

ANALYSIS OF KNOWLEDGE*

J. Adam Carter

University of Glasgow

COGITO Epistemology Research Centre

adam.carter@glasgow.ac.uk

1 INTRODUCTION

The expression ‘analysis of knowledge’ refers – in a wide sense – to a research programme that has been central to contemporary analytic epistemology from around the mid 20th century onwards.¹ The core methodology and aims of that research programme can be captured more or less straightforwardly by reference to the task of filling in the English schema

“S knows that p if and only if ___”

where ‘analysis of knowledge’ also then refers (in a more narrow sense) to whatever goes in the blank.

From the early 1960s to the present day, there has been dispute about just *what* should go in the blank; this is a dispute about what the correct analysis of knowledge is. From 2000 onwards – largely due to the influence of Williamson’s *Knowledge and its Limits* (2000), there is also serious dispute about whether it is fruitful to even seek an analysis of knowledge, i.e., to even *try* to fill in the above blank in the first place.² This latter dispute, roughly between contemporary ‘knowledge-firsters’ and their more traditional opponents, finds the knowledge-first side challenging not only the adequacy of substantive *analyses* of knowledge on the market,

*Thanks to the Arts and Humanities Research Council (AHRC) for supporting this research (Grant No. AH/W008424/1). Thanks also to [...]

¹See Shope (2017) and Ichikawa and Steup (2018) for helpful overviews.

²One complication here, however, is that ‘knowledge-firsters’ (of which more will be said later in this chapter) is itself not a precise term. Williamson is a paradigmatic knowledge firster in that he holds both that knowledge is unanalysable *and*, further, that knowledge should be theoretically explanatory in that we should use knowledge to explain other central notions in epistemology. Some epistemologists have opted for the second ‘explanatory’ leg of the knowledge-first research agenda, without signing on to the first ‘unanalysability’ leg; I mention this because it is probably a mistake to suppose that all knowledge-firsters sign on to the unanalysability component Williamson embraced in *KAIL*.

but the very enterprise proponents of such analyses are engaging in in attempting to provide them.

The second of the two disputes leaves the analysis of knowledge occupying an awkward (if not somewhat paradoxical) place as a research programme in epistemology. That is: on the one hand, it remains, even now, among the first things students are generally taught in introductory epistemology classes (along with the Gettier problem, as part of its core methodology – more on this shortly). And yet, on the other hand, with the emergence of the knowledge-first movement, now over 30% of all respondents in the most recent (2020) *PhilPapers survey*³ hold that knowledge is unanalysable, which implicates that roughly 30% of philosophers think epistemologists should (*qua* epistemologists) *stop* doing that very thing they are taught epistemology by learning how to do.

This chapter has a modest aim, which is not to resolve any philosophical disputes about the analysis of knowledge, but to simply clarify what lies behind them, and to give a brief overview of how the analysis of knowledge has gotten to the curious place it now occupies in epistemology.

Here is the plan. §2 sharpens the task of giving an analysis of knowledge by registering some shared methodological assumptions and then distinguishing between different kinds of analyses: semantic, conceptual, and metaphysical analyses. §3 then offers a brief overview of the tripartite (JTB) analysis of knowledge and ‘Gettierization’, while §4 details some notable ‘post-Gettier’ analyses of knowledge, focusing on the rise of the externalist tradition in epistemology (including modal and virtue-theoretic analyses of knowledge). §5 then brings the knowledge-first critique front and centre, and outlines, and then critically evaluates, two main arguments for rejecting any analysis of knowledge: the Negative Abduction Argument and the Distinct Concepts Argument. §6 concludes by taking stock of the current state of play.

2 VARIETIES OF ANALYSIS

Consider again the familiar schema:

“S knows that *p* if and only if ____”.

Filling in the blank here would require identifying conditions necessary and jointly sufficient for S’s knowing that *p*. This might seem straightforward enough. But different strands of debate in epistemology reveal that there is scope for background disagreement about both (i) what we are in fact doing, when giving such conditions, and (ii) what we should be trying to do. We’ll get to this contentious bit shortly.

³See Bourget and Chalmers (2020).

But let's start by tabling some assumptions that are widely agreed upon. A near-unanimous assumption concerns the range of the variables at play in the schema – viz., that the range includes not only actual thinkers (and propositions they've actually entertained) but also actual and possible thinkers and propositions. This means that identifying conditions that are extensionally adequate to all *actual* thinkers and propositions they know is not enough.⁴ This also means – methodologically – that assessing whether possible thinkers would count as knowing certain propositions under imagined circumstances is entirely appropriate to an assessment of what should go in the blank. That is: the schema (ranging over all possible thinkers/propositions) should be interpreted as holding necessarily.

We can accordingly sharpen the logical form of the familiar schema as follows:

$$\Box \forall s \forall p (K(s, p) \Leftrightarrow (_) (s, p))$$

which says that, necessarily, for all subject s and all propositions p , s knows that p if and only if s stands in a (fill in the blank) relation to p .⁵

Some further shared assumption concern the kind of thing that would go in the blank – which is that it should be (along with extensionally adequate *vis-à-vis* all possible thinkers/propositions) also *informative*, *non-circular*, and *complex*. Regarding informativeness: simply putting a 'K' in the blank:

$$\Box \forall x \forall p (K(s, p) \Leftrightarrow K(s, p))$$

is tautologous and so not informative, even if (obviously) true. So that's not what's sought. And neither is filling in the blank with something ϕ , where ϕ must be defined at least partly in terms of K . Finally, a *nearly* shared assumption⁶ – more substantive than methodological – is that what goes in the blank is (de facto) going to be complex, and so of the form:

$$\Box \forall x \forall p (K(s, p) \Leftrightarrow (_ + _ + _ + \dots) (s, p))$$

A well-known such complex proposal (which we'll discuss substantively in §3) locates *justified*, *true*, *belief* in the blanks opposite knowledge in the biconditional:

$$\Box \forall x \forall p (K(s, p) \Leftrightarrow JTB(s, p))$$

Here, though, is pretty much where shared assumptions end.⁷ Now, for the contentious bit. Suppose you've just filled in the blank, and you think that you've filled

⁴Compare: in seeking conditions necessary and sufficient for justice in *The Republic*, Plato was interested not just in characterising all actual instances of justice (and what so happens to be common to them), but also in all possible instances of justice, and what must be common to those as well.

⁵See Williamson (2008) for a fuller logical characterisation of not only the JTB analysis itself, but also of the conditions under which thought experiments are capable of refuting the analysis.

⁶Meno and Crispin Sartwell (1992) are some exceptions here; as are some analyses we will discuss later in §4 and §5.

⁷For summary: that (i) the schema should be understood as characterising a biconditional that holds *necessarily*; and that what goes in the blank (to complete the analysis) should be at least (ii) informative, (iii) non-circular; and (iv) complex – where 'JTB' is an example of a complex condition.

it in a way where the biconditional holds necessarily, and that you've done this in a way that is non-circular and informative, etc. In doing all of this, have you just *defined* knowledge? If you think the answer is 'yes', then you might also hold a further metaepistemological view, which is that analysis of knowledge is, in the main, a *semantic analysis* — viz., an analysis aimed at making precise the (linguistic) meaning of 'knows'. But there are other things you might think you are doing rather than defining knowledge. Perhaps you take yourself to be engaged (in completing the schema) in *conceptual analysis*, in which case you will take yourself to be analysing the *concept* of knowledge into constituent conceptual components. Or, separate from either of these two endeavours, you might take yourself to be analysing *knowledge itself* – apart from analysing just the meaning of 'knows' or the concept of knowledge; this is *metaphysical analysis*.⁸

At this point, the situation might seem confusing. And rightly so! The question of whether analysing knowledge via completing the schema is *also* any one or more of semantic, conceptual, or metaphysical analysis is a metaepistemological dispute (for discussion, see Carter and Sosa 2022).

Here is a way to hopefully bring a bit of control to the chaos. Suppose, for simplicity sake, that, in embracing a simple 'JTB' analysis of knowledge, you accept the biconditional: $\Box \forall x \forall p (K(x, p) \Leftrightarrow \text{JTB}(x, p))$. As Williamson (2008) has pointed out, you could in principle accept this *without* making any of the following: (i) an identity claim (to the effect that knowledge is identical to JTB); (ii) a synonymy claim (i.e., that the word "knowledge" is synonymous with "justified true belief"); or; (iii) a conceptual claim to the effect that the concept of knowledge is identical to the conjunctive JTB concept. Further, in a similar vein, you could accept the biconditional without embracing any kind metaphysical thesis, including the thesis that JTB is a metaphysical analysis of knowledge. That said, and to emphasise, if *any* of these further claims is true, then the biconditional is also true, and so "a refutation of [the biconditional] automatically refutes each of those further claims too, although not conversely" (2008, 183).

Even with this point in mind, though, we are still left with at least three open questions when it comes to semantic, conceptual, and metaphysical analyses of knowledge. One is descriptive: what, in fact, do most epistemologists take themselves to be doing in addition to be doing (if anything) *beyond* giving necessary and sufficient conditions? Note that an answer here needn't make reference to just one kind of analysis.⁹ A second question is normative: regardless of what we take ourselves to be doing, what *should* we be doing in epistemology? What methodology (e.g., seeking a semantic, conceptual, or metaphysical analysis, or some combi-

⁸For more detailed discussion, see Carter and Sosa (2022, sec. 3.2); see also Sosa (2015).

⁹For example, it's possible to embrace the biconditional – suppose one does this in the case of some version of a JTB account – and then to also embrace JTB-style analyses of (all three of) a knowledge definition in terms of JTB, a conceptual analysis in terms of JTB, and metaphysical analysis in terms of JTB.

nation of them) is appropriate to the epistemologist's project of analysing knowledge? Thirdly, in what way would establishing the truth of one kind of analysis provide evidence for another kind of analysis? For instance, if we learned that the concept KNOWS is identical to the (conjunctive) concept JUSTIFIED, TRUE, BELIEF, then to what extent (if any) would this discovery count in favour of thinking that 'knows' means 'justified true belief', or that knowledge is (metaphysically) grounded in JTB?

Because I take the above three questions to be largely open questions, I won't assume any answers here going forward. Instead, we'll get in to the territory that has most interested epistemologists – which is the *content* of what goes in the schema, and *why* 'JTB' is ill-suited.

3 THE JTB ANALYSIS AND GETTIERIZATION

Consider again the simple JTB analysis¹⁰, which holds just that:

$$\text{JTB-ANALYSIS: } \Box \forall x \forall p (K(x, p) \Leftrightarrow \text{JTB}(x, p)).$$

According to a narrative popularised¹¹ by Gettier (1963), the JTB analysis had been (prior to his 1963 paper) the received view, embraced by Plato, Chisholm (1957), and Ayer ([1957] 2006), among others. This narrative history has been called into question recently by Julien Dutant (2015), who doubts that the JTB-ANALYSIS ever enjoyed the kind of orthodoxy Gettier took it to have. But we needn't get bogged down here. The real action begins with Gettier's refutation of the analysis by counterexample.

Sufficient for refuting the JTB-ANALYSIS would be to establish the following:

$$\text{JTB-REFUTATION: } \diamond \exists s \exists p (\text{JTB}(s, p) \ \& \ \neg K(s, p))$$

And JTB-REFUTATION is what Gettier attempted to show – in a way widely taken to have been successful¹² – through several counterexamples cases depicting possible scenarios of unknown justified true beliefs. Here is one such case (Gettier 1963, 122):

GETTIER CASE: Suppose that Smith and Jones have applied for a certain job. And suppose that Smith has strong evidence for the following conjunctive proposition:

- (d) Jones is the man who will get the job, and Jones has ten coins in his pocket.

¹⁰From here on out, the metaepistemological questions about analysis varieties will be set aside.

¹¹Gettier was not the first to envision cases with this kind of structure – viz., where a belief that is justified is true seemingly by luck. The earliest version of such a case I'm aware of is due to the Indian philosopher Dharmottara, c. 770 CE.

¹²Widely, but not universally. A notable dissenter here is Stephen Hetherington (2016).

Smith's evidence for (d) might be that the president of the company assured him that Jones would in the end be selected, and that he, Smith, had counted the coins in Jones's pocket ten minutes ago. Proposition (d) entails:

- (e) The man who will get the job has ten coins in his pocket.

Let us suppose that Smith sees the entailment from (d) to (e), and accepts (e) on the grounds of (d), for which he has strong evidence. In this case, Smith is clearly justified in believing that (e) is true. But imagine, further, that unknown to Smith, he himself, not Jones, will get the job. And, also, unknown to Smith, he himself has ten coins in his pocket. Proposition (e) is then true, though proposition (d), from which Smith inferred (e), is false. In our example, then, all of the following are true: (i) (e) is true, (ii) Smith believes that (e) is true, and (iii) Smith is justified in believing that (e) is true. But it is equally clear that Smith does not know that (e) is true; for (e) is true in virtue of the number of coins in Smith's pocket, while Smith does not know how many coins are in Smith's pocket, and bases his belief in (e) on a count of the coins in Jones's pocket, whom he falsely believes to be the man who will get the job (1963, 122).

In GETTIER CASE, it looks as though Smith is justified in believing proposition (e), i.e., that the man who will get the job has ten coins in his pocket. And that proposition is true. But, as the thought goes, surely Smith does not *know* that the man who will get the job has ten coins in his pocket. As Gettier puts it: it is '*clear* that Smith does not know that (e) is true' (1963, 123); he is right 'by the sheerest coincidence' (*Ibid.*).

But if *that's* right, then JTB-REFUTATION is established, and by extension, we get: $\neg \Box \forall x \forall p (K(s, p) \Leftrightarrow \text{JTB}(s, p))$, and thus, the falsity of the JTB-ANALYSIS; and given that (as we saw in §2) the falsity of the JTB analysis (asymmetrically) implies the falsity of semantic, conceptual, and metaphysical JTB analyses of knowledge, it seems the import of cases with the structure of GETTIER CASE is theoretically significant.¹³

Let's slow down however, and take note of three important points.

First, although JTB-REFUTATION is logically incompatible with JTB-ANALYSIS, might we not still be a bit quick to jettison what might otherwise be a pretty good theory simply on the basis of an imagined 'counterexample' like GETTIER CASE, and even when, intuitively, that counterexample is compelling? This takes us from logic (which offers no guidance on this question) into the arena of philosophical

¹³Other simpler cases with an analogous kind of structure are often used under the description of 'Gettier cases'; see, e.g., Chisholm's (1977) 'sheep in a field' case.

methodology. According to Weatherson (2003), while counterexamples can be useful in philosophy, we should caution against too readily abandoning theories on the basis of counterexamples that themselves, we must remember, are often underlied by intuitions, some of which may be unreliable. By way of comparison, suppose you have a successful scientific theory with significant explanatory power. You run a test that then gets an anomalous result, one that the theory can't accommodate. Wouldn't you run the test again and inspect it carefully for testing errors before rejecting the theory? This is effectively what Weatherson thinks we should do generally in philosophy, given that the kinds of intuitions we often rely on are far from infallible.¹⁴

This brings us to a second point: what, if anything, accounts for the widespread intuition that the subject (Smith) in Gettier's case, whose justified belief is true by 'coincidence' (1963, 123), lacks knowledge? If we could really get clear about this, then perhaps we can respond to the first (Weatherson-style) point by maintaining that there are in fact principled, theoretical grounds denying knowledge in Gettier cases, grounds which go beyond a mere appeal to intuition that the subject lacks knowledge.

To some extent, substantive attempts to 'solve' the Gettier problem – i.e., attempts to show what *in addition* to JTB is needed for knowledge (see §4) – at least *implicitly* purport to show just this; after all, any theory that takes knowledge to require a JTB plus some 'X' implicates (by advancing that theory as correct) that it is because the satisfaction of condition 'X' is not met in Gettier cases that subject in those cases fail to know.

Nowadays, a popular general philosophical rationale (compatible with various substantive views of knowledge) for what 'goes wrong' in Gettier cases makes reference to 'luck'. As Duncan Pritchard (2015a) sees it, our inclination to withhold knowledge in Gettier cases is reveals our sensitivity to a fundamental platitude about knowledge, which is that it in some way 'excludes luck'. As the thought goes, the Gettier subject, despite getting it right on the occasion at hand, does so luckily – he could very easily have been wrong¹⁵ – and *that's* the general explanation for why knowledge is lacking.¹⁶ A more specific explanation, for Pritchard, would require a more precise articulation of the sense in which knowledge is incompatible with luck, with reference to which it would be claimed that the Gettier subject lacks knowledge in virtue of failing to satisfy that specific anti-luck condition.¹⁷

¹⁴Though the question of how to unpack precisely the kind of reliability that matters for philosophical intuitions to support philosophical knowledge is contentious; for discussion, see Alexander and Weinberg (2014).

¹⁵I'm using this slogan here, as a gloss on luck, for illustrative purposes, as it's the core idea Pritchard appeals to before attempting to sharpen the slogan; he ultimately opts for a specific kind of modal characterisation, which we'll discuss in (§4).

¹⁶See here also Dancy (1985, 134) for similar remarks.

¹⁷In more recent work, Pritchard (2016) has revised his terminology; he now takes 'risk' to play the role previously played by 'luck'. Accordingly, on the more recent thinking, rather than to seek an ac-

Thirdly, it is worth registering an alleged structural barrier to resolving the problem satisfactorily while holding on to a traditional fallibilist conception of epistemic justification, according to which the conditions required for being epistemically justified compatible with the falsity of the target belief.¹⁸ As Zagzebski (1994) sees it, any theory of knowledge as (some kind of) fallibly justified true belief is open to Gettier-style cases. Such theories will have to allow that whatever gap there is between justification and truth could, possibly, be shored up by coincidence. One could of course ‘remove’ any such gap (immunising oneself from Gettier cases) by embracing an infallibilist conception of justification, where a belief’s being justified entails that it’s true. But this, Zagzebski maintains, is too strong. So, there’s no way to escape Gettier problems, as the thought goes, in a theoretically satisfactory way.

Since Zagzebski’s 1994 claim of the inescapability of Gettier problems, there have been two main lines of resistance. There’s been a new-found popularity of infallibilism about justification associated largely (though not entirely) with knowledge-first epistemology; second, virtue epistemologists have developed a rationale for denying that fallibilism in the theory of justification implies the possibility of Gettier cases. We’ll look at the latter in (§4) and the former in (§5).

4 POST-GETTIER EXTERNALISM IN THE ANALYSIS OF KNOWLEDGE

Notice that, in GETTIER CASE, the conditions ostensibly met for the Gettier subject to be justified in believing the target proposition (e) had to do with the subject’s believing that proposition on the basis of adequate evidence. This is very much an ‘internalist-friendly’ way of thinking about epistemic justification, in that it fits snugly with the epistemic internalist thesis that what matters for whether a thinker is epistemically justified in believing a proposition is just her mental states, states which would (in Gettier’s cases) include the evidential grounds of her belief.¹⁹

A major breakthrough in the analysis of knowledge occurred later in the 1960s, when epistemologists – led by Goldman (1967) – first explored what epistemic justification might look like untethered to internalist thinking. Goldman noticed that, in GETTIER CASE the *fact* that actually makes Smith’s belief true (i.e., that the boss hired Jones, who had 10 coins in his pocket) really has nothing to do with anything Smith believes or to his mental states. Perhaps, as Goldman thought, what is im-

count of knowledge that suitably blocks (knowledge-incompatible) luck, we should seek an account that suitably blocks (knowledge-incompatible) risk. The difference between (knowledge-relevant) luck and risk is a minor technical difference (see also 2015) and for the present purposes, it will suffice to follow the previous convention of framing the relevant debate around luck.

¹⁸For an influential recent defence of this position in epistemology, see Brown (2018).

¹⁹Internalists in epistemology make important reference to both evidence as well as to reasons in explaining that status of a belief as justified. For an overview of some of the connections between these issues, see Sylvan (2016).

portantly missing here – and which might explain why knowledge is absent – is an appropriate causal connection between Smith's beliefs and the *fact* that his beliefs concerned.²⁰ This idea inspired his *causal account of knowledge*:

CAUSAL ACCOUNT OF KNOWLEDGE: S knows that p iff p is true, S believes p, and S's belief that p is causally connected to the fact that p in an appropriate way.

Notice that this causal account replaces with a causal condition the kind of internalist-friendly 'J' condition that featured in the JTB analysis Gettier challenged. And this causal condition is not an internalist condition; according to the causal account of knowledge, after all, all that matters for knowledge (that p) is that your belief is both true and appropriately causally connected to the fact that p; you needn't additionally know *that* this causal condition holds, or even have any beliefs about this.

Given that, on the thoroughly externalist causal theory of knowledge, Smith's belief that the man who will get the job has 10 coins in his pocket is not appropriately caused by this fact, it looks like the causal account does *not* seem to problematically 'rule in' GETTIER CASE as a case of knowledge, as the causal condition isn't satisfied. So far, so good. The problem is, there are other cases where it is satisfied and yet it seems as though it is just a coincidence that the agent's belief is true. Consider the following famous 'Fake Barn' case:²¹

FAKE BARN: Henry is driving down the countryside and correctly identifies a barn. He does so despite the presence (unbeknownst to him) of nearby barn facades that would have misled him.

Here's a compelling thought: if we are inclined to think that the sheer coincidence by which Smith's belief is true blocks it from being knowledge in GETTIER CASE, then shouldn't we also think similarly for Henry in FAKE BARN, given that it seems just a matter of luck that he looks at a real barn rather than at one of the fakes that would have tricked him?

Cases like FAKE BARN have led externalist epistemologists to opt for *modal* conditions on knowledge which might seem tailor-made to rule out knowledge in FAKE BARN-style where *prima facie* knowledge-undermining luck persists despite appropriate causation. One such well-known modal condition requires that known beliefs be modally *sensitive* to the truth in the following way:

SENSITIVITY: If S knows that p, then if p had not been true, then S would not have believed that p.²²

²⁰See also, for an overview, Carter and Littlejohn (2021, Ch. 5).

²¹This case (simplified here for presentation) was popularised by A. I. Goldman (1977), who credits the case to Ginet (1975).

²²This is a simplified formulation; it would also need to be relativised to a basis in order to avoid grandmother-style cases due to Nozick (1981).

Another closely related idea is *safety*:

SAFETY: If S knows that p, then S's true belief that p could not have easily been false.²³

Notice that if *either* SENSITIVITY or SAFETY succeeds in capturing a necessary condition on knowledge, then we get the result that Henry lacks knowledge in FAKE BARN. Consider first that if Henry's belief "There is a barn" were false, then he wouldn't have believed it; the nearest world where his belief is false is a world where Henry looks (perhaps just a bit right or a bit left) at one of the facades, and in such a world he ends up with a false belief. Thus, Henry's belief is not *sensitive*. Likewise, his belief is not *safe*; his belief could very easily have been false given that there are near-by worlds where he forms a belief about whether there's a barn in the relevantly same manner but ends up looking instead at one of the many fakes and believing falsely.

It is relatively uncontentious that SENSITIVITY and SAFETY can both successfully rule out knowledge in FAKE BARN, as well as in the original GETTIER CASE. However, it is far more contentious is whether either is capable of *not* ruling out *too much*. In fact, this concern is actually applicable, in the first instance, in the case of FAKE BARN itself. While it is near-unanimously agreed that Smith lacks knowledge in GETTIER CASE, opinion is more divided in FAKE BARN. Suffice to say that if Henry *does* know despite the kind of coincidence or luck that features in the case – a view increasingly supported by experimental studies of folk judgments of fake barn cases²⁴ – then it follows that both SENSITIVITY and SAFETY are too strong.

But setting FAKE BARN aside, independent reasons have been offered for thinking SENSITIVITY and SAFETY, respectively, rule too much genuine knowledge out. This has been a particularly pressing worry for proponents of SENSITIVITY. Consider that if you were a brain in a vat, you would continue to think you were not. So, by SENSITIVITY, you don't know you're not a brain in a vat, and (at least, provided we know those propositions that we competently deduce from what we know) neither would you know anything that you know entails you are *not* a brain in a vat – which, problematically, includes many everyday propositions. An important advantage of SAFETY is that it easily avoids this kind of result; after all, provided BIV worlds are really 'far off' worlds, then the obtaining of the BIV scenario is not something that would *easily* happen. As the thought goes, even though you *could* be wrong (were such a scenario to obtain), you couldn't *easily* be wrong, given the scenario doesn't

²³This is also a simplified formulation. For various proponents of safety as a necessary condition on knowledge, see, e.g., Sosa (1999), Pritchard (2005), Luper-Foy (1984), Sainsbury (1997) among many others. Williamson is also often credited as a proponent of safety; he writes, 'If one knows, one could not easily have been wrong in a similar case' (Williamson 2000, 147). However, Williamson (who rejects the project of analysis – see §5) regards safety as a *circular* necessary condition on knowledge. In this respect, the sense in which he accepts SAFETY does not accompany any further idea that it would be useful in completing some version of a JTB-style analysis.

²⁴See, e.g., Colaço et al. (2014) and Turri, Buckwalter, and Blouw (2015).

obtain in any near-by worlds. This point then indicates a comparative advantage of SAFETY over SENSITIVITY as a prospective modal condition on knowledge.

But even if the proponent of SAFETY can – in short – disregard ‘far-off’ sceptical scenarios, things become more complicated when we consider sceptical scenarios a bit closer to home. For example, according to Sosa (2007), the modal nearness of the dreaming (even if not the BIV) scenario’s obtaining would seem to render our ordinarily perceptual beliefs unsafe even when we’re awake. And so, if knowledge requires safety (in the manner captured by SAFETY), then problematically, it looks as though embracing SAFETY as a condition on knowledge generates the *ex ante* undesirable result that we lack much of the ordinary perceptual knowledge we take ourselves to have. Granted, Sosa’s contention that dreaming renders our ordinary perceptual beliefs unsafe is itself controversial and opens up challenging questions about what constitutes genuine *near-by* error possibilities generally, and under what conditions these are triggered by dreaming scenarios specifically.²⁵

Let’s zoom out: to this point we’ve seen how ‘going externalist’ in our thinking about what makes a given belief epistemically justified in the way that matters for knowledge has several advantages in the context of giving an analysis of knowledge. We saw that a simple causal accounts of knowledge (such as Goldman’s) can handle the original GETTIER CASE; however, by the same token, we saw that additional externalist machinery such as either SENSITIVITY or SAFETY looks needed if we want the result that knowledge is lacking in FAKE BARN and SAFETY. A potential cost of this additional machinery, though, is that it might rule out *too* much, effectively leaving any analysis of knowledge that signs on to such machinery too restrictive.

Over the past 20 years or so, *virtue theoretic* accounts of knowledge have emerged as a popular externalist approach to analysing knowledge – and in doing so dealing with Gettier cases – but without *requiring* any explicit recourse to modal codicils such as SENSITIVITY or SAFETY. The most notable variation of this idea is Sosa’s *aptness account of knowledge*:

APTNES ACCOUNT OF KNOWLEDGE: Knowledge is apt belief; a belief is apt iff it is accurate because adroit.²⁶

Unpacking this terminology: a belief is accurate iff true, and it is adroit iff in form-

²⁵For some recent discussions on this point, see Carter (Forthcoming, Ch. 2) and Greco (2020). Interestingly, Sosa himself offers an alternative reading of dreams in which the nearness of the dreaming scenario cannot in principle render our waking beliefs unsafe; this is his interpretation of dreams as imaginings (2007, Ch. 1).

²⁶The first versions of this thesis are found in the early to mid 1990s in Sosa (1991) and Zagzebski (1996) and received substantial further development in Sosa’s 2005 John Locke Lectures, later published as *A Virtue Epistemology* (2007) and more recently in Sosa (2015) and Sosa (2021). Other variations are found in Greco (2010), Turri (2011), and Carter (Forthcoming). Note that Sosa’s own thinking countenances multiple grades of knowledge, and so the characterisation attributed to him here is in first approximation that is superceded by his fuller view noted in §6.

ing the belief you manifest epistemic competence or skill.²⁷ The core idea of the aptness account of knowledge is that knowledge requires that the accuracy of the belief because of, or manifest, such competence or skill. This is an idea that has straightforward import for Gettier-style cases. For example: Smith's belief that the man who will get the job has 10 coins in his pocket is accurate (it's true), and it is adroit (it is formed in a competent manner) but it is not accurate *because* adroit, but rather, because of luck.

While the APTNESS ACCOUNT OF KNOWLEDGE can rule out GETTIER CASE as a case of knowledge without needing to lean on a sensitivity or safety condition, the situation is different with FAKE BARN; after all, isn't Henry's belief apt? If so, isn't this a bad result?

The difficulty in dealing with cases like FAKE BARN has divided virtue epistemologists sympathetic to the APTNESS ACCOUNT OF KNOWLEDGE into broadly three camps: those who (i) grant that knowledge is not present in FAKE BARN but then insist that Henry's belief is on closer inspection *not* apt; (ii) concede that his belief is apt but then deny that this is on closer inspection a problematic result; (iii) grant that knowledge is not present in FAKE BARN, concede that Henry's belief is apt in this case, and then opt for an extra condition (i.e., safety) to explain why the belief isn't knowledge.²⁸ The first strategy is Greco's.²⁹ The second strategy is Sosa's. And the third strategy is Pritchard's.³⁰ We lack the space here to adjudicate between the three strategies, which represent examples of some of the more popular contemporary attempts to give the kind of extensionally adequate analysis of knowledge that Gettier initially challenged us to find.

²⁷One standard approach to the idea of aptness is developed by Sosa (2015), and which distinguishes an inner-most 'skill' (which one can retain in bad shape or when improperly situated to exercise that skill) from a *complete competence* (i.e., a skill exercised in proper shape and when properly situated); the aptness of a belief requires (by virtue of requiring accuracy because of adroitness) that the belief's correctness derive from the exercise of *complete competence*. Cf., Greco (2010, Ch. 5) for a different approach to the the structure of aptness.

²⁸It's worth noting that cases like FAKE BARN represent just one kind of challenge for proponents of the APTNESS ACCOUNT OF KNOWLEDGE, one that appears – at least for those inclined to deny knowledge in FAKE BARN – to show that the aptness account is too weak. An independent line of argument, advanced by Lackey (2007) and Pritchard (2012) appeals to testimony cases in order to maintain that the aptness account is too *strong*; the central idea here is that we gain knowledge in paradigmatic testimony cases despite relying on others in a way that is in tension with the thought that our own abilities account for our cognitive success. Whereas Greco (2010, Ch. 5) and Sosa attempt to meet this challenge while holding on to the idea that apt belief is necessary (and sufficient) for knowledge, Pritchard's move in response to testimony cases is to 'weaken' the ability condition on knowledge while separately maintaining a safety condition in order to handle Gettier cases FAKE BARN. For critical discussion, see Carter (Forthcoming, Ch. 2).

²⁹Greco has pursued this strategy in several places, most notably in Greco (2009), Greco (2010, Ch. 5), and Greco (2012). Other variations on this strategy are defended in Carter (2016) and Littlejohn (2014).

³⁰See Pritchard (2012); see also Kelp (2013).

5 KNOWLEDGE-FIRST AND ITS CRITICS

As indicated in §1, an influential movement in mainstream epistemology – *knowledge-first* – rejects not only all particular substantive analyses of knowledge (such as those just discussed in §4) but importantly also the legitimacy of the project of analysing knowledge into constituent components. As knowledge-firsters see it, (i) knowledge is unanalysable; and (ii) is useful in epistemology – and more broadly in philosophy³¹ – to explain other phenomena rather than to be analysed in terms of them.

Because this chapter concerns the analysis of knowledge, we'll focus on the *negative* thesis – that knowledge is unanalysable; this section gives an overview and brief critical evaluation of the two main arguments Williamson gives for the unanalysability thesis: the *Negative Abduction Argument* and the *Distinct Concepts Argument*. We'll then close in §6 with a short assessment of the current state of play.

5.1 The Negative Abduction argument

Early on in *Knowledge and its Limits*, a datum that Williamson draws attention to, and reasons from, is the alleged “failure of the extensive post-Gettier research programme over several decades” (2005, 434). The narrative he presents is a familiar one: we find that

a succession of increasingly complex analyses have been overturned by increasingly complex counterexamples (2000, 31).

The trend towards complexity in these accounts (and their counterexamples) is such that Williamson takes it to reveal

[...] obvious signs of a *degenerating research program*. Most of them [the analyses of knowledge], if correct, seemed to make knowledge too grue-like to be worth analyzing (2000, 2, my italics).

Williamson is, in the above passages, effectively drawing an inference (from a particular type of ostensibly unfruitful track record achieved by a methodological strategy) to the purported best explanation of that track record – viz., that the methodological strategy manifest in the bad track record is flawed. There is nothing wrong with this kind of negative abduction argument generally. It is an instance of the type: “Method X has failed many/most/all times in the past” (is evidence for) “X is a failure as a method.”

Moreover, this kind of negative abduction looks compelling when we reflect on the kind of highly complex analyses we find as representative in epistemology in the several decades after the 1960s.³² The *method* characteristically used in these

³¹For discussion of knowledge first approaches within, but also beyond, epistemology, see (eds.) Carter et al. (2017).

³²See Shope (2017) for an overview.

decades (and which Williamson rightly takes objection to) is ‘conquer through increasing complexity’, where increasing complexity was taken to be what was needed to address previous less complex but extensionally inadequate analyses.

For an illustrative example of such complexity, consider – plucked from the heyday of post-Gettier analyses – Swain’s 1981 proposal:

SWAIN’S ANALYSIS OF KNOWLEDGE: *S* knows *p* at *t* if and only if (i) *p*; and (ii) there is some set of reasons, *R*, such that (a) *R* is (or the members of *R* are) that upon which *S*’s belief that *p* is based at *t* (b) *S*’s believing that *p* on the basis of (the members of) *R* is justified at *t*, (c) *S* has *R* as a result of at least one nondefective causal ancestry, and (d) if, at *t*, *S* has [or has had] any other reasons, *R*^{*}, that are relevant to *S*’s justifiably believing that *p*, then *S* would be justified in believing that *p* on the basis of (the members of) the set of reasons formed by the union of *R* with *R*^{*}.

It seems fair to grant that the kind of method which leads to analyses in the neighbourhood of complexity of the above supports a negative abduction. But some caution is needed, though. The ‘conquer through complexity’ method really describes only a partial history of post-Gettier attempts to analyse knowledge. Shortly after Williamson’s negative abduction against knowledge analyses in *KAIL*, in 2000, early versions of the virtue epistemologist’s simple core idea from §4 – that knowledge is type-identical with apt belief – were beginning to take shape, as were modest variations on them.³³

We’ve already seen in §4 that the ‘knowledge is apt belief’ is not without its own potential vulnerabilities. However, we should distinguish substantive objections from methodological objections. Whatever else we want to say about the APTNESS ACCOUNT OF KNOWLEDGE, it is not a continuation of a pattern of more complex, artificial views of the sort targeted by the negative abduction.³⁴ In its simplicity it represents a reversal of that trend. Negative abduction from failures of a methodology of increasingly complex analyses to the failure of a *simple* analysis is not justified. Accordingly, then, even if the Negative Abduction Argument is bad news for gerrymandered analyses, this bad news doesn’t obviously extend to the idea that knowledge is apt belief, or to simple variations on this idea.

5.2 The Distinct Concepts Argument

Williamson has another argument, though. This one doesn’t rely on abduction from a bad track record, but instead directly challenges *any* analysis of the concept of knowledge that equates the concept KNOWLEDGE with a certain kind of con-

³³For discussion of some of these proposals, see, e.g., Carter (*Forthcoming*, Ch. 2) and chapters in (ed.) Vargas (2016).

³⁴For some more detailed discussion on this point, see Miracchi and Carter (2022).

junctive concept.³⁵ Here's how the argument goes:

Distinct Concepts Argument

1. Every standard analysis of the concept knows equates it with some conjunctive concept which has the concept TRUE as a non-redundant constituent.
2. The concept TRUE is not a mental concept.
3. Any concept with a non-redundant non-mental constituent is not a mental concept.
4. So the conjunctive concepts with which the concept knows is equated by analyses of the standard kind are not mental concepts.
5. The concept KNOWS is a mental concept.
6. A mental concept can't be the very same concept as a non-mental concept.
7. So the mental concept KNOWS can't be the same concept as any of the conjunctive concepts with which it is equated by standard analyses.
8. So every standard analysis of the concept KNOWS is incorrect.³⁶

As we saw in §2, commitment to a biconditional such as $\Box \forall x \forall p (K(s, p) \Leftrightarrow (_ + _ + _ + \dots)(s, p))$ *needn't* commit one to any associated *conceptual* analysis of knowledge. Even so, such conceptual analyses are common enough within the tradition of analysing knowledge³⁷ that Williamson's target here is very significant in scope.

Is the argument sound? I think it is not; however – on the other side of the coin – there is a sound argument in the neighbourhood of the Distinct Concepts Argument that plausibly does imperil *many* though not all conceptual analyses of knowledge.

The reason the argument is unsound is that P1, which has all analyses of KNOWS in its crosshairs, is false: *not all*, analyses of KNOWS involve *irredundant* non-mental component concepts. Here are some examples that don't: Brian Ball's (2013) account of knowledge as normal belief, Goldman's (1967) account of knowledge as belief caused by the fact believed, and – of relevance given what we've noted in §5.3 – the *aptness account of knowledge*, which identifies knowledge with apt belief.³⁸ None of these accounts – when interpreted as offering conceptual analyses – equates the concept KNOWS with some conjunctive concept which has the concept TRUE as a non-redundant constituent. And so Williamson's argument is going to be inapplicable to such accounts, *even if* it is applicable to those (many – including many in the post-Gettier decades) analyses that *do* involve irredundant non-mental component concepts.

³⁵For criticism of Williamson's DCA, see, e.g., Goldman (2009) and Cassam (2009). For discussion, see Otero (2020).

³⁶I borrow this particular way of formulating Williamson's argument from §1.3 in *KAIL* from Cassam (2009). For another similar formulation, see Sosa (2015, n. 13).

³⁷For discussion on this point, see, e.g., Ichikawa and Steup (2018).

³⁸See Sosa (2015, n. 13). for discussion.

6 THE CURRENT STATE OF PLAY

Let's briefly summarise some of the key takeaways so far:

- *Scope of an analysis of knowledge (§2)*: When epistemologists – under the description of analysing knowledge – attempt to fill in the blank of the schema “S knows that p if and only if ___”, this schema should be interpreted as applicable to all actual *and possible* agents and propositions; the logical form is the biconditional: $\Box \forall s \forall p (K(s, p) \Leftrightarrow (_) (s, p))$.
- *Two metaepistemological caveats (§2)*: (i) desiderata on what goes in the blank include extensional adequacy (i.e., mere true belief is not enough) but also other theoretical desiderata, such as noncircularity and informativeness; (ii) we can distinguish between embracing (a) a biconditional of the form $\Box \forall s \forall p (K(s, p) \Leftrightarrow (_) (s, p))$; and (b) any of: a *conceptual* analyses, *semantic* analyses, and *metaphysical* analyses; refuting the former refutes (a corresponding version of) any of the latter, but not vice versa.
- *Gettierization (§3)*: A *substantive* desideratum on an analysis of knowledge is that it not ‘rule-in’ Gettier cases as cases of knowledge; what ‘goes wrong’ in Gettier cases is debatable, but one popular explanation makes reference to an alleged incompatibility between knowledge and luck.
- *Post-Gettier externalism (§4)*: Externalism in epistemology, by way of (e.g.) causal, modal, and virtue-theoretic (i.e., aptness) analyses of knowledge, offers resources for ‘blocking’ (certain kinds of) knowledge-undermining luck, e.g., of the sort that features in Gettier cases.
- *Knowledge-first epistemology (§5)*: advances (i) the *negative thesis* that knowledge is unanalysable; and (ii) the *positive thesis* that knowledge, qua theoretical primitive, is particularly suitable for explaining other phenomena of interest in epistemology.
- *Two arguments for the negative thesis* include the Negative Abduction Argument (§5.1) and the Distinct Concepts Argument (§5.2). The former is inapplicable to non-complex analyses of knowledge; the latter is inapplicable to analyses which do not equate the concept KNOWS with some conjunctive concept which has the concept TRUE as a non-redundant constituent.

At this point it is worth zooming out. Taking in a broad view of the analysis of knowledge, its spotty track record, as well as the more recent rise of the knowledge-first programme, how exactly should we view the analysis of knowledge as a project in epistemology *nowadays* – some 60 years since the Gettier problem? Has it – as knowledge-firsters maintain – simply run its course, or is it still deserving of the some or most of significance it has enjoyed? These are difficult questions, but I’ll conclude by venturing some attempted (albeit cautious) answers.

First, it seems that two concessions to the knowledge-firsters should be granted: (i) *regardless* of what we make of the case for thinking knowledge is unanalysable, the

knowledge-firster's *positive thesis* has been developed in various ways³⁹ (in debates about the norms of assertion, action, practical reasoning, evidence, belief, etc) to effect in the decades since *Knowledge and its Limits*, and this is a point traditionalists would be fair to concede; (ii) A second concession the traditionalist should make to the knowledge firster is that – taken together – the Negative Abduction Argument and the Distinct Concepts Argument imperil a large and significant swath of – even if not *all* – projects under the banner of the analysis of knowledge.

But – and here is a second point – the knowledge-firster would be fair to make two concessions to the traditionalist: (i) first, although the Negative Abduction Argument and the Distinct Concepts Argument imperil some, perhaps even many, analyses of knowledge, there remains (as we saw in §§5.3-4) a certain kind of 'sweet spot' that neither argument touches. That is: the Negative Abduction Argument doesn't imperil *simple* analyses of knowledge, and the Distinct Concepts Argument doesn't imperil analyses which *don't* equate the concept KNOWS with some conjunctive concept which has the concept TRUE as a non-redundant constituent. What this means is that (ii) *if* the traditionalist can put their finger on an analysis within this sweet spot that *also* manages to meet the kinds of substantive desiderata discussed in §3 and §4, then the knowledge-firster will have given us no good reason to discount such an analysis.

Question: *is* there such an analysis the traditionalist might point to? Here I will be a bit less cautious. Let's remember that the view that one knows that p iff one has an apt belief that p is neither (i) a complex or gerrymandered analysis; nor (ii) does it involve any irredundant non-mental component concepts. So the aptness analysis of knowledge lies squarely in the 'sweet spot' untouched by the Negative Abduction and Distinct Concepts Arguments. That's an auspicious start. Just how far might it get? As we saw in §4 – the aptness account seems, *prima facie*, overly permissive in fake barn cases; in this respect, the account is, as critics have pointed out, at odds with some of our pre-theoretical intuitions about luck's (in)compatibility with knowledge.⁴⁰ There is, however, a case to be made that the aptness account's verdict in such cases is feature not a bug of the view. After all, if there are multiple *grades* of knowledge – it might be that aptness tracks a lower grade of knowledge compatible with fake-barn-style (even if not ordinary Gettier-style) luck, even while at the same time a higher-grade knowledge, corresponding with a more substantial kind of intellectual achievement, is *not* compatible with such luck. If any analysis of knowledge is to ultimately succeed – (some version of) a multi-tiered aptness account of knowledge *might just be* where we'll find it.⁴¹

³⁹See Benton (2014) for some representative discussion.

⁴⁰Versions of this objection have been leveled in various places by Kallestrup and Pritchard (2012), Pritchard (2009), Pritchard (2015b), and Pritchard (2012).

⁴¹The most recent version of this view advanced by Sosa is his *Epistemic Explanations* (2021). See also Carter (Forthcoming; 2021) for my own preferred variation on this approach.

REFERENCES

- Alexander, Joshua, and Jonathan M. Weinberg. 2014. "The 'Unreliability' of Epistemic Intuitions." In *Current Controversies in Experimental Philosophy*, 128–48. Routledge.
- Ayer, A. J. (1957) 2006. "The Problem of Knowledge." In *Ayer Writings in Philosophy: A Palgrave Macmillan Archive Collection*, edited by Ted Honderich. Palgrave Macmillan.
- Ball, Brian. 2013. "Knowledge Is Normal Belief." *Analysis* 73 (1): 69–76.
- Benton, Matthew A. 2014. "Knowledge Norms." *Internet Encyclopedia of Philosophy*.
- Bourget, David, and David J. Chalmers. 2020. "Philosophers on Philosophy: The 2020 PhilPapers Survey."
- Brown, Jessica Anne. 2018. *Fallibilism: Evidence and Knowledge*. Oxford University Press.
- Carter, J Adam. 2021. "De Minimis Normativism: A New Theory of Full Aptness." *The Philosophical Quarterly* 71 (1): 16–36. <https://doi.org/10.1093/pq/pqaa017>.
- Carter, J. Adam. Forthcoming. *Stratified Virtue Epistemology: A Defence*. Epistemology Elements. Cambridge, UK: Cambridge University Press.
- . 2016. "Robust Virtue Epistemology as Anti-Luck Epistemology: A New Solution." *Pacific Philosophical Quarterly* 97 (1): 140–55.
- Carter, J. Adam, Emma C. Gordon, and Benjamin Jarvis, eds. 2017. *Knowledge First: Approaches in Epistemology and Mind*. Oxford University Press.
- Carter, J. Adam, and Clayton Littlejohn. 2021. *This Is Epistemology: An Introduction*. John Wiley & Sons.
- Carter, J. Adam, and Ernest Sosa. 2022. "Metaepistemology." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Spring 2022. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/spr2022/entries/metaepistemology/>.
- Cassam, Quassim. 2009. "What Is Knowledge?" *Royal Institute of Philosophy Supplements* 64: 101–20.
- Chisholm, Roderick M. 1957. *Perceiving: A Philosophical Study*. Cornell University Press.
- . 1977. *Theory of Knowledge*. Prentice Hall, Englewood Cliffs, NJ.
- Colaço, David, Wesley Buckwalter, Stephen Stich, and Edouard Machery. 2014. "Epistemic Intuitions in Fake-Barn Thought Experiments." *Episteme* 11 (2): 199–212.
- Dancy, Jonathan. 1985. *An Introduction to Contemporary Epistemology*. Blackwell Oxford.
- Dutant, Julien. 2015. "The Legend of the Justified True Belief Analysis." *Philosophical Perspectives* 29 (1): 95–145.
- Gettier, Edmund L. 1963. "Is Justified True Belief Knowledge?" *Analysis* 23 (6). <https://doi.org/10.1093/analys/23.6.121>.
- Ginet, Carl. 1975. *Knowledge, Perception, and Memory*. D. Reidel Publishing Co.
- Goldman, Alvin. 2009. "Williamson on Knowledge and Evidence." In *Williamson on Knowledge*, edited by Patrick Greenough, Duncan Pritchard, and Timothy Williamson, 73–91. Oxford University Press.
- Goldman, Alvin I. 1967. "A Causal Theory of Knowing." *The Journal of Philosophy* 64 (12): 357–72.
- . 1977. "Discrimination and Perceptual Knowledge." *The Journal of Philosophy* 73 (20): 771–91.
- Greco, John. 2009. "Knowledge and Success from Ability." *Philosophical Studies* 142 (1): 17–26.
- . 2010. *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge University Press.
- . 2012. "A (Different) Virtue Epistemology." *Philosophy and Phenomenological Research*

- 85 (1): 1–26.
- . 2020. “Safety in Sosa.” *Synthese* 197 (12): 5147–57. <https://doi.org/10.1007/s11229-018-1863-z>.
- Hetherington, Stephen. 2016. *Knowledge and the Gettier Problem*. Cambridge University Press.
- Ichikawa, Jonathan Jenkins, and Matthias Steup. 2018. “The Analysis of Knowledge.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Summer 2018. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/sum2018/entries/knowledge-analysis/>.
- Kallestrup, Jesper, and Duncan Pritchard. 2012. “Robust Virtue Epistemology and Epistemic Anti-Individualism.” *Pacific Philosophical Quarterly* 93 (1): 84–103.
- Kelp, Christoph. 2013. “Knowledge: The Safe-Apt View.” *Australasian Journal of Philosophy* 91 (2): 265–78.
- Lackey, Jennifer. 2007. “Why We Don’t Deserve Credit for Everything We Know.” *Synthese* 158 (3): 345–61.
- Littlejohn, Clayton. 2014. “Fake Barns and False Dilemmas.”
- Luper-Foy, Steven. 1984. “The Epistemic Predicament: Knowledge, Nozickian Tracking, and Scepticism.” *Australasian Journal of Philosophy* 62 (1): 26–49.
- Miracchi, Lisa, and J. Adam Carter. 2022. “Refitting the Mirrors: On Structural Analogies in Epistemology and Action Theory.” *Synthese*. <https://doi.org/10.1007/s11229-022-03462-y>.
- Nozick, Robert. 1981. “Philosophical Investigations.” *Massachusetts: Belknap Press of Harvard University*.
- Otero, Manuel Pérez. 2020. “Williamson on Defining Knowledge.” *Episteme*, July, 1–17. <https://doi.org/10.1017/epi.2020.27>.
- Pritchard, Duncan. 2005. *Epistemic Luck*. Clarendon Press.
- . 2009. “Knowledge, Understanding and Epistemic Value.” *Royal Institute of Philosophy Supplement* 64: 19–43. <https://doi.org/10.1017/s1358246109000046>.
- . 2012. “Anti-Luck Virtue Epistemology.” *The Journal of Philosophy* 109 (3): 247–79.
- . 2015a. “Anti-Luck Epistemology and the Gettier Problem.” *Philosophical Studies* 172 (1): 93–111.
- . 2015b. “Epistemic Dependence.” *Philosophical Perspectives* 29 (1): 305–24. <https://doi.org/10.1111/phpe.12067>.
- . 2015. “Risk.” *Metaphilosophy* 46 (3): 436–61.
- . 2016. “Epistemic Risk.” *The Journal of Philosophy* 113 (11): 550–71.
- Sainsbury, Richard M. 1997. “Easy Possibilities.” *Philosophy and Phenomenological Research* 57 (4): 907–19.
- Sartwell, Crispin. 1992. “Why Knowledge Is Merely True Belief.” *The Journal of Philosophy* 89 (4): 167–80.
- Shope, Robert K. 2017. “The Analysis of Knowing: A Decade of Research.”
- Sosa, Ernest. 1991. *Knowledge in Perspective: Selected Essays in Epistemology*. Cambridge University Press.
- . 1999. “How to Defeat Opposition to Moore.” *Noûs* 33 (s13): 141–53. <https://doi.org/10.1111/0029-4624.33.s13.7>.
- . 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I*. Oxford University Press.
- . 2015. *Judgment & Agency*. Oxford University Press UK.
- . 2021. *Epistemic Explanations: A Theory of Telic Normativity, and What It Explains*. Oxford University Press.

- Sylvan, Kurt. 2016. "Epistemic Reasons I: Normativity." *Philosophy Compass* 11 (7): 364–76. <https://doi.org/10.1111/phc3.12327>.
- Turri, John. 2011. "Manifest Failure: The Gettier Problem Solved." *Philosopher's Imprint* 11 (8). <http://hdl.handle.net/2027/spo.3521354.0011.008>.
- Turri, John, Wesley Buckwalter, and Peter Blouw. 2015. "Knowledge and Luck." *Psychonomic Bulletin & Review* 22 (2): 378–90.
- Vargas, Miguel Ángel Fernández. 2016. *Performance Epistemology: Foundations and Applications*. Oxford University Press.
- Weatherson, Brian. 2003. "What Good Are Counterexamples?" *Philosophical Studies* 115 (1): 1–31.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.
- . 2005. "Précis of Knowledge and Its Limits 1." *Philosophy and Phenomenological Research* 70 (2): 431–35.
- . 2008. *The Philosophy of Philosophy*. John Wiley & Sons.
- Zagzebski, Linda T. 1994. "The Inescapability of Gettier Problems." *The Philosophical Quarterly* (1950-) 44 (174): 65–73.
- . 1996. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge University Press.