

The Baltic International Yearbook of
Cognition, Logic and Communication

April 2026

Volume 16: *Justification, Evidence, and
Use/Mention Errors*

pages 1-19

DOI: <http://dx.doi.org/10.4148/1944-3676.1133>

RAM NETA

University of North Carolina at Chapel Hill

WHAT ARE WE DOING WHEN WE ARE DOING EPISTEMOLOGY?

ABSTRACT: Jody Azzouni's (2025b) paper "Use and mention with respect to 'know', 'believe', 'evidence', 'justification', 'hypothesis', and so on: A hot mess" makes a number of important and valuable points. In this paper, I try to show that Azzouni's philosophical procedure can survive the worries that he considers, and unsatisfactorily addresses, in Section 16 of his paper.

1. INTRODUCTION

Azzouni's paper "Use and mention with respect to 'know', 'believe', 'evidence', 'justification', 'hypothesis', and so on: A hot mess" (2025b) is as sprawling as its title, and even more impressively accurate. Although I agree with virtually everything Azzouni says in this paper, I do have some reservations about some comments he makes in Section 16 of his paper, entitled "Some Methodological Moralizing". In that section, Azzouni attempts to defend the philosophical procedure of the previous 15 sections from an objection he envisions coming from Adam Leite. In the course of offering that defense, he characterizes his procedure in a way that is less than fully clear to me, and in this paper I'll try to explain why I'm having trouble understanding his characterization. My hope is

to achieve some of the understanding I don't yet have, even if only by provoking some elaboration from Azzouni.

Early in Section 16, Azzouni quotes part of Leite's explanation of how we can refute skepticism without appeal to philosophical theorizing. Here is the portion of Leite's explanation that Azzouni quotes:

We will carefully scrutinize skeptical arguments to see what requirements they impose, and we will consider 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 read a more general description, summarize the results of all this, and the like. (Leite 2024, 10)

Azzouni quotes these methodological remarks from Leite not in order to disagree with them, but rather in order to claim that following the procedure that Leite describes in this very quotation will, contrary to what Leite claims elsewhere, "tip us into philosophical theorizing". But why should we accept this claim? Why can't we effectively rebut skeptical arguments without philosophical theorizing? Consider some of Leite's allegedly "pre-philosophical" responses to skeptical arguments, e.g., you can know that you're not a brain-in-a-vat because you know that such technology doesn't exist yet, or you know that you're not dreaming because you know that dreams aren't like your current experience. Why must we engage in philosophical theorizing either to appreciate the force of such responses, or else to improve upon them?

Azzouni's response to this question goes as follows:

My counter to this way of trying to respond to 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.... [P]hilosophical theorizing invades our thinking

about, nearly enough, everything. (Azzouni 2025b, 60)

How, according to Azzouni, does philosophical theorizing invade our thinking about so many things? Before giving examples to illustrate this phenomenon, Azzouni first makes these two points:

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.

(Azzouni 2025b, 61)

I'm not sure how Azzouni thinks of these latter two points as related to his earlier point that philosophical theorizing invades our thinking about many things, but I'm inclined to think that each of the latter two points concerns one way in which philosophical theorizing invades our thinking. The first way is that philosophy leads us to reject some of our earlier ways of thinking about things. The second way is that philosophy leads us to supplement, or elaborate, how we think about things, and we do this "in order to understand what's going on". Azzouni illustrates the first of these two processes by discussing the liar paradox, which reveals that, contrary to what appears obvious, not every declarative sentence is either true or false but not both. He then moves on to discussing epistemology, and it is here, I assume, that we are meant to find some example of the second process (or at least some example of *the need for* the second process), the process of supplementing our ordinary ways of thinking about things "in order to understand what's going on".

That discussion of epistemology focuses on Leite's remark 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." (Leite 2024, 3)

Azzouni proceeds to argue that the last sentence of this quoted passage is false, and that "fallibilism isn't something that easily 'shown' by how we talk". (Azzouni 2025b, 63). But in the course of arguing for this, Azzouni characterizes fallibilism in a different way than Leite does. Azzouni describes fallibilism as "the view that our knowledge of something P is compatible with the possibility of not-P or with not-P

having a nonzero probability”. (Azzouni 2025b, 63). Before I go on to consider what Azzouni says against the obviousness of fallibilism, I want to make a few remarks about the difference between Leite’s characterization of fallibilism and Azzouni’s. I’ll begin by putting them side-by-side:

L-Fallibilism: We can know things even if our evidence is compatible with various possibilities of error.

A-Fallibilism: Our knowledge of something P is compatible with the possibility of not-P or with not-P having a nonzero probability.

Notice that L-Fallibilism describes knowledge as compatible not with the possibility or probability of anything, but rather with our evidence being compatible with the possibility of error. Thus, L-Fallibilism, but not A-Fallibilism, is consistent with the following view:

When you know that P, it’s *impossible* that not-P. Nonetheless, when you know that P on the basis of some evidence E, then, although your knowing that P is itself inconsistent with the possibility that not-P, your evidence E is consistent with the possibility that not-P (In other words, your knowledge rules out more possibilities than your evidence alone rules out.)

This latter view can itself be spelled out in various ways. For example, the consistency of your evidence E with the possibility that not-P can be understood in different ways: we might think that what is consistent with the possibility that not-P is the content of your evidence E (assuming that, contrary to the view that Azzouni expresses in Section 12, evidence has a propositional content), or we might think that what is consistent with the possibility that not-P is *your having* evidence E. Again, the evidence E might include only that evidence on the basis of which you knowledgeably believe that P, or it might include your total body of evidence at the time at which you knowledgeably believe that P. But however we spell out this latter view, it’s a view that the L-Fallibilist can consistently adopt but the A-fallibilist cannot.

Why does this difference between L-Fallibilism and A-Fallibilism matter? I’m not sure whether it matters to Azzouni’s historical point that just about every epistemologist before the end of World War II

was an infallibilist: for all I know, they were all A-Infallibilists but not L-Infallibilists, and the emergence of fallibilism in the middle of the twentieth century was not a matter of shifting views but rather of shifting attention. But I will not explore this historical question further.

One way in which the difference between L-Fallibilism and A-Fallibilism clearly matters for our purposes is that it affects Azzouni's discussion of lotteries, and Lottery Scepticism. Azzouni writes:

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. (Azzouni 2025b, 64)

What Azzouni says here supports A-Fallibilism. But what is its relation to L-Fallibilism? The L-Fallibilist denies that knowing that P requires you to have evidence that makes it impossible that not-P. From the fact L-Fallibilist denies this latter (putative) requirement on knowledge, it is tempting to infer that the L-Fallibilist *allows* for you to have knowledge even in cases where your evidence makes the probability of not-P nonzero. But this doesn't follow. For one thing, not-P's being possible doesn't entail that the probability of not-P is zero. (Consider a lottery with infinitely many tickets, no more than one of which can be the winner.). For another thing, not-P's being possible doesn't entail that not-P has any defined probability. (Consider a case in which P has a much more fine-grained, indexical content than any of our evidence, e.g., there is a person whose name begins with "S" who is standing precisely 10,000 miles NNW of my current position.). So the L-Fallibilist can happily grant that—at least prior to hearing the announcement of the winning ticket—you don't know whether some particular lottery ticket is a loser.

But why, under those circumstances (when you haven't yet heard the announcement of the winning ticket), *don't* you know that the ticket is a loser? The L-Fallibilist cannot answer this question by appeal to the fact that your evidence is compatible with the ticket's not being a loser. That would still be a fact even after you hear the announcement of the winning ticket, and so even after you come to know that this particular

ticket is a loser. So the compatibility of your evidence with your ticket's not being a loser cannot explain why you don't know, prior to hearing the announcement, that the ticket is a loser. What then, explains your failure to know that the ticket is a loser?

Before I attempt to answer this question, let me point out how peculiar the question is. Sometimes, when we discover that someone doesn't know something that we take to be obvious, at least to them, we seek an explanation of their failure to know that thing. For example, if I learn that my friend doesn't know what the capital of the United States is, I might wonder how they could possibly not know such a thing. Again, if I learn that my friend doesn't know what their own birthday is, I might wonder how they could possibly not know such a thing. In these cases, when I wonder how my friend could fail to know these things, I am surprised by their failure to know, and I seek to understand how such a surprising failure could have arisen. But when I ask what explains why it is that, prior to hearing the announcement of the winning ticket, you don't know that a particular lottery ticket is a loser, I'm not asking a question like that. There is nothing surprising about your failure to know that the ticket in question is a loser. Indeed, anyone who expected you to know, under your current informational circumstances, that the ticket is a loser couldn't reasonably have tried to sell you the ticket.

So when I ask what explains your failure to know, prior to hearing the announcement of the winning ticket, that a particular lottery ticket is a loser, I'm not wondering what explains a strange or surprising epistemic deficit on your part. I'm asking what it is that explains why *anyone* would fail to know, prior to hearing the announcement of the winning ticket or receiving any other such special information, that a particular ticket is a loser. The explanans cannot be something specific to a particular person or time or place, since the explanandum is not specific to a particular person or time or place. When we ask such questions, we assume that the explanans is some general feature of knowledge itself. But what justifies this assumption? More precisely, why think that there is anything about knowledge itself that explains why it is that you don't know, prior to hearing the announcement of the winning ticket, that a particular lottery ticket is a loser? Why not think, rather, that such facts are not all explainable in some uniform way? Could it be that what explains why I don't know that my own ticket is a loser is one thing,

whereas what explains why you don't know that your ticket is a loser is something wholly different? Or could it instead be that nothing explains why we don't know these things?

Such hypotheses will strike most of us as wildly implausible: it could not *just so happen* that each of us cannot know (prior to hearing the announcement of the winning ticket) that our ticket is a loser. There must, we think, be some uniform explanation of this general pattern in what we can and cannot know. And part of the job of epistemology, we think, is to find this uniform explanation. That, at any rate, is what I am tempted to say, and I suspect that Azzouni would agree. Recall Azzouni's statement 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," (Azzouni 2025b, 61). When Azzouni speaks of our effort to understand "what's going on", I suspect that it is precisely this kind of explanatory effort that he has in mind: the effort to discover some general feature of knowledge—or perhaps of our concept of knowledge—that can provide a uniform explanation of such patterns in what we know and don't know.

I've just said that I "suspect" that this is what Azzouni has in mind when he speaks of understanding "what's going on". But why do I merely *suspect* it? Why is it not completely obvious to me that this is what Azzouni has in mind? Indeed, Azzouni seems to call for precisely the kind of explanation I'm describing here when he writes:

[T]he practices and language of the purported pre-philosophical position aren't *transparent*.... 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.

(Azzouni 2025b, 66)

But to say that theorizing is required to answer these questions is not to say that such theorizing can succeed. It is not even to say that these questions have correct answers, whether knowable or not. As Azzouni notes, Hawthorne despairs of "the epistemological project of

extracting an analysis of knowledge from reflecting on use” (Hawthorne 2004, 68). As I understand it, Hawthorne’s despair doesn’t merely concern the Gettier project of providing non-circular and non-trivial truth-conditions for “S knows that P”, but concerns the even more general project of explaining various patterns in what we do and don’t know—for instance, concerning lottery tickets—by appeal to general features of knowledge. Does Azzouni share this despair? Does Azzouni think that these patterns have some uniform explanation (whether knowable or not) in the general features of knowledge? I suspect he does, but given his apparent sympathy with Hawthorne’s conclusion, I’m just not sure.

Should we share Hawthorne’s despair concerning this explanatory enterprise? I don’t think so. To explain why, let me examine the case that forms Hawthorne’s basis for that despair:

- (A) 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... “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 the child was only given the answer by luck.

(Hawthorne 2004, 68–69, quoted in Azzouni 2025b, 65)

I think it will help us to understand what’s going on in (A) if we consider what happens if we vary four features of the example: we vary the nature of the information conveyed so that it is less like public knowledge than like scurrilous gossip, we vary the channel of information conveyance so that it’s not a book but rather something not generally acknowledged as authoritative (at least with respect to the kind of information conveyed), we vary the grammatical form of the question asked so that it doesn’t merely presuppose that at least one of the protagonists has knowledge, and finally we vary the age and circumstances of the protagonists so that they are mature adults who might make some practical use of the information they gain, or assure others of it. Varying Hawthorne’s example in all four of these ways at once produces cases like the following:

- (B) I give six of my adult friends in the foreign service six crumpled-up pieces of paper, and ask them to pick one of the pieces of paper at random. Each paper contains a single messily handwritten sentence accusing Kim Jong-Un of some horrible deed, e.g., “Kim Jong-Un has raped children”, “Kim Jong-Un has killed dogs just for fun”, etc. and all but one of these pieces of paper contains misinformation about Kim’s crimes. I ask my friends to commit the information they receive to memory. I then ask the onlooker who has witnessed the whole sequence of events “Do any of these six foreign service workers *know* that Kim Jong-Un has committed the crime of which he is accused?” Clearly, the onlooker should answer no.

In case (B), I have varied all four of the features mentioned, and I have thereby created a case in which the onlooker should answer the question about knowledge in the negative. But are all four of these variations relevant? And if all four of them are relevant, then why is that? I’ll begin by addressing the first of these questions.

To see that all four of the variations are relevant, at least to some extent, let’s see what happens when we vary our original case (A) one feature at a time. We can start by varying the nature of the information:

- (C) I give six children six books, and ask them to pick one of the books at random. Each book contains a sentence accusing Kim Jong-Un of some horrible deed, e.g., “Kim Jong-Un has raped children”, “Kim Jong-Un has killed dogs just for fun”, etc. and all but one of these books contains misinformation about Kim’s crimes. I ask the children to commit the information they receive to memory. I then ask the onlooker who has witnessed the whole sequence of events... ‘Which one of the schoolchildren knows that Kim Jong-Un has committed the horrible deed of which he is accused?’... The onlooker may, as in case (A), be somewhat inclined to select the child whose book provides true information, but the onlooker’s inclination in this case is likely to be weaker than in case (A).

(C) varies from (A) in only one of the four respects that I’ve mentioned, but this variation slightly but discernibly weakens the inclination to attribute knowledge to the child whose book contains true information.

Now let's vary an additional feature of (A), viz., the channel of information conveyance.

- (D) I give six children six crumpled up pieces of paper, and ask them to pick one of the pieces of paper out of a hat at random. Each piece of paper contains a messily handwritten sentence accusing Kim Jong-Un of some horrible deed, e.g., "Kim Jong-Un has raped children", "Kim Jong-Un has killed dogs just for fun", etc. and all but one of these books contains misinformation about Kim's crimes. I ask the children to commit the information they receive to memory. I then ask the onlooker who has witnessed the whole sequence of events... 'Which one of the schoolchildren knows that Kim Jong-Un has committed the horrible deed of which he is accused?'... The onlooker may, as in case (A), be somewhat inclined to select the child whose book provides true information, but the onlooker's inclination in this case is likely to be weaker than in case (A), and even a bit weaker than in (C).

(D) varies from (A) in two of the four respects that I've mentioned, but this variation considerably weakens the inclination to attribute knowledge to the child who happened to pick the piece of paper that contains true information.

Now let's vary yet a third feature of (A), viz., the form of the question posed to the onlooker at the end.

- (E) I give six children six crumpled up pieces of paper, and ask them to pick one of the pieces of paper out of a hat at random. Each piece of paper contains a messily handwritten sentence accusing Kim Jong-Un of some horrible deed, e.g., "Kim Jong-Un has raped children", "Kim Jong-Un has killed dogs just for fun", etc. and all but one of these books contains misinformation about Kim's crimes. I ask the children to commit the information they receive to memory. I then ask the onlooker who has witnessed the whole sequence of events... 'Do any of the schoolchildren *know* that Kim Jong-Un has committed the horrible deed of which he is accused?'... In this case, the onlooker's inclination to answer that question in the affirmative will be very weak indeed, and their inclination to answer in the negative will be stronger.

Part of what seems to make the negative knowledge verdict in (E) at least somewhat plausible is the emphasis on “know”, together with asking a question in a way that doesn’t even appear to suggest that one of the six children does know. Changing the case in this way seems to me to undermine the appearance of knowledge even more than in (D).

I’ve identified a fourth factor as relevant to the positive knowledge verdict in (A), which is the age and maturity of the protagonists in the case, and their readiness to use this information for some practical purpose. Once we vary this factor in addition to the others, it seems to me that we destroy any residual inclination to attribute knowledge, and this, finally, is what we see in (B).

Recall (B): I give six of my adult friends in the foreign service six crumpled-up pieces of paper, and ask them to pick one of the pieces of paper at random. Each paper contains a single messily handwritten sentence accusing Kim Jong-Un of some horrible deed, e.g., “Kim Jong-Un has raped children”, “Kim Jong-Un has killed dogs just for fun”, etc. and all but one of these pieces of paper contains misinformation about Kim’s crimes. I ask my friends to commit the information they receive to memory. I then ask the onlooker who has witnessed the whole sequence of events “Do any of these six foreign service workers *know* that Kim Jong-Un has committed the crime of which he is accused?” Clearly, the onlooker should answer no.

I hope that these cases I’ve presented so far suffice to show that the positive knowledge verdict in Hawthorne’s case (A) depends upon the four features that I’ve enumerated, but of course it doesn’t depend upon other variables (e.g., the names of our protagonists, the time of day that all this happens, etc.). But why? Is there something about knowledge that could explain this pattern of variation with just these four variables? Or should we despair of finding any such general explanation?

The only basis that I know of for optimism concerning the prospects of such an explanation is that provision of such an explanation itself. And that’s what I want to turn to now. My explanation depends upon an important point that Azzouni makes elsewhere: knowledge requires not merely factivity, but also reliability.¹ But what is reliability? The expression “reliable” patterns with “desirable”, “laudable”, and “risible”. Each adjective indicates that the thing it modifies is worthy of a particular response—it is worthy of being desired, being lauded, being laughed at,

or, in the case of “reliable”, being relied upon. To say that knowledge requires reliability is to say that knowledge requires something that is worthy of being relied upon. But what is this required something?

Process reliabilists like Goldman (1986) will say that this required something is the process that produces potentially knowledgeable belief. Reasons reliabilists like Dretske (1971) will say that the required something is the reason(s) on the basis of which the potentially knowledgeable belief is held. But Azzouni is in no position to say either of these things, since he points out that knowledge doesn’t generally require belief.² Indeed, Azzouni has even challenged the claim that knowledge generally requires representation: simple mechanisms can know things without representing anything.³ So what is the reliable thing that knowledge allegedly requires? The only answer that I can see Azzouni being in a position to give is: it is the knowledge itself. In other words, knowledge itself must be a state that is reliable, i.e., worthy of being relied upon. To know something is thus to be in a state that you can rely on. To know whether the proposition *p* is true is to be in a state that you can rely on when it comes to the truth-value of *P*. To know about *X* is to be in a state that you can rely on when it comes to various crucial aspects of *X*.

When are you in a position to rely on something in a certain domain? This of course depends on various factors, but one of those factors has to do with the range of purposes for which you rely on that thing. Whether you are in a position to rely on Jones to keep a secret might very well depend upon what kind of secret we’re talking about. Even if Jones is generally very reliable, that doesn’t mean that she can be trusted with top secret national security information. Whether you are in a position to rely on your Subaru to get you home tonight depends on what kind of driving conditions you’re facing. Even if your Subaru is generally very reliable, that doesn’t mean that it can be trusted to get you up a steep hill that’s covered in ice. And similarly, whether you are in a position to rely on your knowledge of something might very well depend upon the range of purposes for which you need to rely on this knowledge. Of course, different kinds of information, and different kinds of situation, bring to mind very different purposes. The level of reliability required in order to know whether the lights in the hallway have been left on might be quite different from the level of reliability required in order to

know whether the fuel rods in the nuclear reactor have been replaced.

These considerations can help us to make some sense of the pattern illustrated by our examples A–E. One thing that diminishes the appearance of knowledge in cases D and E is that we don't generally treat messily handwritten messages from anonymous sources, written on crumpled up pieces of paper, as reliable indicators of the truth, especially not when the information concerns some issue on which speculation might fuel misinformation. To believe the content of such a message, and to do so for no other reason than that one sees the message, is to be in a state that one cannot generally rely upon, when it comes to the truth-value of the message at issue. Books, in contrast, are generally treated as reliable, at least when they concern matters of public knowledge such as national capitals.

The emphasis on “know” in the question posed at the end of E seems to do some work in generating our inclination to answer the question of knowledge in the negative. By emphasizing the term, we heighten the standards of its affirmative application. This happens across the board, with many different parts of speech. Suppose I ask you if there are mountains in Iowa, and you mention Hawkeye Point, which is 1670 feet above sea level, in Osceola County. In response, I might complain “that’s not a *mountain*. If that’s the highest point in Iowa then clearly Iowa has no mountains (or not any *real* mountains anyway).” Now suppose I ask you what you’re doing and, while unwrapping your microwave dinner, you reply that you’re cooking your dinner. In response, I might complain “that’s not *cooking*. That’s just microwaving.” Of course, such complaints might be rude, but they’re not unintelligible. We understand them precisely because we understand how emphasizing a term works to heighten the standards of affirmative application of that term.

Finally, knowledge seems to be absent in B because of all these factors, and one more. In order for the foreign service trainees to be able to rely upon their knowledge of Kim’s crimes, they will need to be able to rely upon this knowledge in a variety of professional contexts. But clearly, under the circumstances described in B, they cannot rely upon any such knowledge in any serious professional context. But, given the reliability of knowledge, it follows that if they cannot rely upon any such knowledge, that can only be because they don’t have any such knowledge.

Notice that all of the factors to which we've appealed in order to explain the absence of knowledge in B–E are missing from A. In A, each of six children forms some belief about the capital of Austria by reading it in a textbook. We think of the messages about the capital of Austria being easy to remember, and easy to distinguish from each other. We think of the textbooks as appearing authoritative. We don't think of the kids as having any very serious or high-stakes purposes in relying on their information about the capital of Austria. And so on. Thus, there is nothing to prevent our treating any factive state that results from their committing the book's content to memory as a state of knowledge. Thus, it seems to me that we can explain the pattern exhibited in A–E by appeal to the fact that knowledge must be reliable.

Have we responded adequately to Hawthorne's pessimism concerning the project of explaining various patterns in what we do and don't know by appeal to general features of knowledge? It seems to me that we've still left an important question unresolved. Recall that Azzouni began his discussion of Leite's effort at a "pre-philosophical" response to skepticism by pointing out that at least much of ordinary usage seems to presuppose A-Infallibilism. But the reason Azzouni then turned to Hawthorne's case (A) is because (A) appears to be incompatible with A-Infallibilism. So is there a real conflict in our ordinary knowledge ascribing practices? Or can A-Infallibilism be reconciled with Hawthorne's case (A)? I will now argue that they can be reconciled.

Recall that A-Infallibilism denies that our knowledge of something P is compatible with the possibility of not-P or with not-P having a nonzero probability. Let's begin by focusing on the possibility condition. If one of the children in Hawthorne's (A) knows that the capital of Austria is Vienna, then, according to the A-Infallibilist, this is not compatible with the possibility of the capital of Austria not being Vienna. So now let's consider: is it possible that the capital of Austria is not Vienna?

Of course there is no logical contradiction in the claim that the capital of Austria is not Vienna. And there is also no metaphysical or nomological necessity to the capital of Austria being Vienna: the capital of Austria could have been any number of places, consistently with thing's having the essences that they have, and consistently with the laws of nature. The possibility at issue in A-Infallibilism is epistemic possibility, not logical or metaphysical or nomological possibility. Indeed,

this is obvious from the fact that it is possible to know facts that are logically, metaphysically, and nomologically contingent, such as the fact that, as of this moment, more people now live in California than in Texas. So A-Infallibilism denies that our knowledge of something P is compatible with the epistemic possibility of not-P. So, for the relevant child in Hawthorne's example (A) to know that Vienna is the capital of Austria, it must not be epistemically possible—at least not given that child's body of knowledge—for the capital of Austria to not be Vienna. But why should we think that this is epistemically possible? What grounds could we have for taking it to be epistemically possible, from the child's point of view, that the capital of Austria is not Vienna?⁴

One thing we could do is to ask (or to imagine ourselves asking) the child who's just committed to memory the claim that Vienna is the capital of Austria, whether it's possible that it isn't the capital. But notice, if the child is now sufficiently unconfident in the truth of the claim that she's memorized to answer this question in the affirmative, we would then be willing to say that the child *doesn't* know that the capital of Austria is Vienna. She may have confidently believed it initially when she first read it in the textbook, but now that we are pressing her on whether it might not be true, she no longer maintains her confidence. And even if we don't think that knowledge generally requires confidence, in this case we would be very willing to deny knowledge to the child who's just lost their confidence.

Of course suppose we then follow up by giving the child a multiple-choice test on the capital of Austria, and she answers correctly, even without confidence in the correctness of her answer. Wouldn't her correct answer ground our willingness to ascribe to her knowledge that the capital of Austria is Vienna? What we are witnessing here is another example of the phenomenon mentioned above that our willingness to ascribe or to deny knowledge can shift depending upon the range of purposes that are salient to us in considering whether the putative knower is in a position to rely on their own state. The child who's memorized that the capital of Austria is Vienna is clearly in a position to rely on the resulting state for the purposes of answering a multiple-choice exam. But when we ask in the abstract "is it possible that the capital of Austria is not Vienna?"—and especially if we emphasize the word "possible"—we can bring to mind a much broader range of purposes,

in connection with at least some of which the child cannot rely on her state. The child, hearing out of the blue the question of whether it's *possible* that the capital of Austria is not Vienna, might reasonably wonder whether she can rely on the information that she read. Notice, however, that if the question is raised by someone who obviously (even obviously to the child) has no idea what they are talking about—e.g., the interrogator says “but I heard that Austria is a province of Canada and its capital is Mexico City—isn't it possible that the capital of Austria is not Vienna?”—the child might, in response, reasonably deny that this is possible. Whether the child would grant this possibility shifts back and forth with whether the child would confidently and unqualifiedly affirm that the capital of Austria is Vienna, which in turn (given that the capital of Austria really is Vienna) shifts back and forth with whether we would ascribe to the child knowledge that the capital of Austria is Vienna.

I think Hawthorne's case (A) thus poses no challenge to the view that knowledge that *p* is incompatible with the epistemic possibility that not-*p*. What about the other clause of A-Infallibilism, viz., the view that knowledge that *p* is incompatible with not-*p* having a nonzero probability? Does Hawthorne's case (A) pose a challenge to that view?

Let's begin by first understanding what sort of “probability” is at issue in this view. Clearly, the probability at issue cannot be perspective-independent. The relevant probability is such that we might assign various non-extremal values of it to various hypotheses about some past event, where the objective probability of that event is either 0 (if it didn't happen) or 1 (if it did). So the relevant probability must be relative to a particular agent's perspective.

Now, recall that Azzouni describes Newton, G.E. Moore, and other pre-War philosophers as “infallibilists”. Given his characterization of infallibilism as A-Infallibilism, this implies that these philosophers insisted on the incompatibility of knowledge that *P* with not-*P* having some nonzero perspective-dependent probability. But I don't know of any textual evidence to support this claim. While such philosophers would have been steadfast in denying the compatibility of knowing that *P* with the possibility that not-*P* is true, it doesn't follow that they had any views about the compatibility of knowledge with any particular level of perspective-dependent probability. Indeed, until such post-War

writings as Kyburg 1961 and Chisholm 1966, virtually no epistemologist discussed the connection between knowledge and probability. So I doubt that there is any significant pre-War philosophical tradition that would take any stand on the compatibility of knowing that P with not-P having a nonzero probability.

What about ordinary usage though? Is there something about ordinary usage that would indicate the incompatibility of knowing that P with not-P having a nonzero probability? It may seem that there is, and that Azzouni has already discussed it—viz., our denials of knowledge concerning whether a particular lottery ticket is a loser. But it becomes less clear that such denials reveal any commitment on the part of ordinary usage to the second clause of A-Infallibilism when we consider the fact that we are ordinarily happy to attribute knowledge that the ticket is a loser after the announcement of the winner has been heard: after all, the probability that the announcement was itself erroneous is presumably not zero either. But, if it's not our tacit commitment to the second clause of A-Infallibilism that explains the patterns of what we know and don't know in lottery cases, then what does explain those patterns?

To answer this question, consider some variations on the standard lottery case.

- (F) There is a lottery with infinitely many tickets. The first ticket has a 50% chance of winning. The second ticket has a 25% chance of winning. The third ticket has a 12.5% chance of winning. And so on, where each ticket n has $1/(2^n)$ odds of winning. You find that your own ticket is ticket 18,000,000,000. Do you know that your ticket is a loser? Here, I find our inclination to answer the preceding question in the affirmative is at least as strong as our inclination to answer it in the negative. Neither inclination is decisive though.

If our verdicts about lottery cases were explicable by appeal to the fact that there is a nonzero probability of your ticket's winning, then such verdicts should be just as strong and unequivocal in (F) as they are in the standard lottery case. But they're not. This indicates that nonzero probability alone cannot be what's doing the work in determining our verdict in the standard lottery case. The fact that, in the standard lottery case, the tickets are all more or less *equally* likely to win also seems to be explanatory relevant.

thebalticyearbook.lu.lv

Of course our lottery verdicts might not be the only feature of ordinary usage to suggest the incompatibility of our knowing that P with not-P having a nonzero probability. Perhaps it's also relevant that we are at least often inclined not to credit someone with *knowing* that P when they insist that there is, say, a 5% chance that not-P. For instance, if Jones insists that the chance of rain is 5%, we won't be inclined to describe Jones as *knowing* that it won't rain. Why is this? I think it's because *claims* concerning the obtaining of even small chances imply *claims* concerning epistemic possibility⁵, and knowing that P is incompatible with the epistemic possibility that not-P. So I think that the relevant features of ordinary usage can all be explained by appeal to the first clause of A-Infallibilism, which I've already defended against the objection from Hawthorne's case.

To sum up. I hope I've now explained why I regard all of the following as true.

- (i) L-Fallibilism: We can know things even if our evidence is compatible with various possibilities of error.
- (ii) Simplified A-Infallibilism: Our knowledge of something P is incompatible with the epistemic possibility of not-P.
- (iii) Knowledge is a state that is both factive and reliable.
- (iv) The reliability of knowledge helps to explain our verdict concerning Hawthorne's case (A) and variants of it.
- (v) Hawthorne's case (A) poses no threat to Simplified A-Infallibilism.

In defending these various claims, I've engaged in a lot of theorizing that is obviously philosophical. I've attempted to find the best explanation of our various judgments, and inclination to make judgments, concerning what we do and do not know, and I've sought this explanation by appealing to features of knowledge itself. To the extent that my explanatory effort has succeeded, I hope this allays any Hawthorne-inspired pessimism concerning the prospect of giving an explanatorily illuminating account of knowledge. So, to the extent that Azzouni suffers from such pessimism (and I still am not sure to what extent that is), I hope to have done something to mitigate it.

Notes

¹Azzouni 2025a, 34.

²Azzouni 2025a, 28.

³Azzouni 2020, 65–67.

⁴Albritton 2011 makes vivid how difficult it is to answer such questions.

⁵Perhaps, as Fantl & McGrath (2009) claim, *negligible* small probabilities don't imply epistemic possibility. But we don't neglect such probabilities when we assert that they obtain.

References

- Albritton, Rogers. 2011. 'On a Form of Skeptical Argument from Possibility'. *Philosophical Issues* 21: 1–24. <https://doi.org/10.1111/j.1533-6077.2011.00195.x>.
- Azzouni, Jody. 2020. *Attributing Knowledge: What It Means to Know Something*. Oxford: Oxford University Press, <https://doi.org/10.1093/oso/9780197508817.001.0001>.
- . 2025a. *Challenging Knowledge: How We (Sometimes) Don't Know What We Think We Know*. Oxford: Oxford University Press, <https://doi.org/10.1093/9780197789650.001.0001>.
- . 2025b. 'Use and Mention With Respect to 'Know', 'Believe', 'Evidence', 'Justification', 'Hypothesis', and so on: A Hot Mess'. *Baltic International Yearbook of Cognition, Logic, and Communication* 16: 1–70. <http://doi.org/10.4148/1944-3676.1130>.
- Chisholm, Roderick. 1966. *Theory of Knowledge*. Englewood Cliffs: Prentice-Hall.
- Dretske, Fred. 1971. 'Conclusive Reasons'. *Australasian Journal of Philosophy* 49, no. 1: 1–22. <https://doi.org/10.1080/00048407112341001>.
- Fantl, Jeremy & McGrath, Matthew. 2009. *Knowledge in an Uncertain World*. Oxford: Oxford University Press, <https://doi.org/10.1093/acprof:oso/9780199550623.001.0001>.
- Goldman, Alvin. 1986. *Epistemology and Cognition*. Cambridge: Harvard University Press.
- Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Kyburg, Henry. 1961. *Probability and the Logic of Rational Belief*. Middletown: Wesleyan University Press.
- Leite, Adam. 2024. *How to Take Skepticism Seriously*. Oxford: Oxford University Press, <https://doi.org/10.1093/oso/9780197691175.001.0001>.