support other items of knowledge. This tracks a feature of public evidence, namely that not everything known on the basis of public evidence is itself public evidence. Moreover, given the fallibility of persons and groups, some public evidence may be false. Let us examine these two cases.

2.3.1 K doesn't imply E

Consider an opaque urn filled with an unknown mix of a thousand balls colored either red or black. One draws randomly with replacement from the urn two hundred times. Each draw results in a red ball. On the basis of these draws, one knows the urn contains only red balls. But one's knowledge that the urn contains only red balls is not itself evidence. Rather the evidence is the individual draws.

This case is related to one of Williamson's arguments for $e = k$ that has been criticized by Alvin Goldman and Ram Neta.[44] Suppose there is an opaque urn that has an unknown composition of red and black balls. The first 10 draws are red. Williamson claims that it is consistent with your evidence that the 11^{th} draw is black, but it is *not* consistent with your evidence that the 1^{st} draw was red. Suppose it turns out that the 11^{th} draw is red. Given that the first ten draws are all red, one has a justified true belief that the 11^{th} draw will be red. But this proposition isn't among one's evidence. Hence, justified true belief isn't sufficient for evidence. Williamson intends to suggest that something stronger is required for evidence, viz., knowledge. But Goldman observes that this argument only shows that a justified true belief isn't sufficient for evidence. Goldman's diagnosis is that the evidence must be non-inferential, not merely known.

It is true, of course, that knowledge can be used to extend knowledge. In the urn case, one may attempt to defend $k \rightarrow e$ in the following way. Suppose there are two stores I could have purchased colored balls from. The RG store sells only combination packs of red and green balls, whereas the R store sells only red balls. After coming to know that the urn contains all red balls, this is good evidence that I purchased the balls from the R store, not the RG store. The individual draws don't provide strong evidence that I purchased the balls from the R store. On this line, it is the known universal claim–not the individual draws–that does the evidential work.[45]

[44] For the original case, see Williamson (2000, 200-201). For criticisms, see Goldman (2008) and Neta (2018, 37).

[45] Thanks to John Keller for this case.

Two replies mitigate this objection. First, while it's true that, e.g., the second draw of a red ball isn't great evidence that I purchased the balls from the R store, it is stronger evidence for the R store than the RG store. This follows from the fact that the relevant Bayes factor—

$$\frac{P(\text{red} \mid \text{ R store})}{P(\text{red} \mid \text{RG store})}$$

—is greater than 1. Moreover, the conjunction of claims about the outcome of the individual draws massively supports the R store over the RG store. The universal claim need not itself be evidence to explain the support the individual draws provide for the R-store hypothesis.

Second, although we often reason from generalizations that are supported by our evidence, it is occasionally important to isolate precisely what is the evidence. Consider a meta-analysis of clinical trials on SSRIs. Suppose the analysis concludes there is a small but statistically significant benefit. One might reason to this conclusion by noting that studies 3, 7, and 8 show positive effects, while studies 1, 2, and 5 show none, and the rest are inconclusive. Yet if concerns arise about the reliability of the meta-analysis, it becomes crucial to revisit the primary data–the actual outcomes reported in each individual study. This may reveal problems with the conclusions drawn from the meta-analysis. So, while reasoning from general conclusions is often appropriate, there are contexts–especially ones involving uncertainty or controversy—where identifying the underlying evidence is essential.[46]

2.3.2 False evidence

Second, evidence may be false. A normally reliable testifier may have a slip of the tongue and offer a false report. Jenny may incorrectly report her phone number as 867-5309 after uncharacteristically switching the 3 and the 9. Even so, you have as evidence that Jenny's phone number is thus and so. Defenders of $e = k$ respond by claiming that your evidence is not that Jenny's phone number is thus and so, but rather the truth that Jenny says her number is thus and so. There are two general reasons to resist this move.[47]

First, Juan Comesaña argues that there are cases of rational action based essentially on false premises.[48] He considers a case of a child who acts rationally by reaching for what he

[46] Admittedly, it is a difficult matter to judge what is *exactly* the non-inferential evidential basis for a claim. My considered view is that there is ultimately no such thing independent of a context which specifies what is evidence for what. See Poston (2014) for details.

[47] In addition to the reasons I give below, there is something odd about this move in combination with resisting the idea that some knowledge isn't evidence. In that case, defenders of $e = k$ talk about the inherent difficulty of identifying non-inferential knowledge without returning to some kind of phenomenalist conception of evidence. But here they seem happy to acknowledge that there is a structure where some claims are supported by evidence without being evidence. See also Comesana (2020, 122) where he observes that Williamson's $e = k$ view has a mixed ontology of evidence: Factualism in the good case and Classical Foundationalism in the bad case.

[48] See Comesana (2020).

takes to be a piece of candy on the table. However, unbeknownst to him, it is a bowl filled with decorative glass designed to look like candy. The child rationally believes the false claim that 'the bowl contains candy.' Together with the desire to eat candy, the child rationally reaches for it. Comesaña argues that it is essential to the rationality of the act that this false premise is rational.[49]

Second, Tom Kelly argues that victims in a demon world have false evidence.[50] For such individuals, the world appears just the way it does to someone in a normal environment. On the basis of a perceptual experience of a glass on the table, a normal person rightly believes there is a glass on the table. For the victim of a demon world things are not as they appear. But the victim continues to believe external-world propositions on the basis of those experiences.

Williamson responds to this kind of case by claiming that people do not have as evidence that the world is thus and so (when it is not); rather people have as evidence "the proposition that things appear to be that way."[51] However, as Kelly points out, this response saddles the proponent of $e = k$ with the substantial empirical claim that people have many beliefs about the way things appear.[52] While this objection is not decisive, Kelly uses it as a springboard to argue that the demon world case should be considered as a case of misleading evidence. He ends up advocating a pluralism about evidence on which agents in bad cases can have false evidence, while agents in good cases can possess additional factive evidence unavailable to their demon-world counterparts.[53]

This conclusion has important consequences for the nature of public evidence. One might be tempted to think that if victims of the demon world have false evidence then normal individuals must have the same non-factive evidence. On this view, an account of public evidence would have to be neutral between our world and the demon world. But Kelly rejects this line of thought. As he puts it, "it is surely not a piece of hard philosophical data that our evidence never includes publicly known facts about the way the world is."[54] On his account, individuals in the actual world can have public evidence that includes such facts, whereas victims in bad cases do not.

To illustrate this, Kelly considers the case of climate scientists who are acquainted with the best available evidence for global warming. Such scientists may justifiably believe that global warming is occurring because their evidence includes facts that have been verified by the scientific community—facts about glacier retreat, ice sheet mass loss, and other well-documented phenomena.[55] Importantly, Kelly notes, their epistemic position is not replicated by individu-

[49] See chapters 4 and 5 of Comesana (2020).
[50] See Kelly (2025).
[51] Williamson (2000, 198).
[52] Kelly (2025, 23).
[53] He calls this 'an inclusive pluralism about evidence' Kelly (2025, 37).
[54] Kelly (2025, 36).
[55] Kelly (2025, 34).

als in a demon world who lack access to these facts and instead rely on misleading appearances.

Kelly's inclusive position is compatible with the possibility of false public evidence. This does not require elaborate conspiracy or global deception. Ordinary cases of error, misreporting, or flawed measurement can introduce false claims into the public domain. Nevertheless, the institutional mechanisms of scientific inquiry, peer review, and empirical validation serve as corrective forces. Kelly's view preserves the idea that our evidence often includes facts about the world while making space for the possibility of error.

2.4 Summing up: M-evidence, limitation, and public evidence

The account of m-evidence satisfies **Limitation**. M-evidence constrains what is reasonable to believe. The discussion in this section argued that m-evidence speaks in favor of a claim, has a creditable epistemic standing, and is hosted as true. Moreover, it may be non-propositional and may be false. But I've also highlighted that public evidence is propositional because it serves as the basis of public reasoning. In the following section I turn to **Externality**.

3 Externality & Non-experiential evidence

For an adequate account of public evidence, we need to satisfy **Externality** as well as **Limitation**. The strategy of this section is to identify a class of non-experiential m-evidence that is both externally grounded and capable of satisfying Limitation

3.1 What Monochromatic Mary knows

Frank Jackson presents a thought experiment about a neuroscientist, Mary, who knows all the facts about color perception one can know without ever seeing colors in a normal environment.[56] Mary is raised in a black-and-white room. She knows all the non-experiential facts about color vision. But when she steps outside of the monochromatic room and sees a typical red object in a normal setting, she learns something new–*this is what it's like to see red.*[57]

Jackson says that Mary knows all the physical facts about color vision in the black-and-white room. But a more neutral characterization is that she knows all the non-experiential facts about color vision. The facts that she knows are all the facts she can read about. They are the facts that a person may learn about color vision in a graduate program on the neuroscience of color vision.

[56] Jackson (1982).

[57] The point can also be put in terms of Mary acquiring a new ability (see David Lewis on the ability hypothesis). Jackson uses the case as an argument against physicalism. It's not my intention to come down on a side of that debate.

3.2 Mary & non-experiential information

The Mary case anchors the distinction between evidence that is not essentially experiential and evidence that is essentially experiential. Let us mark this difference by distinguishing between non-experiential evidence and experiential evidence. An item e of evidence is non-experiential if and only if coming to know e does not require the use of a particular sensory route. In the black-and-white room, the facts that Mary learns do not require any particular sensory modality to come to know. She can learn these facts by sight reading, by braille reading, or by being told. But she can't learn *what it's like to see red* by being told. She has to experience it herself.

Compare this with Plato's argument from common sensibles in the *Theaetetus* (184b–186e). Plato argues that there are ideas whose possession requires the use of a particular sense—redness, middle C, texture. But there are other ideas whose possession does not require a particular sense–one, two, sameness, difference, square. These ideas are common sensibles; they apply to multiple sensory ideas–there are two red objects, there are three separate notes of middle C, four surfaces with the same texture. Plato then draws the conclusion that these ideas are supplied by a source distinct from the senses which he identifies as reason.

Plato's identification of ideas of reason shares similarities with the way in which philosophers of mind talk about the information Mary knows in the black-and-white room. Torin Alter says that Mary knows the objective scientific information about color vision. He characterizes this information as follows: "A science is objective if and only if the truths it reveals (or the claims it makes) can be understood independently of having any specific sort of experience."[58] What counts as objective is then different from what can only be revealed by experience.

David Chalmers presents a different take on the information monochromatic Mary knows that isn't tied explicitly to experience. He claims that Mary knows the structural information about color perception. A structural truth is one that can be fully expressed using only mathematical, logical, causal (nomic), and spatiotemporal vocabulary.[59] For example, Mary might know that red light has a wavelength of approximately 700 nanometers and that such light triggers a certain neural pattern in normal human observers—truths expressible in physical and functional terms, independent of subjective experience.

In contrast, a non-structural truth involves aspects of conscious experience that resist such reductive description. Knowing what it's like to see red–its qualitative character–is a non-structural truth, one that can't be captured purely in logical or scientific vocabulary. Structural truths can be acquired by reading, reasoning, or being told; they allow us to understand much of the world without firsthand experience. But non-structural truths—like the

[58] Alter (2023, 4).

[59] Chalmers (2020).

feel of a color or the taste of a fruit—arguably require direct acquaintance.

Importantly, not all knowledge gained through experience is *experiential* in this strict sense. For example, one might learn that there are two pens on the table by seeing them. But this is not necessarily experiential evidence, because no specific sensory route is essential: one could also learn it through touch or hearing (e.g., tapping sounds). The key question is whether a particular sensory route is required for acquiring the knowledge. When it is not, the knowledge qualifies as non-experiential—even if it is learned through experience.

Thus, while much of our empirical knowledge is grounded in observation, a great deal of it is non-experiential in the relevant sense: it does not depend on any particular mode of perception. The success of the mathematical and natural sciences illustrates just how much of our understanding of the world can be built on this kind of structural, non-experiential information.

3.3 Mary & public evidence

The Mary case also serves as a touchstone for thinking about public evidence. The information Mary acquires in the black-and-white room consists of all the publicly accessible facts about color vision—the kind of knowledge one might gain through advanced study in neuroscience or perceptual psychology. This body of information exhibits three key features. First, it has public epistemic standing: it is the sort of information that can be verified, taught, debated, and scrutinized within scientific communities. Second, it is non-experiential: Mary does not need to have any firsthand color experiences to grasp it. She can learn it through books, lectures, or conversation—without ever seeing red. Third, drawing on the distinction between o-evidence (objective evidence) and m-evidence (mental state evidence), the information Mary learns is m-evidence: it consists of knowledge about o-evidence. Thus, Mary's knowledge exemplifies a body of non-experiential, publicly accessible m-evidence that aims to be grounded in objective facts.

We can use these features to characterize public evidence.

> **Publicly Credited Non-Experiential Evidence** (PCNE): An item of evidence e is public evidence if and only if (i) e is m-evidence that aims to be grounded in o-evidence, (ii) e is non-experiential evidence, and (iii) e has a *publicly credited* epistemic standing.

A weaker form of (PCNE) captures cases in which an item of evidence has not yet achieved public standing but is poised to do so.

> **Publicly Creditable Non-Experiential Evidence** (PCNE*): An item of evidence e satisfies PCNE* if and only if (i) e is m-evidence that aims to be grounded in

o-evidence, (ii) e is non-experiential evidence, and (iii) e has a *publicly creditable* epistemic standing.

PCNE* marks a difference in epistemic standing, not a difference in kind: the underlying conception of evidence is the same in both cases. The distinction between PCNE and PCNE* allows us to respect both **Externality** and **Limitation** while accommodating the dynamic way in which evidence becomes public over time.

3.4 Public evidence and common justification

What is it for e to have publicly creditable epistemic standing? I will argue that e has publicly credited epistemic standing if and only if e meets the non-factive epistemic conditions for common knowledge. Let us call this "common justification". For a proponent of $e = k$, the account can be reformulated in terms of common knowledge. I frame the following discussion in terms of common knowledge, but will formulate the explicate account in terms of common justification.

Let's use two cases as a touchstone for thinking about common knowledge. First, it is common knowledge that one drives on the right side of the road in the US. This common knowledge solves the coordination problem of how to structure traffic flow to minimize inefficiency. This knowledge is easily acquired. A child can learn that we drive on the right side of the road. An account of common knowledge should not require the impossible or the very demanding.

Second, just as common knowledge facilitates practical coordination, it also enables epistemic coordination. Consider a scenario where someone is arrested without being read their Miranda rights. If the arrested person remarks sarcastically to the officer, "Oh, we're just ignoring the Constitution now," the significance of this statement depends not merely on both parties individually knowing about Miranda rights, but crucially on each knowing that the other knows this fact as well. This points to the fact that common knowledge serves as a constraint on what is reasonable.[60]

Let's work with this account of common knowledge.

p is common knowledge in a community C if and only if

1. p is known by the vast majority of members of C,
2. (1) is known by most members of C, and
3. If knowledge is needed at level n, reflection on levels 1 to $(n-1)$ generates knowledge at level n.

[60] Thanks to Tim Perrine for this case.

This n-level account of common knowledge doesn't require infinite higher-order knowledge. It tracks the feature that common knowledge is easy to attain. The account restricts what is common knowledge to a community. It's common knowledge among philosophers that it's controversial whether there are any rationally compelling philosophical arguments. It's common knowledge among mathematicians that there are different sizes of infinity. There is common knowledge held by most adult individuals. Call this *the common ground*. Common knowledge held in a specialized community is *specialized common knowledge*.

Putting this together, we then have the following account of public evidence.

> **Publicly Credited Non-Experiential Evidence** (PCNE): An item of evidence e is public evidence if and only if e is m-evidence, e is non-experiential evidence, and e satisfies the conditions for common justification.

Williamson characterizes a thesis about evidence, which he terms *Evidence Neutrality*, that aims to capture the sense in which evidence is a neutral arbiter between rival theories. He explains that the hope for such a conception of evidence is that "whether a proposition constitutes evidence is *in principle* uncontentiously decidable, in the sense that a community of inquirers can always in principle achieve common knowledge as to whether any given proposition constitutes evidence for the inquiry."[61]

Williamson argues that *Evidence Neutrality* is false because "having good evidence for a belief does not require being able to persuade all comers, however strange their views, that you have such good evidence."[62] My characterization of public evidence does not suffer from this problem. (PCNE) doesn't require that you are able to persuade all people that such and such is public evidence. Take the public evidence in mathematics supporting that there are different sizes of infinity. Perhaps some person rejects Cantor's conception of how to measure size of a set by way of a bijective function between the elements of the target set and a base set. That person may be confident that they are right. But misplaced confidence doesn't undermine what is common knowledge among mathematicians.

(PCNE) does not include any dialectical requirements. It implies that among one's evidence we can distinguish what is public and what is not public. It may turn out that being justified in believing a theory requires more evidence than just the public evidence. Perhaps, skills and intuition are required where these are not public items of evidence. So Williamson's objection to *Evidence Neutrality* doesn't threaten (PCNE).

3.5 Two objections to PCNE

This section replies to two common objections to my account of public evidence.

[61] Williamson (2007, 210).

[62] Williamson (2007, 212).

3.5.1 Objection 1: Phenomenal knowledge is common knowledge

PCNE implies that a large body of evidence–namely, the essentially experiential evidence–isn't public evidence. The reason is that phenomenal knowledge is not broadcast in the way required for common knowledge. One might object that this is mistaken: knowledge of what it is like to see red is common knowledge. Children easily acquire this knowledge, and knowing what it is like to see red appears no harder to acquire than knowing which side of the road to drive on. Why, then, should public evidence be restricted to evidence that is not essentially experiential?

There is, however, a significant difference between knowing what it is like to see red and knowing which side of the road to drive on. The latter is easily correctable. It is open to view whether a car is on the wrong side of the road. By contrast, if cases of perfect color inversion are possible, then it is not open to view that a person experiences phenomenal green when looking at a red object under standard conditions.[63] This possibility undermines the higher-order knowledge required for common knowledge.

To illustrate consider two individuals–Clara and Iris–each of whom knows what it's like to see red. They therefore know the same fact and, at least, have mutual knowledge. Does Iris know that Clara knows what it is like to see red? This question must be distinguished from whether Iris knows that Clara has a way of identifying red objects by sight. The latter concerns whether Clara has a way–perhaps different from Iris's way–of visually identifying red objects, whereas the former concerns Clara's phenomenal knowledge. Iris's knowing that Clara knows what it is like to see red presupposes that Iris knows that she and Clara experience red objects in the same way.

The question thus becomes: does Iris know (i) that her own experience of red is the normal way people experience red objects, and (ii) that Clara experiences red objects in that same way? Iris may be able to come to know both claims, but only by inference rather than direct comparison. This already marks an epistemological difference between phenomenal knowledge and knowledge of driving norms. Moreover, Iris lacks knowledge of (ii) whenever she is in danger of acquiring evidence that Clara's experience differs from her own. In such cases, Iris does not know that Clara knows what it is like to see red.

The fact that higher-order phenomenal knowledge requires inferential knowledge about the similarity between individuals shows that phenomenal knowledge is not broadcast in the way as knowledge of ordinary objects is. If we are unsure whether there is a fossil of a particular shape in the ground, we can point to it, describe it, and allow anyone suitably positioned to see it for themselves. PCNE is therefore correct to restrict public evidence to what is not essentially experiential.

Alvin Goldman argues that the requirement that scientific evidence be public conflicts

[63] See Nida-Rʹumelin (1996).

with contemporary cognitive science which accepts as data the contents of a subject's introspective reports.[64] He considers, for example, cognitive research on 'feelings of knowing'.[65] This research investigates a person's ability to predict whether they can correctly answer a question about, e.g., how something works. People often have a feeling that they know how something works before they attempt to answer the question. Goldman observes that this research "presupposes that subjects' introspective reports are accurate as to whether or not they currently have a 'feeling of knowing"'.[66]

Crucially, however, there is a distinction between (a) knowing what it is like for a subject to have a feeling of knowing and (b) knowing that a subject reports a feeling of knowing, whatever its intrinsic character. On my view, the public evidence is (b), not (a). Knowing the intrinsic character of another person's feeling of knowing is arguably more difficult than knowing the intrinsic character of another person's experience of red. The former depends heavily on personal history and psychological idiosyncrasies, whereas the latter depends primarily on shared physiological structure. Consequently, Goldman's objection that public evidence doesn't meet a strong publicity requirement is undermined by close attention to the content of what researchers know in cases of 'feelings of knowing.'

3.5.2 Objection 2: PCNE implies there is little public evidence

Public evidence is abundant in the following sense: for any substantive intellectual debate, there is a body of publicly accessible objective facts that frames the dispute. But this raises a natural worry for my account. Is the public evidence for claims such as anthropogenic global warming really common knowledge? After all, most people do not know, let alone know that others know, the detailed empirical results on which climate science relies. That evidence lies largely within the purview of climatologists.

The worry trades on a failure to distinguish two notions already in play. There is indeed a sense in which the objective facts relevant to climate science–the o-evidence–are public: they are out there in the world, accessible in principle to suitably equipped investigators. But o-evidence, by itself, does not satisfy Limitation. Until it is taken up into cognition in a way that constrains reasonable belief, it does not yet function as evidence in the epistemic sense at issue here. My account of public evidence is explicitly an account of public m-evidence: evidence that both aims to be grounded in o-evidence and constrains what it is reasonable to believe.

Once this distinction is in view, the second reply follows naturally. PCNE does not require that all public evidence belong to *the common ground* shared by nearly everyone. Rather,

[64] Goldman (1997)

[65] See Metcalfe and Shimamura (1994); Nelson (1992)

[66] Goldman (1997, 529)

it allows for *specialized common knowledge*: propositions can be public evidence relative to a community whose members share the relevant background knowledge and stand in the appropriate higher-order epistemic relations. Within the community of climatologists, the relevant claims about temperature records, atmospheric composition, and modeling assumptions are not only known, but known to be known, and so on. They therefore qualify as public evidence in the sense articulated here.

Observations from the sociology of science that theories are often accepted only after entrenched opponents exit the field do not undermine this conclusion. They show, at most, that public evidence need not be rationally compelling for every inquirer. PCNE imposes no such requirement. Public evidence may instead function as a plausibility constraint, delimiting the space of reasonable belief without forcing convergence. Some agents may possess ideological or psychological blindspots from which recovery is not reasonably possible. That possibility is consistent with, and indeed helps motivate, the claim defended in the next section: public evidence constrains justification by shaping what experiential evidence can plausibly support, rather than by delivering knock-down arguments.

4 Why it matters – public evidence provides plausibility constraints

In this final section I suggest one way public evidence may function: it provides a plausibility constraint on the ability of experience to provide justification. We saw above that Kelly notes that "a central methodological norm of science to eschew as inadmissible (e.g.) any alleged episodes of incommunicable insight in considering whether to accept or reject a claim."[67] We also saw in our discussion of m-evidence that experience can provide justification for some beliefs. In the following I look at several cases that suggest that the public evidence provides plausibility constraints on the way experience can justify beliefs.

4.1 Mary again

When Mary steps outside of the black-and-white room she comes to know a new fact–this is what it's like to see red. This new phenomenal knowledge doesn't occur apart from a context that makes it plausible that she will learn something new when she sees a red object in a normal environment. What she knows about color perception and about her new environment makes it plausible that she will come to experience some novel quality and that quality will be *what it's like to see red*.

[67] Kelly (2016).

Mary's case is not a skeptical case, but it can be made one. Suppose Mary is told that she will enter a new room with ordinary lighting but where the colors of typical objects will be randomly changed. She knows that (e.g.,) firetrucks are red, that apples are red, and so on. But now she knows that the colors of the objects in the odd room are not typical. She enters the room and sees a firetruck whose color hasn't changed; it is still red. She learns something new–*that firetruck looks thus*. But she doesn't come to know that *that is what it's like to see red*. Why? Given what she knows about the atypical environment, the public evidence does not make it plausible that her experience will reveal what it's like to see red.

The lesson is that experiential justification is sensitive to public evidence. Whether an experience can justify a belief depends on whether the public evidence makes it plausible that the experience is appropriately connected to the relevant property. It's commonplace that background evidence shapes the justificatory significance of foreground evidence; whether observing a black crow confirms that all crows are black depends on background information about the environment, sampling conditions, and relevant biological regularities. The present point is that the public evidence plays a similar role in influencing the ability of experience to provide justification.

4.2 Expert and Novice

Consider Feldman's case of Expert and Novice.[68] Expert knows a lot about trees and Novice knows next to nothing about trees. They are both looking at a hornbeam tree. They have the same visual appearance, but Expert is justified in believing that that tree is a hornbeam and Novice is not.

What explains the difference in justification? The correct explanation must appeal to differences in their background beliefs. Given what Expert knows about trees, he can identify a hornbeam. But given what Novice knows about trees, he cannot identify a hornbeam. Notice that Expert's knowledge may consist entirely of the public evidence about trees. He's learned the characteristic signs of a hornbeam by reading and now comes to identify that tree as having those characteristics. The public evidence he has makes it plausible that he can visually recognize those features in the present case.

4.3 Reasonable disagreements

Feldman asks whether reasonable disagreement is possible when the parties to a disagreement share a single body of evidence.[69] He argues that epistemic peers who share their evidence

[68] Feldman (2003, 147).

[69] Feldman (2010).

cannot reasonably maintain incompatible beliefs. One option he considers and rejects is that individuals possess "private evidence" in addition to a shared evidential base.[70]

Feldman illustrates the issue with a case in which two detectives consider the same evidence but arrive at different conclusions.[71] Part of the shared evidence is that Lefty embezzled money from his firm; another part is that Righty was having an affair. Let us stipulate that the detectives share all of the relevant public evidence about the crime. The question is whether they can nevertheless be justified in differing beliefs about whodunit.

Suppose that both detectives interview the suspects. One detective, but not the other, has an intuition that Lefty is guilty. This experiential input cannot be shared in the way public evidence can be shared. Although the detective can report the intuition ("I just can't shake the feeling that Lefty is guilty"), reporting it is not the same as having the experience. Still, the relevant question is *not* whether this experiential m-evidence is public, but whether it can justifiably influence belief given the public evidence.

By stipulation, the public evidence makes it plausible that Lefty is guilty.[72] In that context, the detective's experience may permissibly tip the balance of total evidence, rendering it reasonable for that detective to believe that Lefty is the culprit. Importantly, this experiential input does not constitute rationally compelling evidence for others. Rather, it operates within the space left open by the public evidence. The lesson here is that incompatible conclusions can both be plausible relative to the same public evidence, and experiential m-evidence can play a justificatory role without functioning as public evidence.

5 Conclusion

I've offered an account of public evidence–Publicly Credited Non-Experiential Evidence (PCNE)–on which public evidence is non-experiential m-evidence that meets the non-factive epistemic conditions for common knowledge. This account satisfies *Externality* by grounding public evidence in facts about the world that are accessible to multiple agents and subject to public scrutiny, rather than in private mental states. It also satisfies *Limitation* by requiring that public evidence constrain what it is reasonable to believe, through its uptake into mental states that meet the conditions for common justification.

Beyond clarifying the nature of public evidence, the account illuminates its normative role. Public evidence does not merely underwrite scientific justification; it structures the space of reasonable belief more generally by imposing plausibility constraints on experiential justification. Experience can justify belief, but only against a background fixed in part by

[70] Feldman (2010, 210).

[71] Feldman (2010, 205–6).

[72] Of course it also makes it plausible that Righty is guilty.

public evidence. This framework explains how expertise, perceptual judgment, and reasonable disagreement can coexist with a shared evidential base.

The resulting picture preserves the epistemic authority of public evidence while accommodating the legitimate role of experience. In doing so, it offers a unified and publicly accountable epistemology that respects both the intersubjective demands of inquiry and the situated character of individual justification.

References

Alston, W. (1991). *Perceiving God: The Epistemology of Religoius Experience*, Cornell University Press.

Alter, T. (2023). *The Matter of Consciousness: From the Knowledge Argument to Russellian Monism*, Oxford University Press.

Bealer, G. (1998). Intuition and the autonomy of philosophy, *in* M. R. DePaul and W. M. Ramsey (eds), *Rethinking Intuition: The Psychology of Intuition and its Role in Philosophical Inquiry*, Rowman & Littlefield Publishers, pp. 201–240.

Beddor, B. and Dellsén, F. (2025). Inclusive inquiry, *Philosophy and Phenomenological Research* **111**: 796–821.

BonJour, L. (1985). *The structure of empirical knowledge*, Harvard University Press.

Chalmers, D. (2020). Spatiotemporal functionalism v. the conceivability of zombies, *Noûs* **54**(2): 488–97.

Christensen, D. (1999). Measuring confirmation, *The Journal of Philosophy* **XCVI**(9): 437–461.

Comesana, J. (2020). *Being Rational and Being Right*, Oxford University Press.

Conee, E. and Feldman, R. (2008). Evidence, *in* Q. Smith (ed.), *Epistemology: New Essays*, Oxford University Press, pp. 83–104.

Descartes, R. (1641/1996). *Meditations on First Philosophy*, Cambridge University Press.

Earman, J. (1993). Underdetermination, realism, and reason, *Midwest Studies in Philosophy* **18**(1): 19–38.

Feldman, R. (2003). *Epistemology*, Prentice Hall.

Feldman, R. (2010). Reasonable religious disagreements, *in* L. M. Antony (ed.), *Philosophers Without Gods: Meditations on Atheism and the Secular Life*, Oxford University Press.

Fratantonio, G. (2024). Evidential internalism and evidential externalism, *in* M. Lasonen-Aarnio and C. Littlejohn (eds), *The Routledge Handbook of the Philosophy of Evidence*, Routledge, pp. 51–62.

Goldman, A. (1997). Science, publicity, and consciousness, *Philosophy of Science* **64**: 525–545.

Goldman, A. (2008). Williamson on knowledge and evidence, *in* P. Greenough and D. Pritchard (eds), *Williamson on Knowledge*, Oxford University Press, pp. 73–91.

Hempel, C. G. (1952). *Fundamentals of Concept Formation in Empirical Science*, University of Chicago Press.

Huemer, M. (2001). *Skepticism and the Veil of Perception*, Rowman and Littlefield.

Jackson, F. (1982). Epiphenomenal qualia, *Philosophical Quarterly* **32**(127-136).

Joyce, J. (2005). How probabilities reflect evidence, *Philosophical Perspectives* **19**: 153–178.

Kelly, T. (2016). Evidence, *in* E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, Metaphysics Research Lab, Stanford University.

Kelly, T. (2025). The new evil demon problem and the nature of evidence, *in* J. Comesana and M. McGrath (eds), *Knowledge and Rationality: Essays in honor of Stewart Cohen*, Routledge, pp. 18–42.

Kuhn, T. (1962). *The Structure of Scientific Revolutions*, University of Chicago Press.

Kuhn, T. (1977). Objectivity, value judgment, and theory choice, *The Essential Tension*, University of Chicago Press, pp. 320–39.

McCain, K. (2014). *Evidentialism and Epistemic Justification*, Routledge.

McCain, K. and Poston, T. (2014). Why explanatoriness is evidentially relevant, *Thought* **3**: 145–153.

Metcalfe, J. and Shimamura, A. (eds) (1994). *Metacognition*, Cambridge University Press.

Nelson, T. (ed.) (1992). *Metacognition*, Allyn and Bacon.

Neta, R. (2008). What evidence do you have?, *British Journal for the Philosophy of Science* **59**: 89–119.

Neta, R. (2018). Your evidence is the set of facts that are manifest to you, *in* V. Mitova (ed.), *The Factive Turn in Epistemology*, Cambridge University Press, pp. 32–49.

Neurath, O. (1932). Protokollsätze, *Erkenntnis* **1**(3): 204–214.

Nida-R'umelin, M. (1996). Pseudonormal vision: An actual case of qualia inversion?, *Philosophical Studies* **82**(2): 145–157.

Plantinga, A. (1967). *God and Other Minds*, Cornell.

Plantinga, A. (1983). Reason and belief in god, *in* A. Plantinga and N. Wolterstorff (eds), *Faith and Rationality: Reason and Belief in God*, University of Notre Dame Press, pp. 16–93.

Plantinga, A. (2000). *Warranted Christian Belief*, Oxford University Press.

Popper, K. (1959). *The Logic of Scientific Discovery*, Basic Books.

Poston, T. (2014). *Reason and Explanation: A Defense of Explanatory Coherentism*, Palgrave MacMillan.

Pryor, J. (2000). The skeptic and the dogmatist, *Nous* **34**(4): 517–549.

Quine, W. (1969). Epistemology naturalized, *in* W. Quine (ed.), *Ontological Relativity and Other Essays*, Columbia University Press.

Russell, B. (1914/1993). *Our Knowledge of the External World*, Routledge.

Strevens, M. (2020). *The Knowledge Machine*, Liveright Publishing Corporation.

Williamson, T. (2000). *Knowledge and its limits*, Oxford University Press.

Williamson, T. (2007). *The Philosophy of Philosophy*, Oxford University Press.