

Understanding, Vulnerability, and Risk

J. Adam Carter

Abstract: A key project in mainstream epistemology investigates the sense in which beliefs are vulnerable to knowledge-undermining luck and/or risk. This chapter will explore a related but largely overlooked question of how and to what extent our grasping connections between propositions is vulnerable to understanding-undermining luck and risk. The result will be a better view of how our attempts to understand the world are vulnerable when they are, and how to better mitigate against such vulnerabilities.

Key words: understanding, risk, luck, know-how, cognitive vulnerability, ability virtue epistemology

1 Introduction

An idea familiar in recent epistemology is that – in some way – a thinker makes a kind of cognitive gain when moving from (mere) knowledge to *understanding*.¹ I know that quantum mechanics describes microphysical reality but (at least, not that we're yet aware) not macrophysical reality. I know this because experts tell me it is true. But I don't understand *why* this is true; *why* don't the laws that govern the small also govern the large? And vice versa: why doesn't gravity operate on quarks in just the same way as it does on planets and galaxies?

This is a case where it would probably take *quite a lot* to help me move from knowledge to understanding – possibly a lot of subject-specific knowledge, additional cognitive skills (e.g., more advanced mathematics), perhaps even some scientific breakthroughs.

Other cases might require less. I might know that my co-worker, after seeing the waiter, left the restaurant suddenly and inexplicably. I know this because I saw it. But I don't understand *why* my co-worker left so suddenly and seemingly explicably. Here the bridge from knowledge to

¹ For some representative discussion, see, e.g., Kvanvig (2003); Carter and Gordon (2014); Gordon (2017); Greco (2014); Grimm (2012, 2014); Lynch (2016); Pritchard (2016b, 2009).

understanding might not take much to cross at all: in coming to find out that the two were estranged former partners, it now clicks.

The ‘gap’ between knowledge and understanding accordingly varies. This is not unlike the sense in which the gap between *knowledge* and (mere) justified, true belief varies. Standard Gettier cases, for example, seem to be cases of justified true belief that *clearly* fall short of knowledge;² fake-barn style cases, however, seem to be cases of justified true belief that *less clearly* fall short of knowledge.³

One of the central themes in recent epistemology – *anti-risk* epistemology⁴ – aims to clarify the sense in which even a thinker with a justified, true belief (JTB) is nonetheless *vulnerable to ignorance*. Compelling answers will need to make sense of how the risk of being in the position of the Gettier subject differs from the (perhaps even more difficult to eliminate) risk of being like the fake barn subject. One standard approach here identifies two distinct types of knowledge-relevant cognitive vulnerabilities: *intervening epistemic risk* and *environmental epistemic risk*.⁵ The former is (in short) the risk that the fact that made your belief true has nothing to do with that in virtue of which you are justified in believing it is true; alternatively, the causal chain between the two is deviant. The latter is (in short) the risk that your belief could easily have been false, due to near-by error possibilities, *despite* the fact that there is no deviant causal chain between your believing in a way that counts as justified and your belief’s being correct.

If understanding *takes more* to achieve than knowledge, then our vulnerabilities to *its* converse (confusion⁶) are – *ceteris paribus* – going to be greater than our vulnerabilities to mere ignorance, the converse of knowledge. More precisely: suppose I have a justified true belief that X; my vulnerability to ignorance remains insofar as I am vulnerable to intervening or environmental epistemic risk. But it seems that my vulnerability to *confusion* about *why X* might not be merely characterizable in terms of these kinds of vulnerabilities. After all, even if I increase my epistemic position such that these risks are controlled for – suppose even that I *know* this is so – I might still be vulnerable to some confusion: just as I am in the case of quantum mechanics, about which I can know plenty of facts by testimony. What epistemic risks, beyond intervening and environmental

² For experimental evidence that intuition is robust across cultures, see, e.g., Machery et al. (2017).

³ A typical thought is that fake barn cases are not cases of knowledge (Pritchard 2012; Goldman 1977; Littlejohn 2014; Carter 2016); the target belief is neither safe nor sensitive. That said, the requirement of both of these modal conditions on knowledge is disputed (Sosa 2007, Ch. 2). Moreover, while folk judgments adjudicate in favour of denying knowledge in standard Gettier cases, there is a comparative willingness to attribute knowledge in fake barn cases (e.g., Colaço et al. 2014). There is also more theoretical pushback on this point (Sosa 2010; Hetherington 2016; Baumann 2014).

⁴ See, e.g., Pritchard (2012; 2015; 2020a; 2020b; 2022).

⁵ See Pritchard (2005) for an early presentation of this idea in terms of luck; cf., (2016a). See also Navarro (2021) for a recent argument for reconciling the two concepts (luck and risk) in a theory of knowledge.

⁶ For discussion on confusion as the converse of understanding in a way analogous to ignorance as the converse of knowledge, see Carter (2019).

epistemic risks, characterize *this* kind of cognitive vulnerability, one I face so long as I fail to understand?

2 Knowledge, understanding and vulnerability to epistemic risk

It turns out the preceding question is difficult because the range of ways understanding can (and can't) be risky isn't yet well understood. On one line of thought, understanding is actually *less* vulnerable to (certain) risks that knowledge is vulnerable to, given that understanding is compatible with the *obtaining* of these risk events. For example, according to Morris (2012), understanding is *invulnerable* to not only the more contentious environmental epistemic risk, but also to *intervening* epistemic risk – of the sort we find in standard Gettier cases where knowledge-that is uncontroversially not present. Consider here the following case, due to Grimm (2006):

suppose that the CIA slips a hallucinogen into Albert's coffee, and that as a result he 'sees' his dog bump into the table, causing a vase to crash to the floor. Putting things together, he takes himself to understand why the vase fell from the table and crashed to the floor: because the table was bumped by the dog. As it happens, moreover, this hallucination exactly matches the events that are actually unfolding in front of him. (Grimm 2006, 520).

Albert's justification for believing the vase fell from the floor because the dog bumped into it – if indeed we grant that he possesses any⁷ – only deviantly leads to his believing truly that this is so. Albert doesn't know *that* the dog's bumping it is why the vase fell; but as Morris sees it, we should be open to allowing that Albert nonetheless attains understanding-why when he pieces together correct explanans and explanandum which happen to match his hallucination.

Three (increasingly) less permissive takes are given by Pritchard (2009), Khalifa (2013) and Sliwa (2015). Pritchard thinks – as Grimm does – that the subject fails to understand *why* in cases like Albert's hallucination case; Pritchard accordingly takes understanding why something is so to be *vulnerable* to intervening epistemic risk, while nonetheless *invulnerable* to environmental epistemic risk. Consider the following case:

Suppose I arrive home as before to find my house burned down [...] the person I get the information from regarding the reason why this happened is genuinely a fire officer who knows what she is talking about. Suppose, however, that when I arrived back at my house there was a group of people outside who all seemed to be fire officers, and that I chose one of I these people at random to speak to about the cause of the fire. Imagine, though, that nearly all of the people outside my house are merely dressed as fire officers on their way to the fancy dress party noted earlier, and that I just happened to choose the one genuine fire officer among them. Furthermore, had spoken to one of the fake fire officers then I would have been told a false explanation of why my house had burned down, but be none the wiser. (Pritchard 2009, 108)

⁷ Phenomenal conservativists (e.g., Huemer 2006) will accept that Albert is justified in believing what he does here, as well as some other kinds of epistemic internalists; epistemic externalists (including those views that hold that evidence is factive) will reject this.

As Pritchard sees it, understanding why is simply not vulnerable to being undermined by this kind of luck; you *do* understand why the house burnt down, despite the fact that the environmental risk at play, and despite the fact that such environmental epistemic risk is – as he takes it – incompatible with propositional knowledge. Pritchard however thinks our judgment should shift to denying understanding-why if we shift the case such that the fire officer is a fake who only happened to offer a correct explanation (a case like that of Albert); here, then, we should deny both knowledge-that and understanding-why. Given this verdict, Pritchard accordingly takes understanding why to be more vulnerable to epistemic risk (in virtue of being vulnerable to intervening epistemic risk) than Morris takes it to be, even while allowing that understanding-why is not vulnerable to the kind of environmental epistemic risk that he thinks knowledge is vulnerable to being undermined by.

Khalifa offers a somewhat less permissive picture yet of understanding why: Khalifa takes it that *if* understanding is *invulnerable* to environmental epistemic risk (as Pritchard takes it to be), then it would be less valuable on account of this; as he puts it, ‘if understanding is environmentally lucky, it is always inferior to the understanding that a corresponding case of knowledge would provide’ (Khalifa 2013, 1).⁸ With reference to the fire officer case, then, the idea will be that *even if* you understand why the house burnt down when speaking with the real fire officer surrounded by fakes, you would be better off epistemically if also *knowing* this (that the house burnt down because of faulty wiring), something you’d know only if the environmental epistemic risk were eliminated (e.g., if there were no near-by fake fire officers).

Finally, *least* permissive of all is the kind knowledge-reductionist account of understanding why defended by Sliwa (2015); for Sliwa the question of whether you understand why something is so reduces to the question of whether you know why something is so; you understand why p because q iff you know that p because q. Because knowledge is vulnerable to both intervening and environmental epistemic risk, so likewise is understanding.

3 An aporia

The preceding section sketched a picture of the understanding’s vulnerability to risk which the reader may have noticed runs jarringly *contrary* to our guiding insight from §1, which was that if understanding *takes more* to achieve than knowledge by a given thinker, then, *ceteris paribus*, our vulnerabilities to *the former’s* converse (confusion) should be greater than our vulnerabilities to the *latter’s* converse, (mere) ignorance.⁹ And yet, our brief examination of a spectrum of views about understanding why looked to indicate the opposite. We seem faced, then with a puzzle, which it will

⁸ Another recent position on the spectrum is due to Belkoniene (2022); on his view, to the extent that understanding would be compatible with intervening epistemic luck, it would be suboptimal in a way that it would not be when compatible with environmental epistemic luck.

⁹ I am following Williamson (2000) in taking ignorance to be the salient contrast point with knowledge; however, nothing important for my purposes hangs on this; if the reader does not like this contrast, then feel free to replace my phrase ‘risk of ignorance’ with ‘risk of *lacking* knowledge’.

be helpful to now make more precise. Taken together, §§1-2 suggest that we should part ways with at least one of the following independently plausible claims:

- I. For all S, p, understanding why p takes more to achieve than (mere) knowledge why p (viz., than knowledge that p because q).
- II. For all S, for any two epistemic achievements vis-à-vis p, A and B, if it takes more for S to achieve A than B vis-à-vis p, then *ceteris paribus*, one is more cognitively vulnerable to the risk of lacking A than B;
- III. For all S, p, S's cognitive vulnerability to risks of lacking understanding why p is *less than or at most equal to* S's cognitive vulnerability to risks of lacking knowledge why p (viz., knowledge that p because q).

For one who wants to retain (I) – which captures something very *prima facie* plausible about the difficulty of understanding in comparison with merely knowing – then, it looks like the cost will be either rejecting (II), viz., a seemingly analytically true risk comparison principle stating (an instance) of the more general idea that the vulnerability to the risk of *lacking* something is positively correlated with the difficulty of *not* lacking it; or rejecting (III) which seems to encompass the spectrum of views (in §2) on the matter of how we are (comparatively) vulnerable to intervening and environmental epistemic risk when it comes to understanding why versus knowing why something is so.

4 Navigating the aporia

The way out is not by rejecting II. Rather, I think the way forward is to first identify *two* key ambiguities driving the puzzle, understand their sources, and then refine I and III to bring the triad into consistency. Underlying the ambiguities generating the appearance of a puzzle, as I see it, are a cluster of related ideas: first, that understanding, qua intellectual accomplishment, is a genus with (at least) two species of interest in epistemology: *holistic understanding* (i.e., Understanding X, for some subject matter X)¹⁰ and *explanatory understanding* (Understanding why p). Second, holistic understanding very plausibly *takes more* intellectually than explanatory understanding,¹¹ even if both have features with respect to which they're more demanding than knowledge. Third, while intervening epistemic risk and environmental epistemic risk are *representative* cognitive vulnerabilities, they are not *exhaustive* of such vulnerabilities (or perhaps even central – more on this shortly) when it comes to lacking understanding (especially holistic understanding), and despite the fact that intervening and environmental epistemic risk have taken center stage in discussions of epistemic risk in the theory of (propositional) knowledge.

With these points in mind, consider again (I): even if it's true – viz., that understanding why 'takes more' than knowledge – it might not take *very much* more, and not *as much more as* holistic

¹⁰ For a useful discussion of holistic understanding in epistemology, see Kvanvig (2003).

¹¹ For a defence of this point, see Carter and Gordon (2014).

understanding takes. This observation – in conjunction with the point that intervening epistemic risk and environmental epistemic risk are *representative but not exhaustive* cognitive vulnerabilities we face when it comes to understanding – offers a new perspective on (III); consider that even if one’s cognitive vulnerability to *specifically intervening and environmental epistemic* risks of lacking understanding why p is less than or at most equal to one’s cognitive vulnerability to *specifically epistemic and intervening* risks of lacking knowledge why p (viz., knowledge that p because q), it needn’t follow from this that – more generally – one’s cognitive vulnerability to risks (simpliciter) of lacking understanding why p must be less than or at most equal to one’s cognitive vulnerability to risks (simpliciter) of lacking knowledge why p (viz., knowledge that p because q); such risks might still be greater, and greater *even more* in the case of the kind of understanding (holistic) whose superiority to knowledge might have leant initial credibility to (I), framed in terms of understanding-why.

In the remainder of this chapter, I want to first (i) articulate some of the ways in which a given thinker is cognitively vulnerable to *lacking* holistic understanding; and then, (ii) with these in mind, revisit understanding why (the initial starting point) and offer a better picture of the sense in which ambitions to understand why something is so leave us cognitively vulnerable, and in at least some ways that go beyond our vulnerabilities to lacking knowledge.

5 Understanding and attempts

Consider now some structural differences between our *attempts* at propositional knowledge and at holistic understanding, respectively. Beliefs (rather than wishes, or hopes, or other kinds of attitudes) are genuine attempts at propositional knowledge in the sense that – when all goes well by way of belief – what we get is *knowledge*; when it doesn’t, we have not knowledge but *mere* belief.

Question: what stands to holistic understanding as belief stands to knowledge? “A trick question!” You might think. Might it not also be beliefs, *also*? Not quite. Here’s why. An argument for thinking that beliefs are attempts at understanding no less than at knowledge might look like this: suppose I want to understand organic chemistry, which I don’t yet understand. One thing I might do is read a book or go to university to study it, and in doing so, increase (significantly) my stock of beliefs. You might think, then, that at some point on this journey, my large stock of beliefs, if true rather than false (or: known rather than unknown) gives me the understanding of organic chemistry I lacked before. And so, in this way it might seem like I attempt to understand by believing in a way not dissimilar from how I attempt to know by believing.

The problem with this line of reasoning is that it conflates instrumental and constitutive attempts.¹² Beliefs are constitutive attempts at knowing in that believing does *more* than (as mere instrumental attempts at knowing do) put you *in a position* to know. (Compare: collecting evidence about the accused thief might be an instrumental attempt at knowing whether they are guilty).

¹² For discussion, see Sosa (2021, Ch. 2) and Carter (2022a).

Beliefs aren't attempts at knowing like gathering evidence is an attempt at knowing; when all goes well when you gather evidence, you are *in a position* to know; when all goes well when you believe, you *know*.

On a mostly uncontroversial assumption, which I will grant, about holistic understanding, it requires the taking up of an *objectual attitude* (viz., an attitude toward a body of information) rather than of a *propositional attitude relation*. (Holistic understanding is often referred to *as* objectual understanding – understanding *of something*). The job description for an attempt at understanding will need to be some kind of objectual attitude relation (relating a thinker to a body of information) that, when all goes well, results in understanding (and not merely in putting one in a position to understand). When all does *not* go well, when one makes an attempt at understanding, what is residual is a kind of *mere* conception (botched understanding): compare with Williamson's suggestion that *mere* belief is 'botched knowledge' (Williamson 2000, 47).

6 An intervention: aptness structure

In recent co-authored work (see Miracchi and Carter 2022) I've defended the following analogy: a belief (attempt) is to knowledge (its realization) as a *conception* (attempt) is to *understanding* (its realization). Williamson, in the spirit of knowledge-first epistemology, can, and *does*, embrace the former idea (relating belief and knowledge as attempt/realization)¹³; so, signing on to this much is compatible with the idea that 'realisations' are theoretical primitives, not analysable in terms of their corresponding attempts. However, the view defended by Miracchi and Carter goes a step further – adding 'aptness-theoretic' structure to realization/attempt pairs – across both theoretical and practical life. On our view, knowledge is *apt* belief, understanding is *apt* conception. Action is apt intention.¹⁴

Aptness structure is familiar from telic virtue epistemology¹⁵ – on which epistemic normativity is the normativity of attempts *as* attempts. Aptness structure is characterized, most simply, as a structure an attempt has when its success (viz., success in securing the aim internal to that attempt as such) manifests competence. The idea is of special interest in epistemology in part because, among other things, if knowledge is identified with apt belief, then there is a ready-made explanation for why knowledge is lacking in (standard) Gettier cases: in such cases, the subject's *belief* manifests competence, but its *correctness* does not.

Belief's vulnerability to ignorance (on the view that knowledge is apt belief) lines up with a belief's vulnerability to *falling short* of aptness. This idea offers now a useful vantage point for gaining a footing on the ways in which we are, in attempting to understand, vulnerable to risk of

¹³ On Williamson's picture (e.g., 2017) belief is to knowledge as intention is to action; in both cases the former (belief and intention) stands to the latter (knowledge and action) as realisation to attempt (where knowledge and action have reverse directions of fit, mind to world, world to mind).

¹⁴ See also Sosa (2015, Ch. 1) and Carter (2022b).

¹⁵ See Sosa (2021).

lacking such understanding (confusion). The idea, in outline is as follows: a belief (attempt, via a propositional attitude) is vulnerable to ignorance (the converse of its realization, knowledge) in so far as it is at risk of failing to be apt, accurate (qua that kind of attempt) through competence; as a conception (attempt, via an objectual attitude) is vulnerable to confusion (the converse of its realization, understanding) in so far as it is at risk of failing to be apt, accurate (qua that kind of attempt) through competence.

7 Aptness structure extended

Let's take this idea further now. Aptness-structure, as applied to belief, is well understood, but it is not well understood as applied to an objectual attitude, like a conception. Truth is 'accuracy' in the case of belief (and then we explain competence by appeal to dispositions of a thinker to reliably secure that aim, truth, in appropriate conditions).¹⁶ What is mere accuracy in the case of a conception? Working with the idea that believing involves affirming content (of a proposition) in a way analogous to how conceiving (or forming a conception) of a body of information involves not just affirming content but *grasping* relations between propositional content, an accurate conception will then *at least* involve the body of information grasped being either all or at least mostly true; and, further, since one grasps a body of information only if forming some kind of view about how some of that information relates to each other (e.g., this is usually understood as coherence relations or explanatory relations), an accurate conception will implicate also that all or most of what one grasps *about* the relationships between the information must itself be accurate. We needn't get bogged down on the 'all or mostly' point – as the question of whether and how understanding is 'factive' is itself deeply contentious.¹⁷ I will assume at least, though, and contra strong non-factivists about understanding,¹⁸ that you fail to understand something, X, if all or most of the propositions grasped about X are false, or if propositions corresponding with your grasp of the connections (e.g., coherence, explanatory) *between* the relevant X-propositions are not accurate, coherence/explanatory connections between that information. This point then offers a minimal idea of what the accuracy of a conception will involve.

The idea of a *competent conception* can then be understood in terms of the above simple view of an accurate conception. Just as a belief is competent when it is formed via a disposition to believe *truly* (where truth is accuracy, for belief) reliably enough in normal circumstances, a conception is competent iff formed via a disposition to reliably form conceptions that are accurate (that is, now, when the body of information one forms a conception of is all or mostly true along with propositions corresponding with your grasp of the connections between this body of target information). The above, to be clear, is meant simply as an *application* of the general idea that, for any attempt (of which forming a belief and a conception are instances), an attempt is competent just

¹⁶ For discussion of the idea that knowledge-relevant competences are special cases of dispositions more generally, see Sosa (2010).

¹⁷ For some discussion, see, e.g., (Gordon 2017; 2021); Carter and Gordon (2016); Kvanvig (2003).

¹⁸ See, e.g., Elgin (2017).

when it manifests a disposition to reliably enough secure the aim internal to that attempt-type. With reference to this idea, then, if we run a holistic understanding variation on Grimm's case of Albert – suppose that post-hallucinogen Albert forms an accurate conception – the conception formed would not be competent *even if it is accurate*, as forming such a conception in that manner (drug-induced) almost always issues in (mostly false) conceptions, either conceptions directed towards false information or which display an incorrect assessment of explanatory/coherence relations between that information.

Since aptness is a matter of accuracy *through* competence, we now can fill out more substantively what (holistic) understanding would involve construed as an *apt conception*. A thinker in the market for holistic understanding, of some body of information or subject matter, X, will need to first make the kind of attempt apposite to holistic understanding of X, to form an X-directed conception, not by affirming (as we do when forming a belief) but by means of grasping (which I'm not using as a success term here) a body of X-information, which (since understanding is gradient) needn't be perfectly exhaustive of X. This attempt is then *successful* if the relevant information is accurate, in the sense described above, and *apt* when its accuracy manifests the thinker's disposition to *reliably* form such accurate conceptions.

8 Cognitive vulnerability and understanding-relevant risk

With the above idea of (holistic) understanding as *apt conception* in play, we can – from just the minimal substantive gloss just given – already taxonomize some key *varieties of vulnerability* to the kind of confusion (towards some body of information) one faces when at risk in different ways of lacking such understanding. Entirely apart from any vulnerabilities to confusion that would correspond with either intervening and/or environmental risk of the sort discussed in §2, we can identify a thinker as vulnerable to *at least* two interesting kinds of epistemic risks that apply in the case of (holistic) understanding, where an attempt is made via forming a conception, rather than via a belief.

Consider the following two cases:

PARTICLE ACCELERATOR: Ursula, who works at CERN, aims to understand the laws governing a certain kind of particle, P, and so studies the properties of the P-type particles by smashing them against other particles using a particle accelerator. She gathers data, and over time, forms a conception about how all this data fits together: she posits law L as the fundamental law governing the behaviour of P-type particles. She now *thinks* she understands the laws governing P-type particles (namely, that they behave in accordance to laws that are either L itself or derivable from L). Unfortunately, while all her data is correct (and so she has all true beliefs about the behaviour of the particles), *and* her beliefs about how this data fits together through coherence/explanatory relationships are formed competently (they draw from impeccable inductive/deductive principles, etc.), it is just by coincidence that the particles observed happened to be observed under particular configurations such that they appear to behave in a way that would be explained by law L.

They are actually explained by a different law, L^* , something she would have worked out only had she done an unusually large number of experiments for forming that kind of conception.

ULYSSES: Ulysses feels he doesn't understand his new friend, Jeb, whose behaviour has become increasingly mysterious. With determination, Ulysses begins considering what he knows about Jeb and begins trying to piece it together. Unbeknownst to Ulysses, Jeb's rogue therapist (who has considerable time, ability, and resources) has attempted to 'engineer' Jeb's epistemic and moral environments, so as to create physical/social environments where Jeb will be inclined to be ponderous. (The therapist, suppose, believes Jeb needs to be more reflective). Ulysses begins to form a conception of Jeb, on the basis of his observations, which is that Jeb is a very ponderous person, where this conception relies on grasping some coherence/explanatory relations applicable to the data gathered. Jeb *is* in fact very ponderous, it turns out, (robustly so, across a wide set of circumstances), even though *anyone* would have appeared ponderous to an observer who observed the individual in the limited circumstances engineered by the therapist, the only circumstances under which Ulysses observed Jeb, and even though Ulysses wouldn't have formed this conception from the comparatively paltry data set Ulysses would have acquired had the therapist not intervened.

Here are two points to note about these cases. First, note that *mere believers* aren't, through just believing (whether p , on a given question), subject to risks of confusion of the sort we find in these two cases because believers are (as such, through merely forming a belief) *in the market* for such risks; they are not taking up the kind of attitudes that leave one vulnerable to lacking holistic understanding in either ways. The price of entry for such risks is the taking up of an objectual attitude, not merely a propositional attitude.

Second, notice that in the two cases, vulnerability to *different* risks – risks of lacking an apt conception – are present.

The risk of being in the former case is a risk of aptly drawing the *wrong* conception, gaining understanding of the wrong thing, remaining confused about what one aimed to understand, and mistaking understanding the former for the latter. What Ursula *thinks* she understands through her conception is the law governing the behaviour of P-type particles; what she actually understands is a law-like regularity applicable to the behaviour of P-type particles *only in very specific configurations, those they happen to be in*. She lacks an apt conception (understanding) of the right kind of phenomenon; by taking up an objectual attitude, we likewise render ourselves vulnerable in principle to this kind of confusion (lack of understanding), one that might easily not be noticed by a thinker.

The second case bears a kind of 'safety without aptness' structure. Ulysses lacks understanding of the relevant phenomena (Jeb) but for a different reason. The reason is that his *accurate* conception of Jeb is not explained by his relevant competence to draw such a conception, but rather by the fact that the external helper would have made it nearly inevitable that a thinker

observing Jeb in those conditions would *not* draw that conception, if forming any conception about him at all. In short, Ulysses has a *safe* conception but not an *apt* conception; in so far as holistic understanding is an achievement implicating more than just an accurate and safe conception (one that could be engineered externally), Ulysses lacks holistic understanding of Jeb; we, likewise, render ourselves vulnerable in principle to this kind of confusion in taking up objectual attitudes – especially in epistemically paternalistic environments.¹⁹

9 The aporia resolved

Forming conceptions always involves grasping relations (e.g., explanatory or coherence) between a variety of different propositions about a subject matter, not merely affirming some particular proposition (whether *p*). Forming a belief, by contrast, always requires affirmation of a particular proposition.²⁰ But, *sometimes*, forming a belief also involves grasping explanatory relations between multiple propositions – in the case that will now be of interest: the propositions that line up with *explanandum and explanans in a given explanation* of the form ‘A because B’. Take again the target proposition in Pritchard’s faulty wiring example: “The house burnt down because of faulty wiring.” A thinker is a candidate for *knowledge* by forming a belief about this (the house burnt down because X, the house burnt down because Y, etc.), perhaps on testimony. But – and this is a point that has been made by Pritchard and by Hills (2009) – one doesn’t always secure *understanding-why* simply by affirming such content. When the child hears by testimony from a parent that the house burnt down because of faulty wiring, the child might believe on trust, without forming any conception of how the latter might be the sort of thing that would cause the former. More generally: not just children, but adults, can *believe* and plausibly *know* causal propositions of the form A because B, by trusting an expert, without grasping how B is the sort of thing that would cause A. You just know that B *did* cause A.

Returning to Pritchard’s case: a more cognitively sophisticated adult might do *two things* upon hearing the testimony that faulty wiring caused the house to burn down. First, they might (and will) do what the child does, which is to form a belief with that content by *affirming it* on trust. But they might also form a *limited conception* – by way of grasping how two propositions are explanatorily related. I say ‘limited’ conception because, whereas holistic understanding (especially in the case of complex subject matters) involves grasping explanatory and/or coherence making relations between perhaps many different propositions, there is – in the case where understanding why is what’s at issue – really just one very specific explanatory relation the grasping of which is plausibly *necessary* for understanding-why: namely, the relation between explanans and explanandum.

¹⁹ One experimental example that is useful for illustrating this idea is the Duncker candle test. See, e.g., Duncker and Lees (1945).

²⁰ This of course might include tacit affirmation, as in the case of ‘functional’ beliefs (Sosa 2015).

We're now in a position to return to our original discussion, framed in terms of understanding-why, and the puzzle it generated. When one merely *believes* by affirming content of the form "A because B" one is thereby opening up oneself to the standard kinds of epistemic risks (risks of ignorance) one makes oneself vulnerable to just by believing, and this includes intervening and environmental epistemic risk. However, by putting ourselves in the market for not just testimonial knowledge of a proposition of the form "A because B" but in the market for explanatory understanding – by forming a limited *conception* (by attempting a grasp of the explanatory connection between explanandum and explanandum) – one is at the same time making oneself *additionally* vulnerable, vulnerable to additional risks of confusion.

We can appreciate this point by simply bringing into connection (i) a thinker who *not only* affirms that the house burnt down because of faulty wiring, but who forms a conception about the explanatory relatedness of the former and the latter; and (ii) our cases PARTICLE ACCELERATOR and ULYSSES, cases of (two kinds of) vulnerability to confusion distinctive to thinkers who attempt *some kind of conception*, over and above forming a belief. What we can see now is the kind of thinker just described in (i) renders herself vulnerable to a limited version of the kinds of risks that obtain in ACCELERATOR and ULYSSES, even if she would *not* render herself so vulnerable were she to affirm the explanatory proposition, accepting it on trust. That is, just as Ursula draws the *wrong conception*, so might one who believes (truly) that the house burnt down from faulty wiring, but who forms a conception of *how* wiring would have caused that fire that (suppose) is of the wrong level of generality than one takes it to be. Likewise, we can easily imagine ULYSSES-style vulnerability applicable to the thinker who goes beyond just believing the target proposition, but who forms a conception of how the faulty wiring caused the house fire, but where the level of external engineering of the situation would have led most anyone to form this same conception, even in the absence of any competence; and suppose further that absent the engineering, our subject would have continued to believe truly that faulty wiring caused the fire, but whose competence limits would not have furnished them with an appreciation of how this were so absent the deliberate engineering by the helper.

What the foregoing suggests, then, is that the very thing that puts one in the market for not only believing (and if all goes right knowing) an explanatory proposition – namely, a thinker's forming a conception of the explanatory relatedness of explanandum and explanans – *also* renders one vulnerable to additional epistemic risks, risks not only of ignorance (which one *de facto* risks even when only trusting the testifier enough to affirm the causal proposition) but of additional confusion. We've seen now some ways that forming a conception open one up to such risks (not applicable to a mere believer) by looking first at holistic understanding, where the risk of confusion (by lacking that kind of understanding) is even more amplified. More amplified because what one forms a conception of, in cases of holistic understanding, one does so through grasping *many different* explanatory/coherence relations, potentially between many propositions. In short, *more confusion is risked*. Less, but some, such confusion is risked always when one goes beyond mere believing (by affirming) to form a conception *through* grasping explanatory/coherence relations at

all. And that is just what one does when one puts oneself in the market for understanding why. In this way, understanding why does *take more* than mere knowing why; it takes the forming of a conception (that minimally involves attempting to grasp how the cause would explain the effect) in a way that mere believing and knowing, attainable via affirmation through testimony, does not, and it takes that moreover that conception's being *apt*. Furthermore, we can now even more clearly see why III in the puzzle in §3 needed to be modified. That is, even if (and as per §2) one's cognitive vulnerability to *specifically intervening and environmental epistemic* risks of lacking understanding why p is less than or at most equal to one's cognitive vulnerability to *specifically epistemic and intervening* risks of lacking knowledge why p (viz., knowledge that p because q), it needn't follow from this that – more generally – one's cognitive vulnerability to risks (simpliciter) of lacking understanding why p must be less than or at most equal to one's cognitive vulnerability to risks (simpliciter) of lacking knowledge why p (viz., knowledge that p because q).

Acknowledgements

Carter's research is supported by the Arts and Humanities Research Council's *Expanding Autonomy* (AH/W005077/1) and *Digital Knowledge* (AH/W008424/1) projects

References

- Baumann, Peter (2014): "No Luck with Knowledge? On a Dogma of Epistemology." In: *Philosophy and Phenomenological Research* 89 (3), 523–51.
- Belkoniene, Miloud (2022): "Reassessing Lucky Understanding." In: *Episteme*, July 1–15. <https://doi.org/10.1017/epi.2022.23>.
- Carter, J. Adam (2016): "Robust Virtue Epistemology as Anti-Luck Epistemology: A New Solution." In: *Pacific Philosophical Quarterly* 97 (1), 140–55.
- (2019): "Radical Scepticism and the Epistemology of Confusion." In: *International Journal for the Study of Skepticism* 9 (3), 223–37.
- (2022a): "Epistemic Normativity Is Not Independent of Our Goals." In: Ernest Sosa, Matthias Steup, John Turri, and Blake Roeber (eds.), *Contemporary Debates in Epistemology, 3rd Edition*, Malden, MA.: Wiley-Blackwell (forthcoming). <https://philarchive.org/rec/CARENI-4>.
- (2022b): "Trust and Trustworthiness." In: *Philosophy and Phenomenological Research* 1–18. <https://doi.org/10.1111/phpr.12918>
- Carter, J. Adam, and Emma C. Gordon (2014): "Objectual Understanding and the Value Problem." In: *American Philosophical Quarterly* 51 (1), 1–13. <https://www.jstor.org/stable/24475366>.
- (2016): "Objectual Understanding, Factivity and Belief." In: Martin Grajner and Pedro Schmechtig (eds.), *Epistemic Reasons, Norms and Goals*, Berlin, Boston: De Gruyter, 423–442. <https://doi.org/10.1515/9783110496765-020>

- Colaço, David, Wesley Buckwalter, Stephen Stich, and Edouard Machery (2014): “Epistemic Intuitions in Fake-Barn Thought Experiments.” In: *Episteme* 11 (2), 199–212.
- Duncker, Karl, and Lynne S. Lees (1945): “On Problem-Solving.” In: *Psychological Monographs* 58 (5): i-113. <https://doi.org/10.1037/h0093599>
- Elgin, Catherine Z (2017): *True Enough*, Cambridge, MA.: The MIT Press.
- Goldman, Alvin I. (1977): “Discrimination and Perceptual Knowledge.” In: *The Journal of Philosophy* 73 (20), 771–91.
- Gordon, Emma C. (2017): “Understanding in Epistemology.” In: *Internet Encyclopedia of Philosophy*. <https://iep.utm.edu/understa/> (accessed: 4 November 2022).
- (2021): “Elgin on Understanding: How Does It Involve Know-How, Endorsement and Factivity?” In: *Synthese* 198 (6), 4955–72.
- Greco, John (2014): “Episteme: Knowledge and Understanding.” In: Kevin Timpe and Craig A. Boyd (eds), *Virtues and Their Vices*, Oxford: Oxford University Press, 285–302.
- Grimm, Stephen (2006): “Is Understanding a Species of Knowledge?” In: *The British Journal for the Philosophy of Science* 57 (3), 515–35.
- (2012): “The Value of Understanding.” In: *Philosophy Compass* 7 (2), 103–17.
- (2014): “Understanding as Knowledge of Causes.” In: A. Fairweather, A (ed.), *Virtue Epistemology Naturalized*, Dordrecht: Springer, 329–45.
- Hetherington, Stephen (2016): *Knowledge and the Gettier Problem*, Cambridge: Cambridge University Press.
- Hills, Alison (2009): “Moral Testimony and Moral Epistemology.” In: *Ethics* 120 (1), 94–127. <https://doi.org/10.1086/648610>.
- Huemer, Michael (2006): “Phenomenal Conservatism and the Internalist Intuition.” In: *American Philosophical Quarterly* 43 (2), 147–58.
- Khalifa, Kareem (2013): “Understanding, Grasping and Luck.” In: *Episteme* 10 (1), 1–17.
- Kvanvig, Jonathan L. (2003): *The Value of Knowledge and the Pursuit of Understanding*, Cambridge: Cambridge University Press.
- Littlejohn, Clayton (2014): “Fake Barns and False Dilemmas.” In: *Episteme* 11 (4), 369–389.
- Lynch, Michael P. (2016): *The Internet of Us: Knowing More and Understanding Less in the Age of Big Data*, New York, New York: WW Norton & Company.
- Machery, Edouard, Stephen Stich, David Rose, Amita Chatterjee, Kaori Karasawa, Noel Struchiner, Smita Sirker, Naoki Usui, and Takaaki Hashimoto (2017): “Gettier Across Cultures1.” In: *Nous* 51 (3), 645–64. <https://doi.org/10.1111/nous.12110>.
- Miracchi, Lisa, and J. Adam Carter (2022): “Refitting the Mirrors: On Structural Analogies in Epistemology and Action Theory.” In: *Synthese* 200 (9), 1–28. <https://doi.org/10.1007/s11229-022-03462-y>.
- Morris, Kevin (2012): “A Defense of Lucky Understanding.” In: *The British Journal for the Philosophy of Science* 63 (2), 357–71. <https://doi.org/10.1093/bjps/axr023>.
- Navarro, Jesús (2021): “Epistemic Luck and Epistemic Risk.” In: *Erkenntnis*, 1–22. <https://doi.org/10.1007/s10670-021-00387-9>
- Pritchard, Duncan (2005): *Epistemic Luck*, Oxford: Clarendon Press.
- (2009): “Knowledge, Understanding and Epistemic Value.” In: *Royal Institute of Philosophy Supplement* 64, 19–43. <https://doi.org/10.1017/s1358246109000046>.
- (2012): “Anti-Luck Virtue Epistemology.” In: *The Journal of Philosophy* 109 (3), 247–79.
- (2015): “Risk.” In: *Metaphilosophy* 46 (3), 436–61.

- (2016a): “Epistemic Risk.” In: *The Journal of Philosophy* 113 (11), 550–71.
- (2016b): “Seeing It for Oneself: Perceptual Knowledge, Understanding, and Intellectual Autonomy.” In: *Episteme* 13 (1), 29–42.
- (2020a): “Anti-Risk Epistemology and Negative Epistemic Dependence.” In: *Synthese* 197 (7), 2879–94.
- (2020b): “Anti-Risk Virtue Epistemology.” In: Christoph Kelp and John Greco (eds.), *Virtue-Theoretic Epistemology: New Methods and Approaches*, Cambridge: Cambridge University Press, 203–24.
- (2022): “Varieties of Epistemic Risk.” In: *Acta Analytica* 37 (1), 9–23.
- Sliwa, Paulina (2015): “IV—Understanding and Knowing.” In: *Proceedings of the Aristotelian Society* 115, 57–74.
- Sosa, Ernest (2007): *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I*, Oxford: Oxford University Press.
- (2010): “How Competence Matters in Epistemology.” In: *Philosophical Perspectives* 24 (1), 465–75.
- (2015): *Judgment & Agency*, Oxford: Oxford University Press.
- (2021): *Epistemic Explanations: A Theory of Telic Normativity, and What It Explains*, Oxford: Oxford University Press.
- Williamson, Timothy (2000): *Knowledge and Its Limits*, Oxford: Oxford University Press.
- (2017): “Acting on Knowledge.” In: J. Adam Carter, Emma C. Gordon, and Benjamin W. Jarvis (eds.), *Knowledge First: Approaches in Epistemology and Mind*, Oxford: Oxford University Press, 163–81.