

The Baltic International Yearbook of
Cognition, Logic and Communication

August 2025

Volume 16: *Justification, Evidence, and
Use/mention errors*

pages 1-70

DOI: <http://dx.doi.org/10.4148/1944-3676.1130>

JODY AZZOUNI
Tufts University

USE AND MENTION WITH RESPECT TO “KNOW”,
“BELIEVE”, “EVIDENCE”, “JUSTIFICATION”,
“HYPOTHESIS”, AND SO ON:
A HOT MESS

ABSTRACT: Use/mention errors are everywhere in epistemology. They even show up in Gettier’s classic paper. They bedevil philosophical attempts to understand “know”, “belief”, as well as propositional-attitude expressions at the most basic level. What, for example, are propositional-attitudes attitudes towards? Not towards propositions, as it turns out. Use/mention errors confuse philosophers in the most simple of ways: For example, they allow philosophers to think that evidence and knowledge are propositional, in the sense that evidence is true propositions and that knowledge is knowledge of propositions. Similarly, propositional attitudes are seen as attitudes towards propositions rather than what propositions describe. It is shown specifically in detail how Williamson’s influential views about $E = K$ —that evidence is what one knows is flawed because of use/mention errors.

1. THE FIRST INTRODUCTION

Contemporary epistemologists haven’t been paying attention to the subtle ways of certain words—those like “evidence”, “justification”, “hypothesis”, “know(s)”, “believe(s)”, and so on. Specifically, many recent discussions of these words, and many working presuppositions of contemporary philosophers about the relationships, both among these words,

and between these words and what they're taken to refer to, are riddled with unnoticed use/mention shifts, and consequently, use/mention errors. The terminology is Quine's (1951); and the errors I'm speaking of are a family of simple ones: confounding propositions, ideas, hypotheses and beliefs with what those propositions or hypotheses are about (and describe) and what those beliefs and ideas are of.

Here, "proposition", "belief", "idea", and "hypothesis" all refer to quasi-linguistic items. Why "*quasi-linguistic*"? Because, speaking explicitly of beliefs, although the point can be made more broadly, we often attribute beliefs to animals and so we shouldn't foreclose on beliefs *not* being constituted purely of *language*. The essential point is that propositions, beliefs, hypotheses, ideas, and the like have *intentional properties*: that is, they're *about* or *of* the objects that they specifically *attribute* properties and relations to. The objects in turn that propositions, ideas, beliefs and hypotheses are about usually *don't* have intentional properties—pebbles and tables are paradigmatic examples—unless we *endow* them with intentional properties (make various arrangements of pebbles on a table into a symbol system) or unless we're explicitly talking about what has such properties, i.e., if we're talking explicitly about propositions, beliefs, and so on.

So I'm describing use/mention errors more broadly than they're often described—and more broadly than how Quine originally described them. Included are usual cases where we slide from speaking of words to what the words refer to (and vice versa) but also included are cases where we slide from speaking of our idea of something to the something that idea is of. As I said, what I'm describing as a use/mention slide (in either direction) is one from something_a with intentional properties to the something_b that those intentional properties of something_a are directed towards. For example: the idea of Pegasus versus what the idea of Pegasus is of or the word "knowledge" and what that word is supposed to refer to.

One illustration that's given later in this paper of the systematic role of use/mention errors in philosophical arguments about "justification", "evidence", and the like, is Williamson's arguments that $E = K$ —that what we know is the evidence we have. Similarly mistaken is an implicit practice among many epistemologists of taking talk of evidence to essentially amount to talk of justification—what we might call $E = J$ —this

conflation also turns on use/mention errors.

Not all philosophers treat “evidence” as interchangeable with “justification”. Conee and Feldman (see [Feldman & Conee 1985](#) and [Conee & Feldman 2008](#)) explicitly distinguish evidence from justification, and philosophers discussing Conee and Feldman’s work do the same. Conee and Feldman, however, do claim that (2008, 83) “epistemic justification is a product of evidence,” specifically:

S is justified in believing *p* at *t* iff S’s evidence at *t* on balance supports *p*.

That is, S is justified in believing *p* if and only if S has evidence of a certain quality and/or degree for *p*. I discuss their views further in Section 14.

Yet another example of a use/mention error surprisingly occurs in Gettier’s seminal paper: a heretofore unnoticed misuse of variables. (I describe it in Section 9.) Use/mention errors also permeate contemporary philosophical discussions in philosophy of science—but I’ll only touch on this in the current setting.¹ Germaine to *this* paper is that these errors have led to and motivated a number of false but influential conclusions in epistemology.

In pursuit of the arguments of this paper, we’ll trek through some interesting material—specifically, the surprising ways that false assumptions about usage have far-reaching effects on the arguments for philosophical positions. I start (Section 2) with a brief rehearsal of the paradigmatic example of a use/mention error (first described by Quine). Section 3, by way of some illustrations, provides a brief introduction to the tendency in epistemology to conflate evidence and justification. Up to this point, the paper has been engaged with preliminaries.

In Section 4 we start in earnest: a discussion of propositions and that-clauses. I make a sharp distinction between these linguistic (or quasi-linguistic) items and what they describe or name, which in the latter case are *not* propositions—contrary to received wisdom—except when propositions are explicitly what are referred to by those clauses. Section 5 takes up a small clot of complications that arise when we try to characterize what propositions describe or corresponding that-clauses

¹ See [Azzouni forthcoming](#) for a discussion of how van Fraassen and other philosophers commit use/mention errors while defending the semantic view of scientific theories. I’m unconvinced that a cogent defense of the semantic view of scientific theories is possible: all the arguments for this position that I know of trade on use/mention errors.

refer to. There are three cases, (i) where a proposition is true and the corresponding that-clause refers to a worldly conglomeration of objects and relations (e.g., “The Earth still had living beings on it in 2025”), (ii) where the proposition is true but the corresponding that-clause doesn’t refer to a worldly conglomeration of objects and relations (e.g., “Sherlock Holmes is much sharper than Donald Duck”) and (iii) where the proposition is false, and so the corresponding that-clause doesn’t refer to a worldly conglomeration of objects and relations (“Bertrand Russell was alive in 1796”). I suspect that how the nonexistent gets entangled here motivates use/mention errors in this area similarly to how Quine hypothesized certain confusions arise with respect to Pegasus and the idea of Pegasus (discussed in Section 2)—but this isn’t the only motivation behind use/mention errors. The usage surrounding propositional-attitude phrases, talk of beliefs and evidence, etc., invites use/mention errors independently of issues of the nonexistent. In this section I also introduce the jargon, “fact” and “purported conglomerations of objects and relations”, and explain how I’ll use these phrases. In Section 6, I introduce a deflated notion of propositions as *interpreted sentences* (a characterization that’s compatible with the metaphysically-richer views of propositions many philosophers are drawn to). This is to keep the analysis here clean of extraneous matters, primarily complicated metaphysics. I show how it’s natural to run together talk of propositions and that-clauses with what these things describe and refer to. In Section 7, I take up “said” and show how its direct-quotation use corresponds to it being propositions that are said while its indirect-quotation use corresponds to it being the *relatum* of the proposition spoken by the person indirectly quoted. This is a use/mention shift.

This material in place, I turn to “propositional-attitude” expressions and analyze them (in Section 8). The important lesson to be drawn is that the label “propositional attitude” is misleading because it engenders a use/mention error, where we mistakenly characterize a statement of the form $S \phi$ s that P (“Paul believes that Simon is nearby”) as describing a relationship between an agent and a proposition, instead of describing it as a relationship between an agent and what that proposition describes. Importantly, this directly bears on the question of what the objects of knowledge are.

Stumbles between use and mention, as I intimated above, arise

during philosophical attempts to characterize evidence, justification, knowledge, belief, and the like. This is because the words designating these things sometimes semantically operate with respect to use and sometimes with respect to mention. Apparent usage can also be tricky to interpret because we sometimes speak in terms of propositions for convenience when what's meant (and really being talked about) is what those propositions describe. Section 9 focuses specifically on “know(s)”. It's like the other propositional attitudes insofar as the objects of knowledge aren't propositions but instead facts. However, because of the way we speak of what we know being true, we tend to slide very naturally from speaking of what we know as facts to what we know of as truths. Talk of “truths” in turn wallows in use/mention erroring: We understand the noun “truths” sometimes as referring to propositions that are true and sometimes as facts, what those propositions describe—as I illustrate earlier in Section 7, where we dramatically speak of “The Truth” when speaking of facts. (“The truth has finally been exposed”—we're not speaking here of the truth of a *proposition* finally being exposed.)

Similar complications bedevil how we speak of evidence, and what we take evidence to be. “Evidence” is complicated (and, as usage invariably is, misleading). This is the topic of Section 10. We *can* talk of evidence in propositional terms, but such talk is pleonastic and what's intended—“evidence”—is invariably *not* propositional. To make this case, I describe how we speak of evidence and show how evidence, as a result, can't be collections of propositions or the like but instead is often conglomerations of relations and objects. A bunch of bloody clothes in the corner of a hotel room is evidence. “A bunch of bloody clothes is in the corner of the hotel room” isn't evidence; it's a statement *about* evidence. Usage, however, does allow evidence to be regarded as collections of facts, provided we recognize that speaking this way can be awkward because “fact” is a count noun but “evidence” is a mass term. Interestingly, *documents* and the like can be evidence, but it's not the document-type or what a document-type is composed of—propositions—that's evidence. It's the worldly physical items—the actual documents (paper, electronic items) which are the instantiations of document-types—that are the actual evidence. A similar point can be made about testimony when (for example) it's submitted as evidence.

Section 11 takes up Williamson's $E = K$, and follows this (Section 12) with an analysis of Williamson's arguments for evidence being propo-

sitional; I illustrate how stumbles over use and mention vitiating arguments. In Section 13, I consider “justify” and “justification” and compare them to “evidence”.² Despite a lot of overlap in usage, these words bear complex relationships to one another, to our beliefs that we can have evidence for (or which we can justify), and to knowing agents and the knowledge such agents have. In Section 14, I consider Conee and Feldman on evidence, specifically their thesis that epistemic justification is a product of evidence. Their arguments, for the most part, don’t fall afoul of use/mention errors—rather, they engage in what can be called *surreptitious conceptual engineering*: modifying the ordinary understanding of “evidence” to enable establishing their thesis without announcing explicitly that that’s what they’re doing.

Next (Section 15) is a discussion of how debates about scepticism have floundered over failures to be scrupulously clear about usage, although (as I hope this paper illustrates) being scrupulously clear about usage is not only required but very hard to manage. In Section 16, I use the foregoing material against Leite’s (2024) picture of how the sceptic can be responded to. His position is that when we adopt a “pre-philosophical” stance, when we simply employ ordinary epistemic practices (and don’t engage in philosophical theorizing of any sort), we can nevertheless successfully rebut scepticism. Any such pre-philosophical stance is unstable, however: *philosophy begins on the ground floor*. In making this point, I engage in a little methodological moralizing, which I’ll start up with here: We must recognize that ordinary ways of speaking about knowledge, evidence, and the like, are fraught with subtle traps—use/mention errors are only one intricate example of such. Thus, we need to evaluate carefully when we can change natural language (e.g., replace it with partially-technical artifacts that circumvent the traps) and when we can’t—when we must just continue to make do with what we have—but be very careful. In any case, ordinary usage can’t be taken at face value.

Here’s another way of putting the point. In any number of surprising and unexpected ways, discourse—formal and otherwise—is defective or user-unfriendly (misleading). But the defects and ways it’s user-

²I’ve described our talk of justification, both propositional and agential, twice before (Azzouni 2020, Chapter 7, Azzouni 2025, Chapter 4). I’ll be revisiting—and amplifying—some of that material here.

unfriendly can't always be eliminated. This is true of *all* engineering—human designed and evolutionary—so the *conceptual* engineering we attempt by modifying natural language is no exception.

Section 17, finally, summarizes some of the surprising sources of the results unearthed in this paper. I summarize these sources because in that way I not only accentuate the importance of paying close attention to usage, but also illustrate why philosophy shouldn't be composed of specialized studies, the sort of thing we (often) find in the sciences, and that proves to be so valuable for those sciences. Philosophy is a rather different animal. One reason is that philosophy is concerned with the topic-neutral elements of our conceptual home, and those topic-neutral elements are spread across that *entire* home. This is why deductive logic and truth have been central topics of philosophy for so long. It's also why argumentative fallacies are also a central topic of philosophy: use/mention errors, as illustrated here, are among the argumentative fallacies that reasoners are prone to commit.

2. THE SECOND INTRODUCTION

Consider:

Pegasus doesn't exist.

One of Quine's (1980) concerns is that in speaking of Pegasus not existing it seems one is nevertheless speaking of *something* and noting that the *something* doesn't exist. Going beyond Quine's discussion of his example, consider:

Hercules doesn't exist.

It's very hard (I say "impossible"—I can't do it) to avoid the impression—the "intuition"—that these two sentences are describing *two different things* as not existing. But this on the face of it is incoherent. Presuming the sentences are *true*, we're not speaking of *anything* in either case—let alone two things—when we say that neither Hercules nor Pegasus exist.³

³I call these intuitions "aboutness illusions". We—we humans—can't escape the intuition that these sentences are *about things*, and attributing a property (nonexistence) to them. This drives Meinongian positions of one sort or another: there *are* such things—although they don't exist, or maybe (even) they don't have any "being". But they *are*, some

According to Quine (1980, 2), McX (a fictional philosopher) describes Pegasus as “an idea in men’s minds.” The *idea* of Pegasus is what people are talking about when they describe Pegasus as not existing or as being the flying horse of Greek mythology. This is confused (as Quine points out) since what’s being talked about isn’t an idea—the idea (of Pegasus) isn’t *itself* a nonexistent flying horse of Greek mythology. The idea (of Pegasus) is *of* a nonexistent flying horse of Greek mythology. As the grammar indicates, “Pegasus”, the word, doesn’t refer to an idea. It refers, as it turns out, to nothing at all. This example, however, makes it sound like use/mention confusions are *blatant* confusions, and that nothing subtle about us or how we cognize, and nothing deep about usage, drives them. Quine (1980, 2) writes:

McX never confuses the Parthenon with the Parthenon-idea. The Parthenon is physical; the Parthenon-idea is mental The Parthenon is visible; the Parthenon-idea is invisible. We cannot easily imagine two things more unlike, and less liable to confusion, than the Parthenon and the Parthenon-idea. But when we shift from the Parthenon to Pegasus, the confusion sets in—for no other reason than that McX would sooner be deceived by the crudest and most flagrant counterfeit than grant the nonbeing of Pegasus.⁴

However, use/mention errors are rampant in mathematics and, Meingongians say, *nevertheless*. Or (other Meinongians say), they aren’t; but nevertheless the sentences in question are about *them* and the truth values of such sentences are due to the sentences describing the properties such nonexistent things have. The metaphysics here (that is) can get *perverse*—there’s no better word. See Azzouni 2014. Another family of positions motivated by aboutness illusions—not *so* perverse—is that *there are* such objects but they are socially-constructed items. And yet another family of views targets our concept of “existence”—this notion is lightweight, these philosophers say, insofar as *anything we talk about exists*. So Hercules and Donald Duck exist too, at least according to any sense that we can make of our notion of “existence”. Left far behind is the ordinary impression that, of course, when telling a child (debriefing a child) about Santa Claus, a parent is simply explaining in the most natural way that Santa Claus doesn’t exist, and that nothing weird—metaphysically speaking—is going on. Some of the things we talk about exist and some don’t, and unfortunately children (among others) are misled about which is what.

⁴In fairness to Quine, it’s quite likely he wouldn’t understand this as an example of a use/mention error—he seems to have restricted such errors only to the words of artificial and natural languages and their referents.

derivatively, in physics; and what leads to them in those subject areas, whatever it is, aren't concerns with the nonexistent.⁵ Here's an example from a physics text, by Landau & Lifshitz (1976, 5), footnote omitted, italics and bolding theirs:

We can now draw some immediate inferences concerning the form of the Lagrangian of a particle, moving freely, in an inertial frame of reference. The homogeneity of space and time implies that the Lagrangian cannot contain explicitly either the radius vector \mathbf{r} of the particle or the time t , i.e. L must be a function of the velocity \mathbf{v} only. Since space is isotropic, the Lagrangian must also be independent of the direction of \mathbf{v} , and is therefore a function only of its magnitude, i.e. of $\mathbf{v}^2 = v^2$:

$$L = L(v^2). \quad (3.1)$$

Since the Lagrangian is independent of \mathbf{r} , we have $\frac{\partial L}{\partial \mathbf{r}} = 0$, and so Lagrange's equation is

$$\frac{d}{dt} \left(\frac{\partial L}{\partial \mathbf{v}} \right) = 0,$$

whence $\frac{\partial L}{\partial \mathbf{v}} = \text{constant}$. Since $\frac{\partial L}{\partial \mathbf{v}}$ is a function of the velocity only, it follows that

$$\mathbf{v} = \text{constant}. \quad (3.2)$$

Thus we conclude that, in an inertial frame, any free motion takes place with a velocity which is constant in both magnitude and direction. This is the *law of inertia*.

It's impossible to tease out in Landau and Lifshitz's discussion wherein the concern is a piece of notation—or instead wherein it concerns a

⁵I'm leaving aside my nominalism; whether mathematical objects exist or not isn't an issue in this paper.

mathematical object with certain properties that the Lagrangian notation designates.⁶ This, as I said, is common in mathematics; it occurs *harmlessly* everywhere because many properties of mathematical objects are reflected in the properties of the notation designating them, and vice versa. Many beautiful examples occur in abstract algebra where notational concerns, with indices, are treated simultaneously as concerns with the properties of certain algebraic structures that are designated by the relevant notation. We may want to charge physicists and mathematicians with “confusions” of use and mention, or we may want to simply regard it as part of the practice that practitioners seamlessly switch back and forth without notice between considerations of notation and considerations of the properties of the objects the notation designates. (I recommend, of course, the latter position.) Why is this seamless switching back-and-forth without notice possible? The answer: Because (as I noted) mathematical notation encodes methods of manipulation that correspond to properties/relations of the mathematical objects described by this notation.

The point of this paper is to illustrate in fair detail a third domain where practitioners—in this case speakers not belonging to any particular profession—routinely, seamlessly, and systematically slide between use and mention. These slides occur, for example, between the propositions attributing knowledge to agents and the knowledge so attributed (*what it is that we know*). Apart and additional to “know(s)”, sometimes *propositions* and the like are indicated by usage and sometimes it’s what propositions *describe* that’s indicated by usage, and sometimes it’s both. These slippery verbal practices *haven’t* disturbed our ordinary ways of attributing knowledge, justifying our claims, and providing evidence for what we believe; they *have* disturbed (to put it mildly) philosophical attempts to understand what we’re doing when we do these things.⁷

⁶The linguistic phenomenon so-exhibited is the same as in the sentence, “The Marvel hero Giant Man was so-named because of his size.” No error, per se, is involved, although this kind of construction—or ones like it (as we’ll see) can lead to errors.

⁷This isn’t to say that other ways we have of sliding back-and-forth when engaged in knowledge attributions, to ourselves, and others, don’t create problems for non-philosophers. Our ways of sliding between cognition and metacognition, for example, create *many* problems for philosophers *and* for non-philosophers (scientists studying metacognition, for example) as I indicate in Azzouni 2020.

3. THE THIRD INTRODUCTION

Gettier (1963, 121) begins his paper with several characterizations of candidate necessary and sufficient conditions for agential knowledge. Two are:

- (a) S knows that P IFF (i) P is true,
 (ii) S believes that P , and
 (iii) S is justified in believing that P .
- (b) S knows that P IFF (i) S accepts P ,
 (ii) S has adequate evidence for P ,
 and
 (iii) P is true.

Gettier attributes the second characterization to Chisholm; he doesn't give an attribution for the first. The subsequent focus of the gigantic descendant literature has overwhelmingly been on the purported inadequacy of *justified true belief* being necessary and sufficient for knowledge. As far as the epistemic literature is concerned, that is, the second characterization in terms of acceptance and adequate evidence largely amounts to the first—certainly with respect to Gettier counterexamples. Gettier seems to have thought the same since he titled his paper “Is justified true belief knowledge?” and not “Is sufficiently adequate evidence for an accepted proposition knowledge?” More recently, explicit discussion of evidence has emerged, and largely—although not altogether—replaced talk of justification, e.g., Leite 2024; Salow 2018; and Williamson 2000. Leite and Williamson, however (as well as others) don't seem to acknowledge that talk of justification and talk of evidence don't always parallel one another. Nor do they acknowledge that the differences may matter philosophically—specifically with respect to sceptical arguments.

Illustration: Chapter 5 of Leite 2024 is titled “Reclaiming our evidence”. It starts (97) with “the argument from ignorance”—attributed to DeRose—which is couched in terms of “justification.” Leite continues his analysis in terms of “justification” until (suddenly) we read (99):

It is crucially important here that in ordinary life it is entirely acceptable to appeal to considerations about the world as evidence for or against other claims, as reasons for accepting or rejecting suggestions, proposed hypotheses, and the like.

The discussion then continues in terms of “evidence” (although, when discussing Pritchard 2016, Leite switches temporarily to “rational support”, which is the phrase Pritchard uses); no worry seems to arise on his part that the debating ground may have switched along with the terminology. Leite continues to speak in terms of “evidence” until (106) we read:

So your belief is not Sensitive, but you have knowledge nonetheless. It should be obvious that the same point would apply to reasonable and justified belief, as well as to the notion having good or adequate reason to believe. We can thus have excellent evidence in favor of a belief even if that belief is not Sensitive.

The last sentence indicates how “justified”, “reasonable”, and “having excellent evidence” are all being presumed to largely come to the same thing. I really don’t want to sound accusatory here because I’ve sometimes run “evidence” and “justification” together myself in earlier work (see Chapter 5 of Azzouni 2025, for an example). By way of explanation, these words *really do* overlap a lot: one can sometimes be smoothly replaced with the other, and vice versa. Unfortunately, they don’t overlap entirely, and the ways they fail to overlap can affect the trajectories of certain arguments that arise between sceptics and anti-sceptics. That’s one of the things this paper is about.⁸

Genuine synonyms, it must be *stressed*, are virtually nonexistent in natural languages; and sometimes that matters philosophically. It matters, I show here, with respect to “justify” and “justification” on the

⁸How I’ve approached sceptical arguments in Azzouni 2025 sidesteps the differences between “evidence” and “justification”. This is largely because I’m there focused on “justification”, and only tend to use “evidence” when it’s natural in the context—which are cases, as it turns out, where the words are largely interchangeable. (Whew!) As it also turns out (see Section 13), justification is the central notion that sceptics should rely on to mount their knowledge challenges—and not evidence. So I’m being fair to the sceptic in that book by focusing primarily on justification and not evidence. (Again: Whew!)

one hand, and “evidence” on the other. These, it turns out, aren’t close in ways that determine what relationships between these words and “know(s),” “knowledge,” and “belief” are *possible*. The result, as so often happens, is natural language being (yet again) treacherous for those unwary philosophers who don’t pay sufficient attention to it. Specifically, there are interesting and significant differences between these words that bear on Cartesian sceptical scenarios.

A rough sketch of how these words deviate from one another and how this can affect the analysis of when we have knowledge and when we don’t vis-à-vis sceptical scenarios: The evidence of one’s hands is something an agent can have if they’re in a normal world (the one we think we’re in)—and have hands. *Having hands* is evidence that bears directly against the possibility that the agent in question is instead a brain in a vat. Correspondingly, an agent hasn’t the evidence of their hands if they’re a brain in a vat neurologically connected to a computer feeding them the illusion that they have hands). And yet, despite this asymmetry in the possession of *evidence*, the agents may not be *justified* in either case in believing that they have hands. That is, an agent in the normal world having *evidence* for their having hands doesn’t enable the ruling out of the possibility that, unbeknownst to them, they’re a brain in a vat without hands in the way that a *justification* for believing that it isn’t *possible* that one is a brain in a vat *can* rule it out. And this point benefits the sceptic (at least initially) because to meet the sceptical brain-in-a-vat challenge the anti-sceptic has to go after the possibility of being a brain in a vat *directly*. That possibility, if it obtains, changes the *evidence* that the believer in hands has for thinking they have hands. The believer in hands who’s a brain in a vat *thinks* they have evidence that they have hands. But even if the person isn’t a brain in a vat, if the sceptical scenario is a real one, the agent in question has to rule that out in order to be justified in thinking they have hands, even if they have plenty of evidence that they have hands.⁹

⁹Conee & Feldman (2008, 99-100) don’t quite agree. In a case of epistemic twins where one is looking at a tree and the other is looking at an identical-looking facsimile of a tree, they think those twins have the same “ultimate evidence.” By introducing jargon that enables a distinction between “kinds” of evidence, this finesses away the point that, as we ordinarily speak, such twins have *different* evidence. I’ll discuss this in more detail in Section 14.

4. THAT-CLAUSES AND WHAT THEY REFER TO

Consider the sentence:

That John is running again is a wonderful thing.

This isn't an expression of how wonderful the *proposition* "John is running again" is—notice that quotation marks *aren't* required. What's clearly being described as wonderful is an event or a series of events—what's wonderful is that John has, again, taken up an athletic activity. What's being described as wonderful, thus (I'll often use this phrase going forward), is a *worldly conglomeration of relations and objects* intentionally induced by John.¹⁰ Suppose someone instead says:

"John is running again" is a wonderful thing.

Or they say:

The proposition *John is running* is a wonderful thing.

Or:

The proposition that *John is running* is a wonderful thing.

For these locutions to be acceptable, an appropriate backstory is needed according to which the speaker is explicitly comparing various propositions (or sentences)—linguistic items—and describing some of *them* as wonderful things. Upshot: To speak about a proposition, to discuss *its* properties, is one thing (e.g., that the proposition is hard to understand; that the proposition is surprising; that the proposition *itself*—its beautifully-intricate semantic structure—is wonderful). To talk about what a proposition is about—its content, as it's sometimes put—is different (e.g., that John is running again is hard to understand—since his injuries haven't fully healed; that John is running again is surprising

¹⁰Many philosophers are tempted to use the word "fact" to refer to these worldly conglomerations. In fact (ha, ha), many of these same philosophers are tempted to wax metaphysical about facts. I don't think the metaphysics is needed or justified but issues about the metaphysics of "facts" lie apart from the topics of this paper. Nevertheless, we *need* a term like "fact" to label what must be distinguished from a proposition that's *true*—namely what that true proposition *describes*. I'll take this up in Section 5.

because he's claimed so often he'll never run again; that John is running again is wonderful because running is healthy). One lesson of this and the next couple of sections is that *what propositions describe* is often run together with *the propositions themselves*—these are use/mention errors. Another lesson is that it's easy to do this, and there are reasons for why it happens, among them those having to do with ease of communication.

Notice: “that John is running again” operates grammatically like a noun in the above examples; “John is running again” is a declarative sentence. So, the specific worldly conglomeration of relations and objects that's *referred to* by “that John is running again” is the same worldly conglomeration of relations and objects *that's described* by the sentence “John is running again.” For convenience, I'll use the italicized word “*relatum*” to describe the worldly conglomeration of relations and objects that, respectively, what “that John is running again” refers to and what “John is running again” describes. Common philosophical phrases for these are “facts” (as I mentioned in note 10) and “states of affairs”; and there are other phrases philosophers employ as well.

5. NON-EXISTENT CONGLOMERATIONS OF EVENTS AND OBJECTS

Consider again:

That John is running again is a wonderful thing.

Suppose, however, John *isn't* running again. What does “That John is running again” refer to in this case? It's clear that it can't be a worldly conglomeration of relations and objects intentionally induced by John because there's no such thing. However, that doesn't license us to think that the *proposition* “John is running again” is what's being described as wonderful. The distinction between propositions and that-clauses on the one hand, and what they describe/refer to, on the other, *remains*; it's just that in the case of false sentences, *nothing* is described by the propositions expressed and *nothing* is referred to by the corresponding that-clauses.

That-clauses function in these examples like nouns; and the same issue arises with them as arises with noun-phrases, generally—as I indicated in Section 2. Specifically: some names refer, “Barack Obama”, and some don't—“Mickey Mouse”. Nevertheless, one doesn't want to

say that when claiming that Mickey Mouse is more famous than any politician, or when claiming that Mickey Mouse doesn't exist, that one's talking about the word "Mickey Mouse", or the physical drawings, or the concept of Mickey Mouse. We aren't saying of *those things* that they're more famous than any politician and that they don't exist. Or if we are, we're making very different claims from the straightforward ones that *Mickey Mouse* is more famous than any politician or that *Mickey Mouse* doesn't exist. As I mentioned in Section 2, Quine (1980) warns us against making this kind of mistake. It's a use/mention error, in my broadened sense.

It's a tangled longstanding metaphysical question that intrudes into philosophy of language: how is it possible to talk about something that doesn't exist and how is it possible to say something true (or false) about what doesn't exist? I've long discussed this issue¹¹ but for current purposes we need only the following: Whatever it is "that John is running again" refers to (if anything at all), it isn't a proposition; it's either a worldly conglomeration of objects and relations or it's nothing at all, or it's what a worldly conglomeration of objects and relations amounts to when such a conglomeration isn't realized (doesn't exist).¹²

I'll try to encapsulate this with some jargon. When a proposition is true, and it describes a worldly conglomeration of objects and relations, that conglomeration *obtains*, and the corresponding that-clause refers to that very conglomeration. When the proposition isn't true, no conglomeration of objects and relations obtains; instead I use the phrase "purported". I'll also use "purported" when I'm speaking generally about the *relata* of propositions and that-clauses, and where the propositions may or may not be true and the that-clauses may or may not refer. This

¹¹ See, e.g., Azzouni 2010, Azzouni 2004. My view, perhaps notorious, is that although words such as "Hercules," and "Donald Duck" refer to nothing at all, the semantics of English allows sentences containing them to have varying truth values, for example, "Donald Duck is less famous than Hercules" (true) or "the ancient Greeks worshipped Donald Duck" (false). That is, there are sentences of natural languages with such words that differ in their truth values even though the things the sentences are talking about don't exist. Also see Collins 2022, 2025.

¹² What would *that* be? Well, that depends on the metaphysical tastes of one or another particular philosopher. Perhaps it's a *possible* worldly conglomeration of objects and relations that isn't "actualized", or perhaps it's one or another kind of Meinongian object. For current purposes, the details of the metaphysics of what doesn't exist don't matter. We just want to avoid confounding use and mention.

is a different use of “purported”, than the first, but using it both ways won’t lead to problems.

We’re (sadly) not yet free of the complications that the nonexistent plagues us with. Consider the sentence,

Sherlock Holmes is much sharper than Donald Duck.

This sentence is true, and yet there is no worldly conglomeration of objects and relations that this sentence describes. That is, correspondingly, “that Sherlock Holmes is much sharper than Donald Duck” doesn’t refer to a worldly conglomeration of objects and relations, because the only candidate for this purported conglomeration would be the much-sharper relation between Sherlock Holmes and Donald Duck, and since these items don’t exist, no instance of the much-sharper relation is instantiated by them. What is true is that there are many worldly conglomerations of objects and relations that are why “Sherlock Holmes is much sharper than Donald Duck” is true; but these things and their relations aren’t described by this sentence.¹³

Nevertheless, “that Sherlock Holmes is much sharper than Donald Duck” doesn’t refer to anything linguistic. Ordinary language allows us to say that “it’s a fact that Sherlock Holmes is much sharper than Donald Duck”, so I’ll canonize this usage as jargon by saying that “that Sherlock Holmes is much sharper than Donald Duck” *refers* to the *fact that* Sherlock Holmes is much sharper than Donald Duck. Consider a situation where people are arguing over the character traits of fictional characters. It can be said, emphatically:

It’s a fact that King Lear had *three* daughters.

Imagine, that is, a debate about how insightful King Lear is about the people around him, including his family. Someone could make the point that he had *three* daughters, as a way of indicating something obvious that King Lear missed. What allows this is that “It’s a fact that. . .” also

¹³What conglomerations are these? The answer: Ones involving stories, cartoons, movies, and the like. For current purposes, we describe these things as existing, even though—to repeat—they aren’t the topics of the sentence describing Sherlock Holmes as much sharper than Donald Duck. See Azzouni 2017, Sections 4.6-4.8 for further discussion of my theory of *truths*—of what it is about the world that makes various propositions (or sentences) true. Also see Collins 2025 for what strikes me as a similar (or possibly the same) theory of *truths*.

applies to the *relata* of truths, when those *relata* are things that don't exist, and therefore, aren't in instantiated relations.

We must be careful. I wrote earlier:

“that Sherlock Holmes is much sharper than Donald Duck”
refers to the *fact that* Sherlock Holmes is much sharper than
Donald Duck.

This is just to say that “that Sherlock Holmes is much sharper than Donald Duck” is a *fact*. That is, “that Sherlock Holmes is much sharper than Donald Duck” co-refers with “the fact that Sherlock Holmes is much sharper than Donald Duck”. *That is*, Sherlock Holmes is much sharper than Donald Duck is a fact, where the “is,” here, is the “is of predication” and not the “is of identity.” To instead use the “is” of identity, modify the previous sentence: That Sherlock Holmes is much sharper than Donald Duck *is* the fact that Sherlock Holmes is much sharper than Donald Duck.

“Factiveness” is a property being attributed to the *relata* of true statements and not those of false statements. More than that: to say, “that Sherlock Holmes is much sharper than Donald Duck is a fact” is to say something in some respects similar to—but *not the same as*—saying that: “Sherlock Holmes is much sharper than Donald Duck” is true. It's so easy to stumble here, to confuse use and mention.

We should leave for another day the question of what facts are metaphysically speaking.¹⁴ The important point, going forward, is that facts are *factive*. We speak of a that-clause referring to the fact that [...], when the proposition that corresponds to the that-clause is *true*. When the proposition corresponding to a that-clause is false, then the that-clause doesn't refer to a fact. I'm very very tempted to call what a that-clause (“that [...]”) refers to in this second case a *phact that* [...]. But I won't: I'll call these “purported facts that [...]” This jargon isn't

¹⁴Maybe there are just the *relata* of that-clauses, and the *relata* of that-clauses that correspond to true propositions are deemed facts. That's how I'm inclined to think. To speak thus of facts, however, is convenient, and enables the introduction of convenient truths (propositions about facts). That, however, doesn't require us to introduce any sort of complicated fact-metaphysics. Maybe (here's another approach) in the case of true sentences about worldly matters, the corresponding facts exist too—structured conglomerations of objects and relations—although they don't when the true sentences are about nonexistent items. Again: the metaphysics of facts is a matter for another day.

perfect because of the phrase “purported fact”—because “purported” implicates that the *relatum* of the proposition is something that people think could obtain, and people needn’t think that. Really, *phact* would be the best choice for jargon here, but as I said, I’ll stick with “purported fact”.

Notice: We can say, for emphasis, “That John is running again is a fact”. This *isn’t* to say the sentence “John is running again” is a fact, nor that what the sentence “John is running again” expresses—the proposition—is a fact. A proposition isn’t a fact; it can describe something *that is* a fact, in this case the fact that John is running again. The *relatum* of “It’s a fact that John is running again” is the same as the *relatum* of “John is running again” when “John is running again” is true. Neither are propositions, neither are linguistic. Fact-talk parallels truth-talk in interesting ways—primarily that one use of both kinds of talk is for emphasis. We can also use fact-talk the way we use truth-talk to facilitate quantification over what we take to be true or over what we take the facts to be. Just as we can say that “everything Einstein said on September 29th, 1940 is true”, we can also say that “everything Einstein said on September 29th, 1940 is a fact”. Truth-talk (see Section 6) usually involves property attributions to propositions or sentences; fact-talk involves property attributions to what propositions or sentences describe.

6. WHAT IS A PROPOSITION?

Here’s something else to be careful about: what is a proposition? It’s a natural way to speak—one that philosophers have taken up and that I’ve used in Section 5—to talk about propositions as things that declarative sentences “express”.^{15,16} Like “fact,” this bit of usage can (and has) led to a lot of metaphysics—impelled, in part, by the heard implication that what sentences “express” are entities of some sort. It’s best to minimize assumptions about propositions—at least for current purposes—so I here characterize a *proposition* as an interpreted sentence. That is, to

¹⁵E.g., Williamson 2000, 195: “What follows ‘why’ is a declarative sentence, expressing the proposition to be explained...”

¹⁶Use/mention conflation beckons. Another use, in the vernacular, of “express” is directed towards the world, e.g., the truism that “John is running again”, expresses the fact that John is running again.

say that a sentence “expresses” a particular proposition is to say only that the sentence—as a pure syntactic object—is being interpreted a particular way. This I claim (but won’t show in this paper with detailed usage data) is compatible with how we ordinarily use “proposition”.

This deflated notion of proposition is nevertheless sufficient to raise issues that are also faced by richer notions of propositions. The specific point I want to stress here is that propositions and that-clauses, respectively, are especially intimate with what they describe and refer to. They bear—this is a truism—semantic relations to their *relata*. But, as a result, they themselves *have* properties derived from their *relata*; for example, “John is running again”—the proposition—has the property of containing a noun-phrase that refers to John. Strikingly, these semantic relations enable us to describe properties of propositions or that-clauses and *thereby* describe properties of the *relata* of these propositions or that-clauses without mentioning the *relata* themselves.

I’ll illustrate this with an easy and paradigmatic example: the word “true”. “John is running again” describes, as I’ve said, a certain purported conglomeration of objects and relations; when true, the purported conglomeration of objects and relations *obtains*. “‘John is running again’ is true” describes a property of the *proposition* “John is running again”; however, in doing that, that statement indirectly expresses something about a proposition via its equivalence to a statement describing a certain purported worldly conglomeration of objects and relations. This equivalence—a Tarski biconditional (“*John is running again*” is true iff *John is running again*)¹⁷—is rooted in what “true” means, but it’s so natural to rely on the equivalence between the truth of a proposition and what that proposition describes that some philosophers have (mistakenly) identified the attribution of truth to a proposition *with* a description of a purported worldly conglomeration of objects and relations that that proposition describes.¹⁸ These are different. One is a statement about a proposition; the other is a statement about things in the world that aren’t propositions, but we naturally slide from one to the other. As

¹⁷So called because Tarski (1983) singled them out as important for any theory of truth.

¹⁸That is, some philosophers—“truth-deflationists”—have taken the proposition “John is running again’ is true” to be the proposition “John is running again” even though the first proposition is about *John* and that he’s running again, and the second is about a *proposition that describes John as running again*. See Azzouni 2018 for discussion and references.

we'll see in later sections, this sliding from talk of what that-clauses purportedly refer to or propositions purportedly describe to talk of (certain) properties of the that-clauses or propositions themselves is easy to do in the vernacular but leads philosophers to mistakenly think certain notions—notably, *knowledge* and *evidence*—are about propositions or that-clauses when in fact they're about what propositions (purport to) describe or what that-clauses (purport to) refer to.

7. SHE SAID, SHE SAID

Recall that I noted at the end of Section 5 that we can say either that “Everything Einstein said on September 29th, 1940 is true” or that “Everything Einstein said on September 29th, 1940 is a fact.” This indicates that the word “said” has a double role. When what follows *S said* is in quotations, what's attributed to S is the saying of a proposition. When what follows *S said* is a that-clause, what's attributed to S is the *relatum* of a proposition presumably uttered by S. We can say:

(I) Sarah said, “John is running again.”

Sarah said that John is running again.

Sarah told me that John is running again.

Sarah told me, “John is running again.”

All of these are acceptable. It's clear that when quotations are involved, the Sarah-said sentence is false if Sarah did not utter the expression “John is running again.” We can double-check by saying: “Did Sarah *say* that?” or more commonly “Did Sarah *actually* say that?” If Sarah said anything other than “John is running again,” a correction is called for. The item that's said or told is a particular proposition.

Indirect quotation is a little more complicated. First, since what follows “said” or “told” without quotations is a that-clause, above, “that John is running again”, what is literally said or told in those cases is a *relatum* of the that-clause, not a proposition. However, there's a difference here between “tells” and “told”, on the one hand, and “says” and “said” on the other. We can stress the use of “says” and “said” to require something close to an exact quotation. If someone says, “Sarah said that John was running again”, we can say, “Is that what she *said*?”

Or: “Is that what she *actually* said?” And the response can be: “Well, not exactly. She said that John has taken up running exercises again.” Or it can be “Yes, she did. She said that John has taken up running exercises again.” That is, there is some leeway for the speaker with indirect quotation. When we press with, “Is that what she said?” we need something agreeably close to what the person actually said, although it doesn’t have to be *exactly* what the person said. “Told”, “tells”, and the like, allow much *more* leeway. If someone says, “Sarah told me that John is running again”, one can respond with “Is that what she told you?” But in this case, what’s required is something that *amounts to* “John is running again”, there are no requirements on it being close to what Sarah actually said. For example, one can say, “Yes, but she didn’t say it aloud because she didn’t want John to overhear her. So she pantomimed it.”

Interim Conclusion: What can be said or told can be either understood as propositions or as purported worldly items. However, “said” tracks propositions a little more closely than “tells” does. It does so to the extent that, after uttering an indirect quotation, when challenged, “But what did he *say*?” the exhibition of something close to the sentence actually uttered is called for.

Now consider the following:¹⁹

(II) Sarah only tells you the facts.²⁰

Sarah only tells you the truth.

Sarah only tells the truth.

*Sarah only tells the facts.

*Sarah only says the facts.

Sarah only says the truth.

And these:

(III) Everything Sarah says is a fact.

Everything Sarah says is the truth.

¹⁹A common bit of punctuation in linguistics that I employ here is the use of an asterisk to indicate semantic infelicity or ungrammaticality.

²⁰Recall the remark popularly attributed to Joe Friday in *Dragnet*, an old television series: “Only the facts, Ma’am.”

Something Sarah said yesterday wasn't true.

Something Sarah said yesterday wasn't a fact.

When the focus is on the *relatum* of a that-clause, we can't say straight-out, as noted, "Sarah always says the facts," although we can say "Sarah always tells you the facts," or "Sarah will always tell you the facts." Also, we can say, "Sarah always states the facts." And, "Everything Sarah said during her testimony was a fact." Now consider the following exchanges:

(IV) B: "Sarah told me that John is running again. I don't believe it."

C: "If Sarah told you that, then it's a *fact*."

B: "Sarah said that John is running again. I don't believe it."

C: "If Sarah said that, then it's *fact*."

And notice:

(V) *If Sarah said that, then it's a truth.²¹

Better:

(VI) If Sarah said that, then it's true.

If Sarah said that, then it's the truth.

(VI) needs a little explanation. I think there's a genuine slippage between use and mention here that we've simply gotten used to. "What Sarah said" can either be directed towards a proposition or what that proposition describes. "It's true" is directed towards a proposition, but "the truth" (see next paragraph) is directed towards what the proposition describes. Furthermore, the implicit presence of a quantifier allows us to think easily of what's said either way. The sentences of (III) illustrate the same point.

What about "If Sarah said that John is running again then it's the truth"? What is *it*? I think that "the truth" is the culprit. What can be THE TRUTH—described dramatically—is either some propositions that are important but far more naturally it's what those propositions describe.

²¹This sounds off. Some may accept it but many won't.

Consider: “The truth has been revealed—that *you* are a murderer.” Or “The truth has been revealed—*you* are a murderer.” What has been revealed isn’t a proposition but what the proposition expresses: the fact that *you* are a murderer.

What about the awkwardness of the two sentences in (II) “Sarah only tells the facts” and “Sarah only says the facts”? I hypothesize that what’s involved here is that when what’s being said is treated as a *relatum*, it’s to go a step further to state that the *relatum* in question is a fact. We can’t always, therefore, identify what is said as a fact. Thus, although we can’t say, “Sarah only says the facts,” we can say, “If Sarah says that something is the case, it’s a fact.”

All this said, I think with “said” and “says” we’re sometimes slightly more comfortable using “true” than we are using “fact”—despite the complete naturalness of combining “said” with that-clauses. I hypothesize this is because the heard reading of “said” and “says” involves audibly uttered linguistic items—again, despite the complete naturalness of expressions like: “Sarah said that John is running again.” As I mentioned earlier, this discomfort vanishes utterly when quantification is involved, as the sentences of (III) indicate.

8. WHAT PROPOSITIONAL ATTITUDES ARE ATTITUDES TOWARDS

Consider the following sentences that start with propositional-attitude expressions:

Sam believes that John is running again.

Sam hopes that John is running again.

Sam knows that John is running again.

Just as the naked that-clause “that John is running again” has a purported worldly conglomeration as its *relatum*, here too usage impels us to recognize that what’s believed, hoped for or known isn’t a proposition (or a sentence) but something else. As suggested in Section 5, call the something else a fact when the proposition is true, call it a purported fact otherwise. Whatever this something else is, when it exists it’s something actual that’s *referred to* by the clause, “that John is running again,” or *described* by the sentence or proposition, “John is running again.” Otherwise, we can speak of it as well, as referred to or described,

even though it doesn't exist. We shouldn't think, in either case, that it's a proposition (or something similarly intentionally-proprietyed) that's referred to or described unless the subject matter of the sentence following the propositional-attitude expression (the subject matter that the propositional attitude is directed towards) is itself propositional or something similarly intentionally-proprietyed.

Unfortunately, the widely-used nomenclature "propositional attitude" hints that the attitudes in question are directed towards *the proposition that's expressed* by the sentence following the propositional-attitude expression. This isn't the case.²² What Sam believes isn't a proposition or that a proposition is true, but rather, that some purported worldly conglomeration of relations and objects is the case—a conglomeration composed of John, the running-again property, etc.—unless we explicitly say this instead:

Sam believes the proposition, "John is running again".

or,

Sam believes the proposition that John is running again.

These aren't natural ways to write (or speak), but there are other sentences of this form that are *very* natural:

Sam believes the proposition "All men are created equal".

or,

Sam believes the proposition that all men are created equal.

²²What were known as Millian approaches to propositions—broadly inspired by work of Russell's—treated propositions as worldly items, but ones that are structured (set-theoretically speaking, usually) and *containing* the objects they're intuitively about. So, for example, the proposition "John is running" was taken, on some of these views, to be an ordered pair (Running, John). I'm sidelining such approaches: when I speak of the *relata* of certain propositions or that-clauses as worldly conglomerations of objects and relations, I'm implicitly treating those propositions and that-clauses themselves as describing such things (and *not* as containing them). I also assume that what's described by a proposition or a that-clause isn't language-like (unless what's being described by a proposition or that-clause is a piece of language). My points can be made without these assumptions, but only much more prolixly, which I'm trying to avoid. The use/mention issues this paper is grappling with are complicated and confusing enough already; and debates over the ontological structure of propositions are mostly irrelevant to use/mention concerns, in any case.

The phrases “the proposition” and “the proposition that” aren’t semantically pleonastic in these sentences—after all, a proposition is being explicitly referred to and described as such. What, however, can believing a proposition (full stop) amount to? It can be short for, believing that proposition *is true*. One can believe a proposition is true without believing what the proposition *says*. This happens if the proposition is recognized by the agent to be true but isn’t understood by that agent. Many, for example, believe that $E = mc^2$ is true although many of those many don’t know what it means.²³ This possibility aside, believing a proposition, in this case, can come to nothing more than believing what the proposition says. Thus, “the proposition” adds nothing more to the original claim that Sam believes that all men are created equal. It does, however, induce an implicature that Sam strongly or fervently believes this.

9. KNOW(S) IS LIKE THE OTHER PROPOSITIONAL-ATTITUDE EXPRESSIONS

There’s a strong temptation felt by some philosophers to characterize the objects of knowledge as propositions, and there are some cases where we naturally do just this, for example:

How many theorems of number theory does Eric know?

These are special cases, however, because it’s quite unnatural to describe someone’s knowledge about trains or Napoleon or quantum field theory by alluding to a list of propositions that that person purportedly knows.

Furthermore, there’s other compelling usage-evidence that phrases such as “S knows that *P*” actually describe what the proposition *P* is about as what S knows, and not the proposition *P* itself. This is that we routinely attribute knowledge to nonhumans and even artifacts without language skills; and therefore, although propositions are only tools that *we* use to indicate the worldly conglomeration of relations and objects that such animals and artifacts know about—when we attribute

²³Here’s something, at least on the surface, that’s a bit strange. This can be said: “I believe that $E = mc^2$, although I don’t know what it means.” It can also be said: “I know that $E = mc^2$ —I memorized it in school and also learned that scientists that I should trust established it—although I don’t know what it means.” I’ll return to this example in the next section.

knowledge to them—there isn't a requirement that it's propositions that the knowledge of these animals or artifacts amounts to.

Consider, for example,

Gravy knows that the cat food is in the cabinet.

A sentence like this is a very natural thing to say (if one knows the cat is named "Gravy"). Imagine a situation where a guest sees the local cat, Gravy, scratching at the door of a cabinet, and so the host explains why. However, the host can cancel the implicature that Gravy grasps the concept of *cabinet* like so:²⁴

Gravy knows that the cat food is in the cabinet although, of course, she doesn't know what a cabinet is.

Or,

Gravy knows that the cat food is in the cabinet although, of course, she doesn't grasp English and wouldn't understand the sentence, "The cat food is in the cabinet," if you told her that.²⁵

Propositional-attitude cancellations of this sort aren't usually made about cats. That doesn't make such remarks unnatural *tout court* because they *are* made, regularly, about small children, when a parent (for example) is trying to distinguish what a child does understand from what that child doesn't understand.

The cogency of propositional-attitude cancellations is grounded in a propositional-attitude attribution practice that distinguishes propositions and the concepts that are used to construct such propositions from what the propositions are used to describe. Quite simply: What propositions describe can be characterized in many ways using different concepts and different propositions. If a propositional-attitude attribution were the attribution of a (specific) proposition to an agent, this kind of latitude would not be, *prima facie*, acceptable.

²⁴Here I'm reprising material on propositional-attitude concept cancellations from Azzouni 2020 and Azzouni 2025.

²⁵The host might be proudly observing that Gravy, in fact, does understand some words and a couple of short sentences, but not that particular one.

A possible counterexample? Recall from note 23, that someone can say: “I believe that $E = mc^2$, although I don’t know what that means.” If someone doesn’t know what a proposition means then, so the thought can go, they can’t be saying of whatever it is that the proposition describes (that $E = mc^2$) that they believe *that* because they don’t know what *that* is. *Response:* Why not? We can say, “I don’t know what’s out there, but I believe it’s bad.” Ignorance of what a proposition means doesn’t force a speaker to only speak about the proposition itself.

Nevertheless, it’s easy to slide from speaking of what one knows being a fact or a state of affairs to speaking of what one knows being a proposition. That is, it’s easy to slide among and between these:

Sam knows John is running again.

Sam knows that John is running again is true.²⁶

Sam knows that “John is running” is true

Sam knows “John is running” is true.

The factivity of “know(s),” coupled with the Tarski-biconditional equivalence of sentences and the truth of those sentences, is what makes the thought that these all come to the same thing an easy reach.

Still, knowing a proposition and knowing what a proposition describes isn’t the same thing, and they’re distinguished by grammar. If one says,

John knows the proposition “ $2 + 2 = 4$ ”.

that can be distinguished from knowing that $2 + 2 = 4$. We can describe someone who memorizes lots of propositions (for a quiz show, say) as knowing those propositions, but not knowing what these propositions describe. In such a case, the person need not even think the propositions are true. So this is distinct from knowing that $2 + 2 = 4$. Still, in practice, we come close to running these together, as I indicated in note 23. Someone can be taken to know that $E = mc^2$ even though they don’t know what it means. It’s also natural to protest: They don’t *know* that

²⁶Some, especially if philosophically trained, will hear (or see) this as infelicitous. Some won’t.

$E = mc^2$, they just memorized it.²⁷ This isn't fatal to the knowledge attribution if one can respond with "they didn't *just* memorize it, they *also* realized it's true."²⁸ I don't, therefore, take this example to show that the knowledge attitude is invariably directed towards propositions. Rather, that "know(s)" is factive coupled with our easy manipulation of Tarski biconditionals allows us to be sloppy about the verbal distinction between knowing a proposition, where what's literally meant is knowing *of* a proposition, P (and not what it says) and cases of knowing that proposition is true, or knowing that P . In both cases, though, despite how we speak we're usually not interested in the knowing relation to propositions but instead the knowing relation to what propositions describe. All this allows us (I dare say, even *motivates* us) to be sloppy about use and mention in these cases.

That the objects of knowledge aren't propositions is also illustrated by other grammatical locutions using "know(s)". For example, the command "Know that I love you" doesn't ask the person spoken to acquire knowledge of (or about) the proposition "I love you" (in the context it's uttered) or even to acquire knowledge that the proposition "I love you" is true. The fact depicted by the sentence is what's indicated—even though it's an easy inferential reach to the truth of the proposition itself.

It might sound as if placing factivity as a condition on "know(s)" requires what an agent knows to be propositions. "S knows that P implies P is true"—here I'm using variables just the way Gettier originally did—recall the opening passages of this paper. But this is a sloppy and misleading use of variables on Gettier's part, as we can see as soon as we replace them with sentences:

Sam knows that John is running again implies John is running again is true.

This sentence is ill-formed because we need quotation marks like so:

²⁷Notice the use/mention shift: "they don't know that $E = mc^2$ ", that is, they don't know what " $E = mc^2$ " describes. "They just memorized *it*." That is, they just memorized the proposition, " $E = mc^2$." Notice another use/mention conflation: We memorize propositions. But we often say things like: "She memorized all the facts about Watergate."

²⁸Here we've shifted to continuously talking about the proposition, " $E = mc^2$ " but *also* indicating what it describes by using "is true".

Sam knows that John is running again implies “John is running again,” is true.

And thus “S knows that P implies P is true” involves an illegitimate quantification into quotation marks because the first variable stands in for items that are used but the second stands in for items that are mentioned.²⁹

Let’s briefly consider the usage-relations between “know(s)” and “facts”. We can say,

John knows all the facts.

Here, “all the facts” describes everything about an implicitly given subject area. And it’s being asserted that John knows all of that. Notice again, “everything” about a subject area isn’t a proposition: it’s what a proposition *describes*.

Interim Conclusion based on the considerations raised in Sections 8 and 9: Propositional-attitude expressions describe relations between agents and the *relata* of the propositions that follow the word “that”. They don’t describe relations between agents and propositions, except when the propositions themselves are the *relata* of the propositions expressed by the sentences that follow the word “that”.³⁰

10. WHAT EVIDENCE IS

Let us turn to “evidence”: My claim, established in this section, parallels the ones made in Sections 4 and 5, with a rider. “Evidence” usually

²⁹Some philosophers have similarly sloppily characterized Tarski biconditionals this way: “ P is true iff P ”. How to characterize a generalization that governs Tarski biconditionals is an issue that the truth literature has been grappling with at least since the time of Tarski’s work. What blocks the expression of the generalization, if one uses ordinary logical tools, is precisely that the second occurrence of the variable occurs in a sentential position although the first one doesn’t. See *Azzouni 2018* for discussion and references; see *Azzouni 2006* for the presentation of anaphorically-restricted quantifiers, a logical regimentation that circumvents the problem.

³⁰*Mea culpa*: I have at times relaxed into describing the objects of knowledge as propositions in *Azzouni 2020* and *Azzouni 2025*. I (fervently or strongly) believe that the arguments and positions described in those books aren’t damaged by this particular mistake. I (fervently or strongly) believe, that is, a close revisit of those arguments and positions will show that I’m in the same position there that mathematicians and physicists are in when *they* slur over use and mention. Recall the discussion of this in Section 2.

holds of something purportedly worldly—more generally, it holds of the relata of “fact that” clauses, whatever those are. Or: in the case of facts like “Sherlock Holmes is much sharper than Donald Duck,” whatever we take the *relatum* of “that Sherlock Holmes is much sharper than Donald Duck,” to be. What evidence *isn't* is bodies of discourse or propositions. Additionally, “evidence” is a mass term, not a count noun—this will matter later when we come to evaluating the relationship between evidence-talk and fact-talk. The rider: the concerns discussed in Section 5 about cases with false propositions and their corresponding that-phrases don't arise with “evidence”: What “evidence” holds of—facts that . . .—always obtains. This will play an important role later.

I'll start with some sentences containing “evidence” that read strangely, and some that don't:

*Those truths are certainly evidence that he committed the crime.

*That proposition, “the smoking gun was in the hotel room,” is admissible evidence.

That the smoking gun was in the hotel room is admissible evidence.

We often *have* good evidence. Or, the evidence we have may be (unfortunately) misleading. As the sentences above hint, we *can't* speak of our evidence (or any evidence) being true or false:³¹

*That's true evidence.

*That's false evidence.

Evidence, therefore, is never the propositions describing something (unless propositions are the topic). And, further, evidence is nicely captured by the use of that-clauses:

That John is running again is good evidence that John is sensible.

³¹But Salow (2018, 702, note 16) writes: “Here and elsewhere I assume that only truths can be evidence. Williamson (2000, Ch. 10) influentially defends this claim.” I examine Williamson's defense of this evidence-are-true-propositions claim in the next section. I thank Helena Fang for drawing my attention to Salow's article. Salow's footnote was the seed from which all this grew: when I read it I stumbled over the phrase: “only truths can be evidence.”

That Mickey Mouse is more famous than any politician is good evidence for your theory of the central role of entertainment in American culture.

We thus identify evidence with the *relata* of that-clauses, subject to a necessary condition: that the propositions corresponding to those that-clauses must be true. For consider:

You don't have the evidence you think you have. There were no bloody clothes in the hotel room.

That's just not evidence for your theory: King Lear had *three* daughters, not *four*.

This suggests, since facts are what that-clauses refer to when those clauses correspond to true sentences, that perhaps we can identify facts with evidence. Can we? One obstacle that usage offers in opposition to this is that "fact" is a count noun and "evidence" isn't:

*That evidence are facts.³²

*Those facts are good evidence for your position.³³

What does work is this:

That fact is evidence.

Imagine someone gesturing at the bloody clothes in the corner of a hotel room, saying at the same time, "That fact is evidence against you." "That fact" individuates the bloody clothes as one thing, a fact. But "evidence" treats the same bloody clothes as stuff. Nevertheless, we can identify them this way. Similarly, those who can hear "Those facts are good evidence for your position" as acceptable are circumventing the clash (that I hear) between "Those facts," counting the facts, and "good evidence", which doesn't count anything. The same problem arises if someone points at the ocean and say, "Those molecules are water." I have trouble hearing that as acceptable; some won't have a problem.

It's important to recognize that words and their semantic properties are often contingent—words, that is, are contingent and historical; and

³²Equally bad: That evidence are the facts. The facts are that evidence.

³³Some may hear this as acceptable; see below.

so the properties correctly attributed to them can *change*, and at times these changes can be introduced deliberately. Although this isn't true of "know(s)" and "knowledge"—these are lexical universals that appear in all natural languages with the same properties³⁴—this *is* true of the words "fact," "evidence," and "justification". These words emerged in lexically complicated ways over the centuries during the scientific revolution. Specifically, "fact" didn't settle into its current semantic role—where facts can only be described by true propositions³⁵—until sometime after Newton's work had been disseminated.³⁶ The point is that these words still seem to be capable of semantic mutation, and so it's also possible, so it seems, to explicitly engage in conceptual engineering vis-à-vis them—something not true of "know(s)". I'll return to the issue of conceptual engineering when I discuss Conee and Feldman's views on evidence in Section 14.³⁷

Let us now consider cases where evidence is destroyed or hidden, and ask the straightforward question: *What exactly* is destroyed or hidden when evidence is destroyed or hidden? The accounting firm Arthur Andersen, for example, started shredding documents to allegedly hide certain facts about unlawful accounting practices they were allegedly engaged in.³⁸ Donald Trump allegedly hid classified documents in his bathroom, near a toilet. In both cases, what was *literally* being hidden were physical items—documents. When describing these cases it's important not to slip between talk of *types* and talk of *tokens*. We can describe a *document* as composed of propositions. In this case, we're describing documents-as-types: they're composed of propositions or of interpreted sentences. But that's not what was being hidden (that was not the evidence that was supposedly being hidden) by Arthur Andersen or by Donald Trump. The evidence being hidden is composed of tokens: actual physical items. We can recognize this by realizing

³⁴For evidence of this, see Azzouni 2020, Introduction; Azzouni 2025, Appendix.

³⁵This is put prolixly. There's a strong temptation to slur use and mention by writing instead, "where facts must be *true*."

³⁶See Shapiro 2000 and Wootton 2015 on this. It's ironic that the 2016 political maneuvering of the meaning of "fact"—the "your facts aren't my facts" bullshit—was in fact a revival of an earlier understanding of "fact" in the English legal tradition, one that originally appeared along with the word itself.

³⁷I thank Corey Dethier for reminding me of Conee and Feldman's work on evidentialism.

³⁸I owe the question of how one is to think of evidence being destroyed or hidden to Eric Dean; I also owe being reminded of this first example to him.

that propositions can't be shredded or hidden in boxes near a toilet. Only physical items can be treated this way. In the case of Donald Trump, the point is transparent because the document-tokens were marked "classified" and it's those markings (on the document-tokens) that were the nub of contention between those who claimed that Donald Trump broke the law and the lawyers defending him. The propositions expressed by the sentences in the document-tokens aren't relevant.

Perhaps the Arthur Andersen case is different. After all, the content is relevant: the documents being shredded (so one might argue) express propositions that the shredders are trying to hide. This is a mistake: they're not trying to hide the *propositions*, what they're trying to hide is *evidence* that certain propositions are true. What evidence is that? Well, the actual documents that are being shredded.

I used the phrase "to hide certain facts" above. What facts are these? Again: "the fact that the Arthur-Andersen accounting firm had engaged in such-and-such activities", or "the fact that Donald Trump had in his possession certain classified documents". One, in both cases, was trying to hide things. But one wasn't, by doing so, trying to hide propositions. One was trying to hide the *relata* of those propositions. That is, one was trying to prevent legal authorities from learning what certain propositions describe, coming to *know* these things.^{39,40}

The same point can be made about testimony. When testimony is submitted as evidence, it's not the testimony-type—a set of propositions—that's being submitted. What's being submitted is a token. A particular *instantiation* of a set of propositions is being submitted. Indeed, it's likely that what's being submitted as evidence is (evidence of) a (previous) *speech-act*. Evidence for the claim I've just made is that the testimony can be disallowed on grounds that don't involve the propositions expressed in the testimony being false. Submitted testimony is what, for example,

³⁹We confuse types and tokens all the time. I discuss this, and attempt to explain why in Section 0.1 of *Azzouni 2013*. Type/token confusions *aren't* use/mention confusions. And aboutness illusions (recall that neither Hercules nor Pegasus exist) are yet a third type of cognitive failing. Alas, we have many. And so many of these cognitive failings—although not all—make philosophy *hard*.

⁴⁰Again, notice the temptation to slur over use and mention, to enable putting things more neatly. One would like to write, "one was trying to prevent legal authorities from learning certain truths, coming to know them." More: to speak of *truths* is *sometimes* to speak of propositions, and *sometimes* to speak of what those propositions describe. Recall the discussion of this in Section 7.

someone heard, and so is evidence that certain things were said.

11. UNDERCUTTING E=K

Evidence is worldly. What we know is worldly as well. So, consider Williamson's influential suggestion that $E = K$, that evidence is knowledge. Is it?

Well, usage is unfriendly to this suggestion:

*Sam knows that evidence.

Acceptable, though, is this:

We already know about that evidence.⁴¹

And this:

I'm sure there's evidence for this hypothesis although no one has discovered it yet.

A piece of evidence is a purported worldly conglomeration of objects and relations.⁴² What we know is also a purported worldly conglomeration of objects and relations. Can they be identified as Williamson and philosophers influenced by him would like? Consider the evidence that an agent has for something. What's natural to say, as the first item above indicates, is that we know *about that*, where the "that" in question refers to the evidence (e.g., via a gesture towards bloody clothes in the corner of a hotel room). What work is "about" doing here?

⁴¹We sometimes drop "about" in locutions like this, writing: "We already know that evidence." But this sounds off precisely because of the missing "about". Far more natural is: "We already know about that [evidence]," where a gesture towards something worldly (for example) replaces the word "evidence". Still, one may worry because this sounds natural: "We already heard that evidence." But more natural still is instead saying: "We already heard about that." This second sentence reveals a use/mention slurring in "we already heard that evidence", since the two sentences are understood to come to the same thing.

⁴²"A piece of" is an awkward phrase that one reaches for when using a mass term and trying to count-noun it. The problem has been with us (in English) for a long time. Hume (1977, 72) writes: "But though animals learn many parts of their knowledge from observation, there are also many parts of it, which they derive from the original hand of nature." Words that help us individuate portions of what a mass term designates are more natural in other cases, "chunks of meat", "cups of water", and so on.

We talk about knowing about things, kinds of things, and worldly conglomerations of objects and relations. For example: “She knows all about trains,” “They know about the scandal developing in the White House,” “He knows everything there is to know about staplers.” When items of evidence are worldly conglomerations of objects and relations, it’s natural to describe someone as knowing about them, when they do know about them. Thus we can describe someone as knowing about *this* evidence for why John committed the crime but not *that* evidence. Evidence is something we know about when we *do* know about it. It’s, therefore, not itself knowledge. Knowledge is always knowledge that *someone* has. We don’t say, “I’m sure there’s knowledge for this hypothesis although no one has discovered it yet.” Knowledge is always had by agents but, as noted in Section 9, knowledge is worldly, not propositional. Evidence is only similarly relative to an agent when it’s evidence that’s *had* by that agent. So although evidence may be had by an agent, knowledge *must* be had by some agent or other. This is why the second sentence above is cogent. Evidence for something may not be had by anyone. That is, there can be evidence that no one knows about. $E \neq K$.

We should be careful about someone “having evidence”, because of the having relation. We often speak of someone having evidence *e* for *p* because they *know* that *e* is evidence for *p*. But someone can “have” evidence *e* for *p*, because *e* is in their possession although they don’t know that *e* is evidence for *p*. For example, a detective may have evidence that Lefty was at the scene of a crime because he has Lefty’s fingerprints on a bottle. But he doesn’t know the fingerprints are Lefty’s. He doesn’t knowingly have evidence that Lefty was at the scene of the crime, although he has that evidence in his possession.⁴³ When we describe someone as having evidence for something, it’s usually understood that they know it’s evidence. But that can be cancelled: “John had evidence that Lefty was at the scene of the crime, but he didn’t realize it.” I return to this aspect of our use of “evidence” in Section 14.

Evidence is always evidence *of* something, or for something. First consider:

⁴³This possession-without-knowledge is a common trope in detective novels, and sometimes the reader is duped too. The reader is given all the clues about who did the crime—the reader has all the evidence. But the reader doesn’t realize it. By the way, “clue” deserves a parallel analysis to the one I’m giving “evidence”; I can’t do it at this time.

S has evidence that *P*.

S knows that *P*.

These don't work the same way. For one thing (using the variable "*P*," sloppily, as Gettier did), S knows that *P* implies *P* is true; but S has evidence that *P* doesn't imply that *P* is true.⁴⁴ Now consider:

S has evidence of *P*.

S knows of *P*.

These don't operate the same way either. S has evidence of *P*, usually, when *P* is worldly, not propositional. For example, a geologist looking at a rock formation can say that the rock formation is "evidence of a volcanic eruption 100,000 years ago."⁴⁵ But if S knows of *P*, it's usually—if not always—the *proposition P* that S knows of. Also: "evidence of ..." or "evidence for ..." is *mandatory*: evidence is *always* evidence "of" something or "for" something—and the something is different than the evidence itself. Knowledge, on the other hand, is never really "for" anything; and although knowledge can be "of" something, the something is what's described by the proposition following the knows-that clause.

If an agent has evidence, and they know it to be evidence (for something) then that agent knows something *about* that evidence and in doing so, they usually know a number of specific things related to that evidence. I know about the bloody clothes in the hotel room. And I know that the bloody clothes in the hotel room are good evidence that John committed murder. I know that there are bloody clothes in the hotel room *now*; I know that the clothes weren't there yesterday. Similarly, the geologist knows that the rocks he's looking at are igneous rocks, and that knowledge is crucial to the rock formation being evidence *for him* of a volcanic eruption 100,000 years ago. Otherwise, it's still evidence of that, but not *for him*, since he doesn't know the rocks are igneous rocks. Knowledge of an agent affects whether something worldly (bloody clothes in the corner of the hotel room or a rock formation) is evidence *for him* of something or not. Indeed, evidence can be evidence

⁴⁴Compare Chomsky's old examples that "John is easy to please" and "John is eager to please."

⁴⁵My thanks to Eric Dean for the example and for illuminating discussion about the example.

for many different things, and so, because of the different knowledge about that evidence that people can have, it may be evidence of one thing for one person but evidence of something else for another. The bloody clothes in the corner of the hotel room may be evidence for one person that John was murdered, evidence for someone else that *someone* was murdered, and for a third person, evidence that the hotel staff hasn't been conscientious (because the third person doesn't realize blood is involved but does know certain facts about the hotel staff). None of this concomitant knowledge enables an identification between an agent's evidence and their knowledge. None of it makes the evidence for something (for an agent) that agent's knowledge nor does it make into evidence what the agent knows.

That S knows *P* is always a relation between what *P* describes (a worldly conglomeration) and the agent S. Evidence is something that an agent can knowingly have, and then there too is a relation between the evidence that *P* and the agent. But it isn't the same relation: to knowingly *have* evidence that *P* although deeply affected by what an agent knows about that evidence doesn't come to an identification of the evidence someone has with the knowledge they have. One, rather, knows *about* the evidence and thereby knows a lot of things that are what that-one-has-evidence-for is due to. This is as intimate as the relationship between the evidence an agent knowingly has and an agent's knowledge gets: quite close, but not an identity.

Once a philosopher mistakenly thinks that knowledge is of propositions and that evidence, too, amounts to propositions, it's possible for that philosopher to unwisely consider identifying the two, despite usage militating so much against this, and despite the delicate fact that although both words are relational—they're used to describe relations between agents and other things—the relations in question are *different*. Given these background mistakes, it's natural to further think that evidence is (composed of) true propositions, again despite this sounding so bad when said aloud (as I've illustrated). Having evidence of *P* and knowing that *P* are relations that agents have to facts (what true propositions describe); they aren't the *same* relations although they're closely intertwined.

A last point: It's hard to see why we should conceptually-engineer (stipulate) a notion of "evidence" that corresponds to Williamson's mis-

taken characterization of “evidence”. What would it be able to do that our current notion is unable to do? What would it be able to do in relation to our talk of knowledge that our already-in-place notion “evidence” can’t do? It’s hard to see what we would need such a notion for. The ways that “evidence” differs from “knowledge” are valuable; we need to keep them.

12. WILLIAMSON’S ARGUMENTS FOR $E = K$

Williamson is sometimes dismissive of usage. Regarding “evidence,” and his suggestion that all evidence is propositional, he writes (2000, 194), “How can ‘All evidence is propositional’ do more than stipulate a technical use for the word ‘evidence’?” Answering his own question, he continues: “Indiscriminate description of the ordinary use of a term and arbitrary stipulation of a new use are not the only options. We can single out theoretical functions central to the ordinary concept *evidence*, and ask what serves them. That strategy is pursued here.”

The three theoretical functions that Williamson considers (to avoid “indiscriminate description of ordinary use”) are inference to the best explanation, probabilistic confirmation and the ruling out of hypotheses on the basis of evidence. Let us see if any of them support the idea that the mass term “evidence” holds of propositions as Williamson thinks he shows.

I start with inference to the best explanation. As Williamson notes, this is a hypothesis-practice that’s not restricted to the sciences. We often, in ordinary life, choose between hypotheses that explain something by evaluating which (194) “would explain the evidence better than any other one would, if true.” Williamson’s example, however, is a scientific one: the fossil record. One or another hypothesis (about the existence of a kind of animal at such-and-such a time) along with other assumptions explains why certain fossils (located at such-and-such a place) are as they appear to be, and some hypotheses do a better job at this than others. Williamson (195) writes: “Thus evidence is the kind of thing which hypotheses explain. But the kind of thing which hypotheses explain is propositional. Therefore evidence is propositional.”

Williamson’s own example betrays him. That the fossil record is such-and-so isn’t a proposition—recall Section 4. That John is running

again is explained by his resolve and grit. But that John is running again isn't a proposition but the *relatum* of the that-clause, "that John is running again" and the proposition "John is running again" (notice that the first use of "that John is running" is missing quotation marks although the second one has them). That John is running again is a worldly conglomeration; his resolve and grit are also worldly; they're properties *he* possesses. So something worldly *can* explain something worldly, and propositions don't seem relevant.

But it looks like we can bring hypotheses in, if we wish, and we talk in a way that makes it seem that hypotheses are propositions that can explain things. We can say, for example, that the laws of gravity explain the motions of the planets. Gravitational laws of motion are general propositions; the motions of the planets aren't propositions, let alone general ones. So it looks like hypotheses, just like laws, can explain worldly conglomerations of relations and objects. But this isn't right either. It isn't the actual laws or hypotheses—the *linguistic items*—that explain; it's what those laws or hypotheses *describe* that explains. The laws of gravitation describe how massive objects affect and are affected by other massive objects. And that massive objects act in these ways on other massive objects (and vice versa) is what explains why the planets orbit the Sun in the ways that they so far have done. There's a systematic practice of slipping from speaking of what propositions say to instead speaking of the propositions themselves, and treating the propositions themselves as doing the explanatory work. But a proposition is a quasi-linguistic object, and a quasi-linguistic object doesn't explain (it can't explain!) why physical objects move as they do. If something general explains this, it must be something general that the universe exhibits: A way that the universe is isn't a proposition—it isn't a quasi-linguistic item that states that the universe is that way. The use/mention errors involved here are systematic and show up everywhere—not just in discussions about evidence and justification.

Williamson (195) writes further that:

Inference to the best explanation concerns why-explanations, which can be put in the form '— because ...', which is ungrammatical unless declarative sentences, complements for 'that', fill both blanks. We cannot simply *explain Albania*, for 'Albania because ...' is ill-formed. ... What follows 'why'

is a declarative sentence, expressing the proposition to be explained—*that* Albania exists, or *that* it has the distinctive feature.

The proposition expressed isn't to be explained; the fact or state of affairs that the proposition describes is what's to be explained. We want to explain why, say, Albania has a surprisingly large number of bunkers and a surprising love for Mercedes vehicles.⁴⁶ We don't want to explain anything about the proposition "Albania has a surprisingly large number of bunkers and a surprising love for Mercedes vehicles." If we wanted to explain something about the *proposition*, we might want to explain why, for example, the proposition "Albania has a surprisingly large number of bunkers and a surprising love for Mercedes vehicles" has fifteen words. We might want to explain how human beings grasp propositions (and how they realize what propositions mean). These are things about *propositions* that we might want explained. These are certainly things about *propositions* that philosophers of language have wanted to explain.

Again, Williamson (195) writes that:

... the sensation in my throat is evidence for the conclusion that I am getting a cold in the sense that the hypothesis that I am getting a cold would best explain why I have that sensation in my throat. The evidence to be explained is that I have *that* sensation in my throat—not just that I have a sensation in my throat.

For sure: that I have *that* sensation in my throat is a worldly conglomeration of relations and objects; to point to *that sensation* in my throat is to point to a specific conglomeration which is to be explained. The hypothesis that I am getting a cold is a hypothesis about me and the world. What it says about me and world is what explains *that* sensation in my throat. A *hypothesis*, which is a kind of quasi-linguistic object, doesn't explain things of this sort. It's what a hypothesis *is about* that explains things of this sort.

We *do* often slide from speaking of the content of propositions "to be explained" to the propositions themselves "to be explained" when we

⁴⁶I typed "surprising facts about Albania" into Google. I got these, among other things.

still mean what the proposition describes and not the proposition itself to be the target of explanation. So we often speak of propositions being explained, where the explanatory focus isn't on the actual propositions but on what they're *about*. Imagine, for example, that someone writes a shocking proposition on a blackboard. An audience member can ask pointedly, "Can you explain that proposition?" A response, however, won't focus on the proposition *itself* but instead about what the proposition is *about*. (That, yes, standard economic theory predicts X but in fact Y invariably occurs.) Philosophers shouldn't be taken in by the fact that we sometimes speak of propositions when we're talking about what those propositions purport.

I wrote in the last paragraph about someone writing a shocking proposition on a blackboard. "Shocking proposition": Here what's meant—and this is usually so with phrases like this—isn't that *the proposition* is shocking (whatever would that mean—that it's written in bright candy-cane colors?) but that what the proposition *says* is shocking. Notice that the points made in Section 6 are coming up here. What's shocking, say, is that America is no longer a democracy. We can say (instead) that the proposition "America is no longer a democracy" *says* something shocking. That is, the proposition itself has the property of saying something shocking; and thus, we can slide to saying that "America is no longer a democracy" is a shocking proposition. Via the semantics of "America is no longer a democracy"—its property of saying a particular thing—we can attribute the property of shock directly to that proposition. The derived property of a proposition is described (that it says something shocking); and this comes to the same thing as a property of a worldly conglomeration of objects and relations being shocking—that America is no longer a democracy. The further slide to directly saying the proposition itself is shocking is like a waitperson saying: "the ham sandwich wants his check", when what's meant isn't that a ham sandwich (on a plate) would like the check, but that the *person* who ordered the ham sandwich would like the check.

Williamson (2000, 195) writes: "One can use an hypothesis to explain why A only if one grasps the proposition that A. Thus only propositions which one grasps can function as evidence in one's inferences to the best explanation." But, of course, what only follows from the first sentence quoted is this: Only propositions which one grasps can

function to convey what the proposition is about which in turn can function as evidence in one's inferences to the best explanation. If we don't grasp the proposition then we don't understand what evidence that proposition describes. This truism isn't what Williamson needs.

Two factors contribute to the common slide from talking of the *relata* of propositions to the propositions themselves. First, we (relatively) rarely intend to talk about propositions in ordinary life; we (relatively) rarely intend to talk about language, that is. Certain specialists do often talk about propositions, of course, and intend to: Some philosophers (of language), for example, do this, linguists do also, and of course those who explicitly study literature also often talk about the propositions expressed by sentences or the sentences themselves. The rest of us almost never do this. So, talking about propositions and in that way really speaking of what the propositions are about doesn't lead to interference (because of possible misunderstandings) because the listener won't worry—won't even think to worry—that propositions themselves might be the actual topic of conversation.

Second, to focus on the *relata* of propositions is to focus on what I've been awkwardly describing as worldly conglomerations of objects and relations, or purported worldly conglomeration of objects and relations. But, recall why "purported" is important: we're often using sentences to convey *falsehoods*—deliberately or not. And, as Section 5 indicated, this is when things get subtle. What are we talking about in these cases? Just the same thing we're talking about when we're talking about Hercules or Donald Duck: Nothing at all. There is no worldly conglomeration of objects and relations that's the *relatum* of a false proposition—that's why the proposition is false. This doesn't mean, however, that we're instead talking about something else, the proposition! To assume this is to make the mistake Quine (1980) accuses McX of: confusing the idea of Pegasus with Pegasus. Some propositions describe worldly conglomerations of relations and objects—the true propositions do that. Some don't—the false ones. And the false ones, thus, describe nothing at all.⁴⁷ There is no word in English that applies to the *relata* of false propositions as "fact" does to true propositions. As I mentioned before, I'm tempted to use "phact"—where phacts are the *relata* of false propositions, but I've held

⁴⁷Or, if you wish, "John has taken up running again," describes John—who exists—falsely as running again. Putting the matter this way doesn't affect the argument here.

back. Without this jargon, we naturally fall back on sliding between talk of what a proposition purports to be about to the proposition itself—because we're in the grip of the mistake Quine points out.

Another point about this should be made. General propositions and singular propositions are both kinds of propositions—quasi-linguistic items. But although it's easy to describe what a singular proposition *describes*, that John is running again, for example; it's often not easy to describe what a general proposition describes. Consider "All men are mortal": does that describe a property all men have? Maybe that's a relatively easy thing to say with respect to this generalization, but it's not easy when it comes to other generalizations, those about gravity, for example. We can take refuge in vagueness, characterize a generalization (when true) as describing "the way that the world is"—I tried something like that above. It's much easier to fall back on the propositional vehicle (which, after all, does describe what's doing the explaining) and talk about *that* as doing the explaining the same way that we describe a person who ordered the ham sandwich as a ham sandwich, and say that it wants the check.

Williamson's points about probabilistic reasoning are similar to what he writes about inference to the best explanation; and so they involve the same use/mention error. Williamson (196) writes: "But what has a probability is a proposition; the probability is the probability that . . ." But here, again, are the same set of issues: "that", and what follows it in the phrase above, is a that-clause and what's being described as having a probability isn't a proposition but what the that-clause purports to name.

The last theoretical function for "evidence" that Williamson raises seems to bring in a fresh issue. Williamson (196) writes:

More straightforward uses of evidence also require it to be propositional. In particular, our evidence sometimes rules out some hypotheses by being *inconsistent* with them. For example, the hypothesis that only males have varicose veins is inconsistent with much medical evidence. But only propositions can be inconsistent in the relevant sense. If evidence *e* is inconsistent with an hypothesis *h* in that sense, it must be possible to deduce $\neg h$ from *e*; the premises of a deduction are propositions. Moreover, the subject who

deduces $\neg h$ from e must grasp e .

Although it may not look that way, we've been here before. The *mechanism* by which we recognize that a purported *situation* is inconsistent with another situation or with what a hypothesis describes is via propositions that describe the situation(s). So, of course, an inference drawn via a proposition that describes the purported situation is the method by which an inconsistency between what a hypothesis describes and a situation is revealed. Consider the hypothesis that all frogs are happy. An unhappy frog is inconsistent with what that hypothesis purports to describe. An unhappy frog, however, isn't a proposition. It is, I hope this isn't news, a frog.

This is hardly to deny that propositions (or interpreted sentences) themselves don't have logical properties. They obviously do. They are inconsistent with one another; some propositions imply other ones, and so on. But this is entirely compatible with our talking of purported worldly conglomerations *also* being inconsistent with other ones or with hypotheses. Again: we talk about worldly conglomerations that don't obtain just as we talk about entities that don't exist. And so, to make sense of this doesn't require us to fall back on treating propositions as the objects of focus because they depict or describe state of affairs (or what have you) that don't obtain. No more need we fall back on words or concepts being the objects of concern when we wonder whether Donald Duck is more famous than Sherlock Holmes or not.

Notice how we speak very naturally of situations (or states of affairs, or whatever) *implying* other situations or states of affairs:

That John is running again implies that John is running.

That one plus one equals two and that two plus two equals four implies that one plus one plus one plus one equals four.

That John is healthy implies that he doesn't have cancer.

In each one of these cases, what's (grammatically speaking) being characterized, as in implication relations, are purported states of affairs—certainly *not* propositions. The mechanism by which we *exhibit* the implications between these described situations is, of course, via propositions—we exhibit the implications between these described situations via the logical structure of propositions or via the semantics of

the words that appear in the sentences that express these propositions. But the lessons of usage are clear: It's perfectly straightforward for purported worldly situations to imply other purported worldly situations, or to be inconsistent with them, and so on.

If this weren't the case, we would face a serious problem of explaining how logic can be applied to the world. The simplest story to tell is one that describes the (logical) structure of propositions, what it is about propositions that enables them to bear logical relations to one another, as matching the ways that purported worldly conglomerations of objects and relations bear logical relations to one another.

To get mildly technical for a moment—in order to develop the thought of the last two sentences a little: this is the lesson of Gödel's completeness theorem. There's a correlation between the logical structure of the sentences of a formal language and the way that possible states of affairs (models) are. If we understand the sentences of a formal language as interpreted, then what they're talking about are the models, subject to those sentences, of course, being interpreted by the models in question.⁴⁸

I'll make one last point about how use and mention can be confounded when talking about evidence. Just as we talk about the facts via propositions that describe them, we also talk about evidence via propositions that describe it. Leite (2024, 99) writes (as I quoted him earlier):

It is crucially important here that in ordinary life it is entirely acceptable to appeal to considerations about the world as evidence for or against other claims. . .

This is a slightly arch and formal way of speaking, but understood literally it sounds as if “considerations” can be evidence. To speak more precisely (less sloppily), however, is to recognize that “considerations”

⁴⁸Corey Dethier (email 6/26/2025) puts the point plainly: “While our logical rules operate on linguistic entities—we can't apply logical rules to (say) cats—that doesn't mean that cats are free from the laws of logic. . . So it's true that if we are going to derive conclusions (contradictions) from evidence, we're going to need to describe that evidence linguistically. But to then conclude that evidence is propositional is like concluding that the property 'bald' applies to Tim's name because if we're going to describe Tim as bald, we need some way of referring to him.”

point to evidence or bring evidence to our attention; it's not that they *are* evidence. Indeed,

*Those considerations are evidence.

is completely unacceptable. Leite (2024, 127) says further "... the considerations that you appeal to as evidence are true." Are considerations true or false? We *can't* straight out say:

*Those considerations are true.

Nor can we say:

*Those considerations are false.

Or

*The considerations you've raised are all true.

What we can say is:

The considerations you've raised are all valid.

where this pretty much means that the considerations in question are *relevant* or *pertinent*. We can't say:

*The evidence you've raised is all true.⁴⁹

nor,

*The evidence you've raised is all valid.

When we raise considerations we do so via propositions. But it doesn't seem that the considerations *are* the propositions which are provided (and therefore truth vehicles). Well, then, are they what the propositions *describe*? That doesn't sound right either. We don't say "John knows those considerations," or "the considerations in question are the facts." The problem, I think, is talk of "considerations" is derivative from consider-imperatives; the noun "consideration" doesn't refer or even purport to refer to anything. Recall that I wrote above,

⁴⁹This isn't because "evidence" is a mass term. Compare. We *can* say, "The meat you brought is all rancid."

Consider the hypothesis that all frogs are happy.

Here what's being suggested is that we consider a hypothesis, although in doing so, what we'll consider is what the hypothesis says, not the hypothesis itself. Not always, of course. It can be that one is considering something about something propositional, we might consider something about consider-imperatives for example. But we can also be asked to directly consider something that's non-propositional:

Consider a situation in which John is running again.

My suggestion: Speaking of "considerations" is derivative from these sorts of imperatives, which are open-ended insofar as one can be asked to consider *anything*: To say, "here are some considerations..." is to invite the listener to consider... But because "consideration" is a noun, one can mistakenly assume that there are things, considerations, that the attention of the listener is being drawn to, and then we can ask what those things are, propositions or what propositions describe, or worse (as Leite does) just assume they're propositional, and in fact truth vehicles.

Analogy: Suppose someone really tried to determine what "sakes" are on the basis of usage. The result would be flummoxing, just as it is if one seriously sets about trying to determine what considerations are. Also see [Anscombe 1981](#) where she describes the difficulties one gets into if one tries to determine the ontological status of a "direct object". I'm suggesting the problem is the same one.

13. JUSTIFYING AND JUSTIFICATION

Let us turn to "justify" and "justification".⁵⁰ To start, usage makes it clear that "evidence" and "justification" have entirely differently semantic roles. We can ask the following question:

Are we justified in thinking we have evidence for this claim?

We can't, however, flip this sentiment. We can't say,

⁵⁰Here, in part, I'm revisiting and amplifying my work from [Azzouni 2020](#), Chapters 7 and 8 and [Azzouni 2025](#), Chapter 5.

*Do we have evidence for this justification?

Or,

*Do we have any evidence for this justification?

Nevertheless, there is (as I said) substantial overlap. An agent who has a justification for a belief, and can offer it, *often* also has evidence for that belief. S can be justified in thinking that R committed the crime. S's evidence can be the bloody-looking clothes in a hotel room where that apparent blood seems to match the blood that's on R's hands. However, the evidence that justifies someone's belief can be misleading; for example, the clothes in the hotel room look bloody, but that evidence is misleading because the red stuff isn't blood. Someone seeing them, nevertheless, may still be justified in thinking that R committed a crime.

Justifications are justifications for beliefs, or for what an agent thinks, or for claims that an agent makes—not for what agents *know*. “What justifies that claim?” we can ask, or “What justifies that belief?” We never say, “What justifies that knowledge?” So what are the objects of justifications, that is, what exactly gets justified? Again, we should tread carefully. I claim (say) that John is running again. Suppose someone describes my claim that John is running again as ridiculous. Is it my *claim* that's ridiculous or is it that John is running again that's ridiculous?

Well, it depends. We might think that John is running again is ridiculous because John isn't healthy enough to run again—not yet anyway. So maybe it's that the purported worldly state of affairs is ridiculous. But, really? Isn't what's ridiculous is that one *thinks* John is running again or *believes* that John is running again? This is only one way of understanding John's running again being ridiculous. It can also be ridiculous because John is unhealthy and he shouldn't be trying to run again. In this second case, we can thus describe the worldly conglomeration of objects and relations as *itself* ridiculous because it includes John's intentions, and his intentions are ridiculous.

But in the first case, we're actually talking about the thought or the belief. Given what the thought is (a thought that John is running again) or what the belief is (the belief that John is running again), the thought itself or the belief itself is ridiculous. So in the case of a ridiculous-attribution, it can *either* be the belief or thought itself that's ridiculous or (in certain circumstances) it can be what the belief or thought is of that's

ridiculous. Either the belief or thought itself can be the target of the attribution of ridiculousness *or* it can be what the belief or thought is of that's the target of the attribution of ridiculousness. So, when it comes to something being ridiculous, there are two possibilities: the purported worldly conglomeration of objects and relations is itself ridiculous or the belief or language vehicle that describes that purported worldly conglomeration of objects and relations is ridiculous.

What about the case where we wonder if a belief is justified? What are we focused on in that case? This seems not to be open to both the possibilities that ridiculous-attributions are open to. If we ask if a belief is justified, what we're asking is whether a *belief that P* is justified. That's not to ask whether *P* is justified; indeed, asking that seems to make no sense. We can't, that is, ask this:

*Is that John is running again justified?

*Is it justified that John is running again?

The second item almost sounds acceptable. But it sounds that way only because we try to hear it as short for: "Is John justified *in* running again?" Consider: Is it justified that inflation is rising again? This sounds short for: Is there a good reason that inflation is rising again? Good reasons are explanations. And explanations, as we've seen, are provided by what's worldly—or purportedly worldly. (I've avoided weighing in on whether explanations have to be, as it's sometimes put, *factive*.)

We can ask,

Is it justified to believe that John is running again?

Or, more naturally:

Is your belief that John is running again justified?

So too, we can ask:

Is the proposition that John is running again justified?

But unlike earlier, notice, we can't eliminate the phrase, "the proposition"; the result of doing *that* is "Is that John is running justified?" which we've already noticed is unacceptable.

So, contrary to what was shown about “evidence,” earlier (in Section 10), what gets justified (or not justified) are *propositions*. Correspondingly, what gets justified (or not justified) are *beliefs*. This is striking, especially since it was shown earlier that the objects of knowledge *aren’t* propositions, but instead, what’s described *by* propositions. To repeat, what gets justified aren’t the *relata* of propositions or that-clauses; what gets justified are beliefs and propositions themselves. This, anyway, is what ordinary usage seems to dictate.

So, let us recollect the title of Gettier 1963, “Is justified true belief knowledge?” The answer seems to be this. A justified true belief is something that lives on the propositional (the intentionally-propriety) side of the proposition-world relation. But knowledge lives on the worldly side of this relation; knowledge is of conglomerations of purported relations and objects. We can know that John is running again, for example, when John *is* running again.

This doesn’t mean that an agent having a justified true belief that *P* can’t be a necessary and sufficient condition for that agent knowing *P*. But this, if it were the case, *wouldn’t be* to say that a justified true belief *is* (a piece of) knowledge. That *is* ruled out by the foregoing. And, of course (and in any case), what rules out an agent having a justified true belief that *P* being a necessary and sufficient condition for that agent knowing *P* are the Gettier cases that Gettier discovered.

14. CONEE AND FELDMAN ON EVIDENCE

Recall that Conee & Feldman (2008, 83) write:

S is justified in believing *p* at *t* iff S’s evidence at *t* on balance supports *p*.

This isn’t to necessarily identify justifications as a kind of evidence (“sufficient evidence”). It can be understood as claiming that an agent having sufficient evidence for *p* is necessary and sufficient for that agent being justified in believing *p*. This necessary-and-sufficiency thesis doesn’t slur over use and mention so if it’s to be rejected, it must be rejected on other grounds.⁵¹

⁵¹Conee and Feldman mostly don’t slur over use and mention to establish their thesis. At one point, however, they write (2008, 87), “Sometimes a person cites one belief as

Conee & Feldman (2008, 86) quote Austin (1962, 115), and I'll do so as well:

The situation in which I would properly be said to have *evidence* for the statement that some animal is a pig is that, for example, in which the beast itself is not actually in view, but I can see plenty of pig-like marks on the ground outside its retreat. If I find a few buckets of pig-food, that's a bit more evidence, and the noises and the smell may provide better evidence still. But if the animal then emerges and stands there plainly in view, there is no longer any question of collecting evidence; its coming into view doesn't provide me with more *evidence* that it's a pig. I can now just *see* that it is, the question is settled.

Conee & Feldman (2008, 86) comment: "We endorse the view that Austin denies here. We think that seeing the pig does provide additional evidence. Indeed, it is this additional evidence that settles the question."

To maneuver the view they endorse, they distinguish between *scientific evidence* and *justifying evidence*. Scientific evidence, they say (84), "is publicly available". I take "scientific evidence" to be how we *ordinarily* use "evidence". On the other hand, for an agent to have *justifying evidence* for *P* is for them to have the appropriate background beliefs to enable them to see that the evidence in question is evidence for *P*. Furthermore, Conee and Feldman's claim about sufficient evidence being necessary and sufficient for justification is one about what they call "justifying evidence," not "scientific evidence". Indeed, they write, about a case where a detective has fingerprints that are evidence that "Lefty" was at the scene of the crime, but who doesn't know those fingerprints are Lefty's that:

One could restrict use of "evidence" to scientific evidence and say that the detective does have evidence [that Lefty was at the scene of the crime]. But then one would have to say that merely having evidence supporting a proposition

a reason for another. We take evidence to be what provides epistemic reasons. Thus it may seem that one belief can be evidence for another belief. We see no strong reason for resisting this claim..." This is to slur over use and mention; I discuss this explicitly shortly.

is not enough to make one at all justified in believing that proposition. It is clearer and simpler to preserve the connection between having supporting evidence and justification, and thus to say that merely knowing the characteristics of the fingerprints but not having the connecting information is not having evidence for the conclusion.

But consider how we normally speak. We do say of a detective in this case that he had in his possession evidence that Lefty was at the scene of the crime, but that he didn't realize that—recall the distinction (Section 10) between having evidence for *P* but not knowing it's evidence of *P* and otherwise knowing that the evidence is evidence for *P*. Detective novels regularly describe situations in which there is evidence of such-and-so but that it takes a Sherlock Holmes to recognize that the evidence is evidence of *P* or *Q*. Two points, therefore. First, Conee and Feldman's use of "clearer and simpler" indicates that they're engaged in conceptual engineering. They're recommending a change in usage vis-à-vis our ordinary talk of evidence. But, second, their justification for this revision seems *only* to be that in this way a connection between evidence and justification—their thesis about justification and evidence—is "preserved".⁵² That is, understanding "evidence" their way enables their thesis between evidence and justification. I suggest that we need more of a reason for a revision of our ordinary use of "evidence" than this.

Conee and Feldman, further, distinguish between *ultimate evidence* and evidence that isn't ultimate, writing (87), that "we hold that experiences can be evidence, and beliefs are only derivatively evidence." And later (88) "... all ultimate evidence is experiential."

Again, one can legislate, change what "evidence" refers to. But let's first get clear about how we use "evidence". Conee and Feldman write (87):

Some philosophers have argued that only believed propositions can be part of the evidence one has. Their typical ground for this claim is that only believed propositions can serve as premises of arguments. Our view differs radically from this one. We hold that experiences can be evidence,

⁵²But "preserved" is the wrong word. The right word is "forged".

and beliefs are only derivatively evidence. Examples intuitively support that we have experiences as evidence. Your evidence for the proposition that it is warm where you are typically include your feelings of warmth, your evidence for the proposition that you are frustrated by being stuck in the heat in a traffic jam typically includes your feeling of warmth, your evidence for the proposition that the car in front of yours in the traffic jam is red typically includes your visual evidence of how the car looks, and so forth. It is not just other propositions that you believe that contribute to your justification. The experience itself contributes.

This description of how we ordinarily use “evidence” is hardly intuitive (if usage is any guide to “intuition”, anyway). Recall the quotation from Austin. What’s natural, once the pig trots into view, is that “evidence-gathering” is at an end. We don’t need evidence that a pig is thereabouts because the pig has come into view. And Austin indicates further that usage is on his side in a footnote on the same page (and the next), writing: “I have, it will be said, the ‘evidence of my own eyes’. But the point of this trope is exactly that it does *not* illustrate the ordinary use of ‘evidence’—that I *don’t* have evidence in the ordinary sense.” Exactly right. And consider a situation in which someone is denying to you that a pig is about. You’ve noticed the pig-food, the smell, etc. “No, no,” the person protests, “it’s skunks, weird-smelling skunks. Your evidence is misleading.” And then the pig trots into view. “Well,” you say, “if that’s not evidence of a pig, then I don’t know what is.” This is a *joke*. And what makes it a joke is precisely that one no longer needs *evidence* that a pig is about. Denying that a pig is about is futile.⁵³

The other “intuitive” examples that Conee and Feldman describe, my “evidence” that the car in front of me is red (when I see it), the “evidence” that the temperature is high (because I feel warm) are all *not* examples of what we would ordinarily describe as evidence. Importantly, though, they are all cases where the speaker would be described as *justified*. I’m justified in thinking the car in front of me is red (when I see it) precisely because I do see it. I’m justified in thinking the temperature is high

⁵³Well, maybe not. American politicians, these days especially, often deny that a pig is about when the elephant-sized pig is right in the room in front of them (and in front of everyone else). *And they get reelected anyway.*

because I feel hot. So these are cases where we take an agent to be justified in their beliefs even though they may not have evidence.^{54,55} In the quotation from Conee and Feldman above, notice they write: “It is not just other propositions that you believe that contribute to your justification. The experience itself contributes.” This is absolutely right. Unfortunately, no one normally describes either what’s believed or experienced as evidence, even though both are understood to provide justifications.

Conee and Feldman’s coinage, “ultimate evidence”, is an attempted revision of the ordinary word “evidence” with the aim of aligning our ordinary talk of *justification* with our ordinary talk of *evidence*. Our uses of these words come apart, however, in just the kinds of cases that Conee and Feldman understand “ultimate evidence” to apply to. It’s natural to say, “That I feel hot justifies my belief that it’s hot out.” It’s just not natural to say, “That I feel hot is evidence that it’s hot out” except in special circumstances—where one has reason to distrust one’s senses. At that point, one might describe, in a tentative way, that one’s sensations are evidence. But just as likely, one won’t. For example, if someone realizes, because of a disease they have that when they feel hot it’s sometimes the case that it’s actually not hot out, they generally won’t demote their experience of feeling hot to being only evidence of it being actually hot, or just a clue of that.

This reveals another discrepancy between the ordinary notion of evidence and the one that Conee and Feldman are creating. Evidence is never “ultimate”, or “final.” It’s almost always circumstantial, in the sense that there’s an implicit contrast between the evidence we have for something having happened, for example, and what makes that it happened *definitive*. This is the point that Austin is making in the quotation given of his. Seeing the pig under normal circumstances isn’t evidence of a pig precisely because it’s definitive.

Last point about all of this. As our practices with testimony make clear, although a person’s experiences (and memories) aren’t evidence *for the person having those experiences (and memories)*, they are evidence for a third party. This exchange is very natural: “John murdered Sam.”

⁵⁴Conee and Feldman also describe memories as “evidence”. I omit discussion of this; the same points can be made that I’m making about experiences generally.

⁵⁵With the following rider: This is how we ordinarily speak when we’re not sceptical-scenario raising philosophers. See Section 15 for further discussion.

“What’s your evidence?” “Sally saw him do it, and Sally doesn’t lie.”

15. SCEPTICAL ARGUMENTS

How do the differences between how we talk about justification and how we talk about evidence bear on certain sceptical arguments? Let us consider a modified version of one of Gettier’s original definitions of knowledge:

- (c) S knows that *P* IFF
- (i) *P* is true,
 - (ii) S believes that *P*, and
 - (iii) S has adequate evidence for *P*.

I’ve modified Gettier’s (b), quoted in Section 3, in order to focus on how evidence and justification play against each other; so I’ve replaced “accepts” with “believes” (and I’ve changed the order of the clauses as well). Let us drop the aspirations of (c) to be a definition or a set of necessary and sufficient conditions, and instead regard it only as a list of necessary conditions; indeed, let us primarily focus only on (i) and (iii) being necessary conditions.⁵⁶

The focus, here, is on the fact that both Leite and Williamson address Cartesian sceptical arguments through the lens of “evidence” rather than through the lens of “justification” (although neither of them realizes that the shift in terminology might make a difference to the arguments they’re considering). However: Because evidence is worldly, what evidence an agent has changes, depending on whether that agent is in a sceptical scenario or not. Let us turn to this, and see how it bears on knowledge. Specifically, let us focus on what role adequate evidence versus justification plays vis-à-vis knowledge when considering Cartesian sceptical scenarios.

A preliminary that I’ll note and then set aside: hallucination scenarios—when deployed by sceptics—should be taken seriously, and in a way that the dreaming argument, the evil demon argument, or even brain-in-a-vat scenarios shouldn’t be taken seriously. This is because hallucinations, as knowledge-challenging possibilities go, are simply

⁵⁶Belief isn’t a necessary condition for knowledge; see Azzouni 2020 and Azzouni 2025.

better because they're *real* possibilities. By this I mean that there are drugs and medical conditions that induce lucid visual hallucinations—that is, hallucinations that are sometimes rather global in nature but which one can learn to recognize (under *most* circumstances) are hallucinations.⁵⁷ I'm not convinced we're anywhere near brain-in-the-vat possibilities, but I'll focus on these instead of hallucinations for the sake of exposition (that is, because they're the sceptical challenges that are most familiar to philosophers). So we're assuming that being a brain in a vat is a genuine possibility, and we're evaluating the question of how this possibility—when taken to be genuine—affects the evidence and the possible justifications of the agents for what they believe in such situations.

So let us imagine that two agents, otherwise identical in what they experience, are, in one case a person S living an ordinary life (the sort of lives most of us think we're living) and in the other a person S_V who is a brain in a vat. We can say the following: S has a whole lot of evidence for their beliefs that S_V doesn't have. For example, S has a lot of evidence for their belief that they have hands. For one thing, they *have* hands. But there's a great deal of evidence, apart from that, as well. S does (and can do) all sorts of things because they have hands. All of this is worldly, and thus can be described in evidential terms. S_V , on the other hand (sorry about that joke), doesn't have any of *this* evidence: they haven't hands, and indeed, almost anything else that amounts as evidence for S having hands isn't shared by S_V .

What *do* S and S_V have in common? Well, there's the rub. What they have in common are their experiences, and only their experiences. That, after all, is how the thought experiment is set up. Do they therefore have *evidence* in common? As I said, the evidence in common (if it is evidence) can only be what the two of them are *experiencing*. We can try to put it this way. Some of the evidence for S and S_V having hands is that it *seems* to them that they have hands. *That evidence*—if it's evidence—is shared by S and S_V .

We never speak of seeing our hands as *evidence* that we have hands;

⁵⁷ See Azzouni 2025, Section 7.3. I also especially recommend the eminently readable Sacks (2012). An example: Sacks quotes a correspondent: "In the presence of white noise such as running water or a central air conditioning system, I frequently hear music or voices. I hear it distinctly (and in the early days, often went searching for the radio that must have been left on in another room)..." (70).

relatedly, we never speak of our experiences as evidence except in unusual cases—contrary to what Conee and Feldman suggest. Consider someone saying, “I’m not angry.” The respondent might say, “my evidence that you’re angry is that you’re throwing things around.” The original claimant won’t say: “Well, *my* evidence that I’m not angry is that I don’t *feel* angry”—unless they’re being sarcastic. (They might say this angrily, and the respondent may use their angry speech act as further *evidence* that they’re indeed angry.)

According to (c), S_V doesn’t know that they have hands. Furthermore, according to (c), S does know they have hands because, after all, S doesn’t just have evidence they have hands; they clearly have *adequate* evidence for this. So there’s an evident asymmetry here, due to the asymmetry in the *evidence* that’s available to both agents. And this asymmetry, given that evidence relates to knowledge, along the lines of (c), leads to the view that these agents differ in what they know. Thus, even if being a brain in a vat is a real possibility for an agent, as long as that agent isn’t *actually* a brain in a vat, that agent’s knowledge won’t be undercut because they have evidence that their twin S_V doesn’t have.

Both Leite and Williamson are tempted by an argument in this neighborhood. If justification is assimilated to evidence the way that both of them are also tempted to believe, then these considerations look even stronger: S is justified in believing they have hands; S_V isn’t. And thus, even if we demote (c) to only providing necessary conditions for knowledge, we still get the result that S_V doesn’t know they have hands on grounds that don’t apply to S .

Knowing, as we do now, that justification and evidence aren’t to be assimilated, should we be comfortable in thinking that S is justified in thinking they have hands? Again, assuming that being a brain in a vat is a real possibility, here are some considerations against the idea that S is justified. Recall from Section 13 that although we can ask “Are we justified in thinking we have evidence for this claim?” we *can’t* ask “Do we have evidence for this justification?” And this bears directly on the scenario that S might be in. For, given the possibility that S might be a brain in a vat, and given that such a possibility is *close*, we can say (quite naturally) that S isn’t justified in thinking they have the evidence they think they have.⁵⁸ *Justification trumps evidence*, precisely because

⁵⁸Williamson agrees that the nature of the evidence we have isn’t transparent to us:

evidence is worldly, and what's worldly is precisely what's challenged in sceptical possibility-scenarios. The result: S isn't *justified* in thinking they have hands although they've got plenty of evidence that they've got hands. And so, the sceptic argues, S doesn't know they have hands because they haven't ruled out a nearby possibility. Conclusion: Whether they have adequate evidence or not for what they believe simply doesn't bear against Cartesian scepticism.

16. SOME METHODOLOGICAL MORALIZING

Until now, I've been describing the use/mention confluences surrounding our talk of "justification" and "evidence", and illustrating a general perspective about the role of philosophy (and specifically, the philosophical topic of epistemology) in our ordinary intellectual lives. Leite (2024, 10) opposes this perspective, arguing that there's a "pre-philosophical" stance that our ordinary epistemic lives occur within. This stance includes "epistemic claims and principles to which we are committed in ordinary life and science." He takes it that sceptical arguments can be challenged successfully from within this pre-philosophical stance—that is, that we can refute sceptical challenge without having to develop epistemological theories to enable us to do this. Instead:

We will carefully scrutinize skeptical arguments to see what requirements they impose, and we will consider whether those requirements are among the commitments of ordinary life and science. To do this we might do such things as articulate epistemic claims and principles to which we are committed in ordinary life and science, reflectively describe those commitments and what they involve, reflectively describe aspects of our ordinary epistemic practices, utilize clear examples to show that particular requirements or principles are not correct, draw certain kinds of generalizations from such examples, put two and two together to reach a more general description, summarize the results of all this, and the like.

evidence can be misleading without our realizing that it is. That this plays against assimilating evidence to justification isn't realized by him.

But wait—don't these instructions tip us into philosophical theorizing? (They actually do, as I illustrate in the rest of this section.) Let's start with following sceptical challenges:

Perhaps you are a brain in a vat.

Perhaps an evil demon is fooling you.

Perhaps you are dreaming.

Pre-philosophical responses look like this:

I know I'm not a brain in a vat because brain-in-a-vat technology doesn't exist.

I know I'm not being fooled by an evil demon because I know there are no evil demons.

I know I'm not dreaming because my dreams aren't like waking life. Dreams like waking life aren't possible.

Here we're (supposedly) using our ordinary understanding of what's possible and what's not possible to reject the possibilities that the sceptic is raising.⁵⁹

My counter to this way of trying to respond to the sceptic is that there is no stable pre-philosophical position—not with respect to our “ordinary” epistemic methods and (for that matter) not with respect any aspect of our ordinary intellectual life. So what I'm claiming here is that philosophical theorizing (if we're moderately thoughtful) invades our thinking about, nearly enough, anything.

The reason that there is no stable pre-philosophical position is because where we begin is in a human-created world of language, concepts and methods—the “things” we've learned. This is a conceptual structure that we grow up in (grow into). Key is something that (Quine 1960, 3)

⁵⁹This is a common interpretation of Moore's responses to scepticism in Moore 1939 and Moore 1959. The idea he's taken to share with Leite (on this interpretation) is that Moore, from within the “commonsense” framework, responds to the sceptic by utilizing—and only utilizing—our ordinary epistemic practices. Moore (1939, 145-146) looks like he's illustrating this point when he compares his proof that he has hands with how one shows there are misprints on a page of a book. I don't interpret Moore this way but as instead as making a series of burden-of-proof moves against the sceptic. See Azzouni 2025, Chapter 2.

noticed when he wrote: “Neurath has likened science to a boat which, if we are to rebuild it, we must rebuild plank by plank while staying afloat in it. The philosopher and the scientist are in the same boat.”

The point I want to stress here is that where we begin is dynamic, and philosophy (recognized or not) has been involved in the changes we find ourselves having to make for a very long time. A related point is that even when we don’t need to make changes in how we do things “pre-philosophically,” we still often find that we must ascend to philosophizing in order to understand what’s going on.

What drives us to philosophize is, often, something we stumble across in ordinary life, which raises puzzles which we try to handle the way that Leite suggests (“utilize clear examples to show that particular requirements or principles are not correct, draw certain kinds of generalizations from such examples, put two and two together to reach a more general description, summarize the results of all this, and the like”). What drives us to *philosophy* is always the same thing: no easy resolution of the complex oddities of our ordinary practices is obvious. *That’s* when philosophy shows up.⁶⁰ (And sometimes, when the work of philosophy is done, that it was needed in the first place becomes invisible because we have to work at remembering our history—for obvious reasons.)

So consider something *utterly basic*. We often try to recognize what’s true and what’s false. Which propositions, that is, are true and which are false. It’s also clear that we don’t think on occasion: maybe some propositions are *both* true and false. That’s not a possibility for us. So putting 2 and 2 together about our practices in science and in life gives us what looks like a pretty obvious example of an ordinary general principle, in science and in life, about truth:

Every declarative sentence is true or false and not both.

Oops, what about “The sentence appearing between these very quotation marks is false”? What about (1)?

⁶⁰Thus it’s no surprise that philosophy has had such a presence in our collective attempt to understand the family of quantum-mechanical theories—right from their origins and continuing until today. Attempts at this aren’t circumscribed within academic boundaries: members of physics departments *and* members of philosophy departments are both at work on this.

(1) (1) is false

(And, of course, there are many other ways to produce paradoxical sentences like this.)

Philosophers have worried about these paradoxes since the beginning of philosophy, through the Middle Ages, and up until now. Why should they worry? Or rather, since not all philosophers worry, I'm asking, rather, why should *any* philosopher worry? Or anyone else? Well, because there's an apparent problem right on the surface of our ordinary ways of speaking about what's true and what's false. Those ways of speaking lead to contradictions. And we know contradictions are bad.⁶¹

The rest of my examples will be drawn from epistemology.⁶² Leite, as quoted earlier, assumes that it's relatively transparent what our pre-philosophical practices are. Correspondingly, he assumes that the usage properties of "know(s)" are relatively transparent as well. Nothing about how we talk about who knows what and who doesn't generates questions that lead to philosophical theorizing on Leite's view. He also assumes that the generalizations that govern epistemic language are easy to reach for and evaluate. No "philosophical theorizing" is called for to get to them.

That this is evidently false I hope has been made clear by the previous sections of this paper. How "evidence" and "justification" and "knowledge" are used, and how they are (and aren't) semantically intertwined isn't obvious and can trip up the unwary. Consequently, exactly what

⁶¹Why are they bad? Because they trivialize whatever we say by making everything follow? Actually, no. That's too sophisticated. They're bad because it needn't be transparent from a version of a contradiction how the contradiction is being used—what results theorems, concepts, etc., are being drawn from the contradiction that contradict other theorems, results, etc. that are also being drawn from the contradiction. Call this the *spread* of a contradiction. Sometimes a contradiction hasn't much spread. The reason that so much set theory survived the fact that naïve set theory was contradictory was because the spread of the actual contradiction wasn't wide. Unfortunately, and this is connected to Church's result about there being no decision procedure for consistency, there is no general method for measuring the spread of any contradiction. And the liar paradox, at least at first sight, should produce alarm because it relies on such a basic and widespread notion, "true". As it turns out (shown first by Tarski) liar paradoxes don't have much spread either. The point is that none of this can be shown or resolved "pre-philosophically": what it requires is quite sophisticated.

⁶²But examples show up *everywhere*. See Azzouni 2024 for discussion of how often conceptual puzzles arise in what's here being called pre-philosophy, and how widespread they are.

the epistemic principles are that govern these words isn't obvious either. I go on to give several more examples, all of which are in the epistemic literature.⁶³

You look out the window and see a bird. Here's a possibility: It's not a real bird but a stuffed bird. Leite (2024, 3) writes:

Such worries fall by the wayside once we recognize that fallibilism is true. As is shown by our best scientific practice as well as ordinary life, we can know things even if our evidence is compatible with various possibilities of error.

Is fallibilism—the view that our knowledge of something P is compatible with the possibility of $\neg P$ or with $\neg P$ having a nonzero probability—to be read easily off our best “scientific practice as well as ordinary life”? Well, Newton didn't think so. Newton was an infallibilist about knowledge.⁶⁴ Indeed, pretty much every epistemologist was an infallibilist until after World War II. G.E. Moore was an infallibilist (he couched his infallibilism in terms of “certainty”). Fallibilism in epistemology then emerged, and relatively suddenly. Malcolm thought “know(s)” was ambiguous—two meanings: a fallibilist one and an infallibilist one. Austin, however, was a militant fallibilist, and by the time of Gettier 1963, pretty much every epistemologist was. No one thinks any longer that “know(s)” is ambiguous. It's a good (historical) question why this happened. Certainly it can't be because fallibilism is just obvious, given our usage in science and ordinary life.

Why did it take so long to see that fallibilism is compatible with how we speak and our science; why is it so hard to see what's “shown by our best scientific practice as well as ordinary life? Because it's not obvious, not even a little bit—fallibilism isn't something that's easily “shown” by how we talk.⁶⁵

⁶³One family of examples I don't give involve probability. It's striking, or ought to be, that probability is crucial to pre-philosophical knowledge gathering, and yet we are (collectively and as individuals) completely awful at probabilistic reasoning. That's why casinos are always a great investment. (This is advice for those unsavory oligarchs who are looking for somewhere to park their money.)

⁶⁴See Azzouni 2023 for discussion.

⁶⁵One problem (discussed in Azzouni 2020) is that it's hard to even express fallibilist sentiments about one's own knowledge easily. Consider “I know P but P could be false.” This just sounds wrong, and here I'm making a point about ordinary usage. That's why

Let us consider another example, lotteries. An advertising slogan that was used for many years: “You never know.” This slogan strikes most of us as *right*. That is, our ordinary epistemic practices license the thought that you never know you’ve got a losing ticket, as long as the possibility that your ticket is a winning ticket has nonzero odds.⁶⁶

So here are some easy steps to *Lottery Scepticism*:

I think I’m going to finish typing this sentence before the ceiling above me collapses and kills me in the next half hour.

There’s a nonzero probability that I’ll win a lottery even though the odds of my winning are lower than the ceiling above me collapsing and killing me in the next half hour.

Therefore: I don’t know that I’m going to survive the next half hour because I don’t know the ceiling above me will collapse on me and kill me because I never know that I don’t have a winning lottery ticket regardless of the odds of it winning.

Recall that Leite was quoted above as writing:

We will carefully scrutinize skeptical arguments to see what requirements they impose, and we will consider whether those requirements are among the commitments of ordinary life and science. To do this we might do such things as articulate epistemic claims and principles to which we are committed in ordinary life and science, reflectively describe those commitments and what they involve, reflectively describe aspects of our ordinary epistemic practices, utilize clear examples to show that particular requirements or principles are not correct, draw certain kinds of generalizations from such examples, put two and two together to reach a

David Lewis (1999, 419-420) writes:

If you are a contented fallibilist, I implore you to be honest, be naïve, hear it afresh. ‘He knows, yet he has not eliminated all possibilities of error’. Even if you’ve numbed your ears, doesn’t this overt, explicit fallibilism *still* sound wrong?

⁶⁶This is a pretty good characterization of infallibilism, by the way.

more general description, summarize the results of all this, and the like.

There doesn't seem to be anything in the steps to Lottery Scepticism that go beyond the pre-philosophical position, given how Leite has described it. A solution to Lottery Scepticism—if there is one—will require philosophy, specifically an analysis of what it is about how “know(s)” works, that allows such an easy argument for Lottery Scepticism to occur, as well as a philosophical theory about knowledge to show that this easy argument is unsound.⁶⁷

I've been pointing to aspects of how we ordinarily understand knowledge that point towards infallibilism; and leaving implicit the aspects of how we ordinarily understand knowledge that point towards fallibilism. The latter, of course, arise because of all the ways that we ordinarily attribute knowledge to ourselves and others that we can clearly be mistaken about. But ordinary usage actually points towards a rather extreme fallibilism that's as bewildering as the ways that usage points towards infallibilism. Here's an example due to Hawthorne (2004, 68-69):

I give six children six books and ask them each to pick one of the books at random. All but one contains misinformation about the capital of Austria. I ask the children to look up what the capital of Austria is and commit the answer to memory. One child learns 'Belgrade', another 'Lisbon', another 'Vienna', and so on. I ask an onlooker who has witnessed the whole sequence of events (or someone to whom the sequence is described) 'Which one of the schoolchildren knows what the capital of Austria is?'... It is my experience that those presented with this kind of case will answer, not by saying 'None of them', but by selecting the child whose book [reads] 'Vienna'—even though that child was only given the answer by luck.

Indeed, the extreme liberality with which we're willing to attribute knowledge to someone (like in this case) makes Hawthorne (2004,

⁶⁷ See Azzouni 2020, Chapters 10 and 11. On my view what's required is a characterization of “know(s)” that treats it as simply lexically vague. This, I claim, is what's needed to respond to lottery scepticism: full-blown epistemic theorizing.

68, footnote 48), “despair of the traditional epistemological project of extracting an analysis of knowledge from reflecting on use.”

Note what this despair amounts to. Contrary to what Leite says, there’s nothing obvious or even reasonable-looking about our ordinary practices of knowledge attribution.

What is it about our purported pre-philosophical epistemic practices that makes them require philosophical intervention? There are (in my view) two factors that destabilize the pre-philosophical stance (in *all* walks of life and in *all* the sciences)—these two factors make philosophizing *essential* to all walks of life (unless someone insists on just muddling through without worrying too much about it).

First, the practices and language of the purported pre-philosophical position aren’t *transparent*. (This point is stressed by the later Wittgenstein.) Is our language really trivially inconsistent, as the liar paradox seems to indicate? Are there practically no constraints on a child knowing something as Hawthorne’s example seems to indicate? And our epistemic practices—are they infallibilist or not? *None of this is obvious*.

To figure any of this out, one has to theorize. Some of the required theorizing is linguistic. Some of it is about methodology. All of it involves (or soon involves) philosophical theorizing. It certainly involves what philosophical theorizing looks like in the contemporary setting.

We can’t take for granted what our contemporary epistemic practices look like. (Leitean complacency will not help us.) For one thing, epistemic practices (like all our practices—recall Neurath’s boat image) are a moving target: they evolve. (And sometimes under the pressure of philosophical analysis or insight.)⁶⁸

Second: There are flaws, defects, and infelicities in how we speak and in our practices—both epistemic ones and conceptual ones. So what’s sometimes needed is *conceptual engineering*. One sometimes has to replace or modify what we inherited. (A lot of philosophy, *knowingly or not*, is engaged in that. It’s how Quine, long ago, understood philosophy.)

For the record, I think we need to look at usage carefully to see how we get tripped up when we talk about what we know and what we don’t, and how we handle challenges to our knowledge claims—sceptical ones and even ordinary ones. I don’t think any repairing or modifying of how

⁶⁸Recall note 63. Not only the mathematics needs to be worked out in order to get clear about probabilistic reasoning; philosophical issues must be resolved as well.

we talk about knowing or knowledge is called for (Azzouni 2020, 2025). But this process (of being reassured about our epistemic practices) calls for a lot of philosophical theorizing—along with linguistics, philosophy of science, science itself, and so on. (*Philosophy* is a special science.)

Also, for the record, I think liar’s paradoxes, as they emerge in ordinary language, are another matter entirely. Natural languages really are trivially inconsistent (Azzouni 2006). We thus need to engage in a little conceptual engineering, as well as explaining how we can successfully function in a language that’s trivially inconsistent. Again: philosophy is called for—along with linguistics, logic, and other tools. (*Philosophy* is a special science.)

17. WHAT’S SURPRISING

I want to close this paper by describing, in a cursory way, some surprising elements of the analysis of justification and evidence in the first fifteen sections of this paper. First, issues with use and mention have been heretofore pretty uniformly restricted to topics in logic, philosophy of language, and metaphysics—mainly in relation to how we psychologically handle referring to the nonexistent, but also about its widespread use in mathematical exposition.⁶⁹ It was a surprise (to me) that it comes up extensively with respect to propositional-attitude attributions and in such fairly complicated ways; and that it plays out differently for “justification” and “evidence”. (I didn’t see this coming.)

In any case, it’s hard not to notice that use/mention errors are ubiquitous and hard to avoid. I’ve noticed they’re something students have trouble with, and non-philosophers too. And here I’m talking about the relatively simple phenomenon of nouns that don’t refer (that, for example, “Pegasus” doesn’t refer to the idea of Pegasus). But the use/mention errors I’ve been describing in this paper are far more subtle and complicated, especially because sometimes we really do mean, even in ordinary discourse, to be talking about propositions and not what they describe. The use-and-mention back-and-forth I’ve been describing in the course of this paper, to repeat, is far more subtle than anything

⁶⁹There is a surprising reference to use and mention by Chisholm (1964, 90). He complains that the distinction hasn’t always been respected by epistemologists. But this reference is cryptic as well as lonely (it’s not clear what Chisholm is complaining about), and, in any case, the remark doesn’t seem to have been taken up by anyone else.

that happens with Pegasus. No wonder it's been overlooked and no wonder it's done so much philosophical damage.

I mentioned that it was also a surprise (to me) that something that's usually a matter of pure metaphysics and philosophy of language—what's going on when we (in ordinary language) refer to what doesn't exist—also shows up directly in the analysis of propositional-attitude talk. In philosophy anyway, I think these illustrate the benefits of being aware of neighboring debates and issues. Some mistakes have long tentacles.

ACKNOWLEDGEMENTS

My thanks to Eric Dean, Corey Dethier, Noelle Gentile and Jeffrey McConnell for suggestions, caveats, and corrections which improved this paper. My thanks as well to Jurgis Škilters and Daniel Skibra for creating the occasion for me to write this paper and also for useful suggestions and comments.

DISCLAIMER ON FOOTNOTES

The editorial policy of *The Baltic International Yearbook of Cognition, Logic, and Communication* specifies the use of endnotes. In the case of the present article, an exception has been made to use footnotes instead, reflecting the way the notes are closely integrated into the text.

References

- Anscombe, G. E. M. 1981. 'The intentionality of sensation: A grammatical feature'. In 'Metaphysics and the Philosophy of Mind', vol. II, 3–20. University of Minnesota Press.
- Austin, John. 1962. *Sense and Sensibilia*. Oxford University Press.
- Azzouni, Jody. 2004. *Deflating existential consequence: A case for nominalism*. Oxford University Press, <https://doi.org/10.1093/0195159888.001.0001>.
- . 2006. *Tracking Reason: Proof, Consequence, and Truth*. Oxford University Press, <https://doi.org/10.1093/acprof:oso/9780195187137.001.0001>.
- . 2010. *Talking about Nothing: Numbers, Hallucinations, and Fictions*. Oxford University Press, <https://doi.org/10.1093/acprof:oso/9780199738946.001.0001>.
- . 2013. *Semantic Perception: How the Illusion of a Common Language Arises and Persists*. Oxford University Press, <https://doi.org/10.1093/acprof:oso/9780199967407.001.0001>.

- . 2014. 'Freeing Talk of Nothing from the Cognitive Illusion of Aboutness'. *The Monist* 97, no. 4: 443–459. <https://doi.org/10.1093/monist/97.4.443>.
- . 2017. *The rule-following paradox and its implications for metaphysics*. Springer, <https://doi.org/10.1007/978-3-319-49061-8>.
- . 2018. 'Deflationist Truth'. In Michael Glanzberg (ed.) 'The Oxford Handbook of Truth', 477–502. Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780199557929.013.18>.
- . 2020. *Attributing Knowledge: What It Means to Know Something*. Oxford University Press, <https://doi.org/10.1093/oso/9780197508817.001.0001>.
- . 2023. 'Smith, Smith and Seth, and Newton on "Taking to Be True"'. In Marius Stan & Christopher Smeenk (eds.) 'Theory, Evidence, Data: Themes from George E. Smith', 1–19. Springer. https://doi.org/10.1007/978-3-031-41041-3_1.
- . 2024. 'A White Horse is Not a Horse: Markus Gabriel's "New Realism"'. In Jan Voosholz (ed.) 'Markus Gabriel's New Realism', 33–73. Springer. https://doi.org/10.1007/978-3-031-69526-1_3.
- . 2025. *Challenging Knowledge: How We Sometimes Don't Know What We Think We Know*. Oxford University Press, <https://doi.org/10.1093/9780197789650.001.0001>.
- . forthcoming. 'Scientific Pluralism and the Language-Based Approach to Scientific Theories'. In Maria Martínez-Ordaz & Otávio Bueno (eds.) 'From Contradiction to Defectiveness to Pluralism in Science: Philosophical and Formal Analyses', Springer.
- Chisholm, Roderick M. 1964. 'Theory of Knowledge'. In Roderick M. Chisholm, Herbert Feigl, William K. Frankena, John Passmore & Manley Thompson (eds.) 'Philosophy', 233–344. Prentice-Hall.
- Collins, John. 2022. 'The Null-Hypothesis for Fiction and Logical Indiscipline'. *Analytic Philosophy* 65, no. 2: 131–144. <https://doi.org/10.1111/phib.12274>.
- . 2025. 'World and Truth in Conflict'. In Richard Gaskin (ed.) 'The Question of Linguistic Idealism', 111–137. Oxford University Press. <https://doi.org/10.1093/9780191968617.003.0005>.
- Conee, Earl & Feldman, Richard. 2008. 'Evidence'. In Quentin Smith (ed.) 'Epistemology: New essays', 82–104. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199264933.003.0005>.
- Feldman, Richard & Conee, Earl. 1985. 'Evidentialism'. *Philosophical Studies* 48, no. 1: 15–34. <https://doi.org/10.1007/BF00372404>.
- Gettier, Edmund. 1963. 'Is Justified True Belief Knowledge?' *Analysis* 23: 121–123. <https://doi.org/10.1093/analys/23.6.121>.
- Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford University Press.
- Hume, David. 1977. *An Enquiry Concerning Human Understanding / A Letter from a Gentleman to His Friend in Edinburgh*. Hackett.
- Landau, L. D. & Lifshitz, E. M. 1976. *Mechanics*. 3rd ed. Elmsford, NY: Pergamon Press.
- Leite, Adam. 2024. *How to Take Skepticism Seriously*. Oxford University Press, <https://doi.org/10.1093/oso/9780197691175.001.0001>.
- Lewis, David. 1999. 'Elusive Knowledge'. In 'Papers in Metaphysics and Epistemology', 418–446. Cambridge University Press. <https://doi.org/10.1017/CBO9780511625343.026>.
- Moore, G. E. 1939. 'Proof of an External World'. In 'Philosophical Papers (1962)', 126–148. Collier Books.
- . 1959. 'Four Forms of Skepticism'. In 'Philosophical Papers (1962)', 193–222. Collier Books.

- Pritchard, Duncan. 2016. *Epistemic Angst: Radical Skepticism and the Groundlessness of Our Believing*. Princeton University Press.
- Quine, W. V. O. 1951. *Mathematical Logic*. Revised ed. Harvard University Press.
- . 1960. *Word and Object*. MIT Press.
- . 1980. 'On What There Is'. In 'From a Logical Point of View', 1–19. Harvard University Press, 2nd ed. Revised.
- Sacks, Oliver. 2012. *Hallucinations*. Vintage Books.
- Salow, Bernhard. 2018. 'The Externalist's Guide to Fishing for Compliments'. *Mind* 127, no. 507: 691–728. <https://doi.org/10.1093/mind/fzw029>.
- Shapiro, Barbara J. 2000. *A culture of fact: England, 1550-1720*. Ithica: Cornell University Press.
- Tarski, Alfred. 1983. 'The Concept of Truth in Formalized Languages'. In John Corcoran (ed.) 'Logic, Semantics, Metamathematics', 152–278. Hackett.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford University Press, <https://doi.org/10.1093/019925656X.001.0001>.
- Wootton, David. 2015. *The invention of science: A new history of the scientific revolution*. New York: Harper.