

The Value of Knowledge and The Test of Time

The ‘Problem’

The fast growing literature on the value of knowledge stems from a compelling Pre-theoretical Intuition: Knowledge is more valuable than mere true belief. This Pre-theoretical Intuition gives rise to the Value Question: What makes knowledge more valuable than mere true belief? And that question, finding no immediate answer, gives rise to the Value Problem: The problem we can seem to have in answering the Value Question. Our primary difficulty in answering the Value Question is that when we look at any standard example of a mere true belief, and compare its value with the value of the correlative knowledge state, it is not immediately clear that knowing *p* is any more valuable than merely truly believing *p*. Let’s rehearse a standard sort of example. You wake up in the night to the loud bleeping of the smoke alarm. You form the belief that there’s a fire; so you immediately get everyone out safe and dial 999. As it happens, your belief is true, for there is a fire in the basement; but the smoke alarm is faulty and went off at random. You have a true belief, but lack knowledge. So what? What greater value would a state of knowledge have been? You got everyone out and dialled 999. The value bestowed on a mere true belief by the fact that it is true seems to exhaust the value of the counterpart knowledge. Here we confront the Value Problem.

It all started with *Meno*. Socrates and Meno have been discussing whether a person’s being good is a matter of knowledge or not, and Socrates is proposing that being good, and being able to show others the right path, might rather be a matter of true opinion:

Socrates: Look—suppose someone *knew* the way to Larissa (or wherever) and was on his way there, and showing other people how to get there; obviously he’d be good at showing them the right way?

Meno: Of course.

Socrates: And what about someone who had an *opinion* on how to get there – a correct opinion – but who’d never actually been there, and didn’t know how to get there; wouldn’t he be able to show them the way as well?

Meno: Of course.

Socrates: ...With his true belief, but without knowledge, he’ll be just as good a guide as the man with the knowledge?

[Meno agrees.]...

Socrates: So in other words, a correct opinion does just as much good as knowledge?

This last question inspires some fleeting resistance from Meno, but soon gives rise to Meno’s famous question about the value of knowledge, a question which has inspired much of the recent literature.

Meno: Except in one respect, Socrates. If you have knowledge, then you’ll *always* be dead on target; but if you only have a correct opinion, sometimes you’ll hit, and sometimes you’ll miss.

Socrates: What makes you say that? If you’ve always got the correct opinion, won’t you always be ‘on target’ as long as you’ve got your correct opinion?

Meno: Yes, good point...it seems that must be right; which leaves me wondering, Socrates: If that’s the case, why on earth is knowledge so much more valuable than correct opinion, and why are they treated as two different things?¹

The way Meno puts it, in his conjunctive question at the end here, suggests that whatever makes knowledge more valuable than correct opinion is the same thing that crucially differentiates the two. Some version of this idea is surely right, but I shall argue that the particular way in which the idea is played out in the literature helps to distort the debate, and effectively conceals at least one of the most fundamental aspects of the value of knowledge. My principle aims here will be to identify two key presumptions that together effect the distortion and concealment; and to give a positive account of what I take to be one of the most basic values of knowledge—a

¹ Plato, *Protagoras* and *Meno*, trans. Adam Beresford (London: Penguin, 2005); 129.

value that Socrates points to in the answer he goes on to give to Meno's question, but which can only be missed or misconstrued within the confines of much of the current debate.

The Diagnosis: Two Unwarranted Presumptions

In the literature we see the value problem crystallizing into a highly specific shape. And the contributions are partisan in terms of the general epistemological team that the contributor is on. The value problem seems to present itself to most who tackle it as a challenge and an opportunity to advance whatever particular epistemological theory they espouse. Indeed, the value problem—very distant now from its origination in Meno's epistemologically innocent value question—has become something of a modern epistemological football. This has two disadvantages: any proposed solution is hostage to epistemological fortune in that it stands or falls along with the particular analysis of knowledge that issues it; and it encourages players to look for the value of knowledge in something that distinguishes their theory of knowledge from their competitors' theories, when in fact the basic value of knowledge may be better explained by reference to something less epistemologically specific. Spectators to the literature have seen a movement away from the most basic reliabilist line, and a surge in the general direction of credit accounts of one or another stripe. Given how the ground-rules of the game have developed, credit accounts come to seem admirably well kitted out to solve the problem. They are; but I believe that the way the ground-rules have developed distorts the natural philosophical question, so that we have ended up with a somewhat artificial game. In order to explain what I mean, I shall describe the general trajectory of the literature, and then give my diagnosis of the pressures that give it the peculiar shape it now has.

There is a range of different credit accounts, but the common idea is that what gives knowledge its special value is the credit that is transferred to the knowledge state from the agent for achieving his true belief in the manner requisite for knowledge. Quite what that manner is depends on the particular stripe of the credit account. (On John Greco's agent reliabilist view, for instance, the subject's true belief must be due to some stable trait of cognitive character; on Ernest Sosa's view, the true

belief must be ‘attributable’ to the knower as his own doing; and Duncan Pritchard argues, in this volume, in favour of an agent reliabilism supplemented by a safety condition.²) At the virtue epistemological end of the spectrum is the view, advanced by Linda Zagzebski, that the agent’s credit worthiness is a matter of her good epistemic motive, most fundamentally, her love of truth. I shall focus on Zagzebski’s account³ because it provides a good illustration of both how satisfying an account of the value problem can be within the framework of the current debate, while simultaneously exposing the features of that framework that I want to highlight and reveal as unduly limiting the range of answers we might give to the value question.

She sets up the issue by considering and rejecting reliabilist responses to the value question. Reliabilism says that a true belief arrived at by a reliable process or faculty is more valuable than a true belief arrived at in any other way, and that added value is the value of knowledge. But, argues Zagzebski, this answer does not work, because reliability is only as valuable (or disvaluable) as that which it produces. Reliability *per se* has no value. She invokes an example to bring the point home: a great espresso made from a reliable espresso machine is no more valuable than one made from an unreliable machine. A great espresso is a great espresso; a true belief is a true belief. This argument is justly challenged by Pritchard⁴, who points out that it assumes there are only two kinds of value—intrinsic and instrumental—whereas in fact there is a third category of value, sometimes called ‘final’ value. If something has

² See, in particular, John Greco, ‘Knowledge as Credit for True Belief’ in Michael DePaul and Linda Zagzebski (eds.) *Intellectual Virtue: Perspectives from Ethics and Epistemology* (Oxford: Clarendon Press, 2002); Ernest Sosa, ‘The Place of Truth in Epistemology’ in DePaul and Zagzebski eds.; and Duncan Pritchard, ‘Knowledge and Value’, this volume.

³ I shall focus in particular on Zagzebski, ‘The Search for the Source of Epistemic Good’, in Michael Brady and Duncan Pritchard (eds.), *Moral and Epistemic Virtues* (Oxford: Blackwell, 2003); but see also her earlier paper, which makes similar negative arguments against forms of reliabilism, though is less worked out in terms of her own position: ‘From Reliabilism to Virtue Epistemology’ in Guy Axtell (ed.) *Knowledge, Belief, and Character* (Lanham, Maryland: Rowman and Littlefield, 2000).

⁴ Duncan Pritchard, ‘Knowledge and Value’. For related criticisms, see also Philip Percival’s response to Zagzebski, ‘The Pursuit of Epistemic Good’, in Brady and Pritchard (eds.), *Moral and Epistemic Virtues* (Oxford: Blackwell, 2003).

final value, we value it to that extent for its own sake (and so non-instrumentally) but not in virtue of its intrinsic properties. Whereas intrinsic value is possessed in virtue of intrinsic properties, and instrumental value accrues in virtue of what something is a means to, final value is possessed in virtue of other relational properties. Granted that reliability in itself has no value, still the reliabilist could claim that a true belief reliably produced is valuable for its own sake in virtue of certain relational properties. In the case of agent reliabilism, for instance, the relational property in question might be that of being produced by a stable trait of intellectual character. Certainly that looks like a plausible claim of value, and it is one not catered for by Zagzebski's line of attack. Given the existence of final value as a species of value, then, Zagzebski is not entitled to assume that reliability's lack of intrinsic value means it is impossible that some kind of reliability in how true beliefs are generated could constitute the value of knowledge; for the value of knowledge might yet turn out to be owing to relational properties associated with epistemic reliability. But I will not dwell on this, as my main purpose lies elsewhere.

In Zagzebski's discussion, having dispensed with reliabilism, she goes on to press the positive case for her virtue epistemological solution to the value problem. Seized by the question how a component of knowledge can *transfer* value to the knowledge state itself, she pursues the idea that just as, in general, good motives add value to the acts that they produce, so do good epistemic motives add value to the acts of belief that they produce. A true belief motivated by a good epistemic motive thus acquires the added value of the good motive: and that's the special value of knowledge. But, she observes, there can of course be cases where the true belief achieved is in itself not worth having, for the content of a true belief might be trivial, or in various ways bad. Illustrating trivial true belief, she invokes Sosa's example: 'At the beach on a lazy summer afternoon, we might scoop up a handful of sand and carefully count the grains...' (Sosa, 2003, 156).⁵ Illustrating bad true beliefs, she mentions 'knowing exactly what the surgeon is doing to my leg when he is removing

⁵ Ernest Sosa, 'The Place of Truth in Epistemology' in Michael DePaul and Linda Zagzebski eds. *Intellectual Virtue: Perspectives from Ethics and Epistemology* (Oxford: Clarendon Press, 2003); 156.

a skin cancer; knowing the neighbours private life'.⁶ Still, argues Zagzebski, in all such cases, the agent gains a certain credit for the good epistemic motive that led her to acquire the belief, and so that which renders her true belief knowledge is *admirable*. This admirability is to be distinguished from *desirability*, which is a matter of the content of one's cognitive state being worth having (not trivial or worse than trivial). Not all knowledge is desirable; but all knowledge is admirable. A particularly valuable kind of knowledge concerns true beliefs that are both desirable and admirable—knowledge worth having; and the best kind of knowledge (a 'great good'⁷) is when not only the admirability but also the desirability of the true belief can be credited to the agent—knowledge acquired by the agent because it is worth having.

Given the way the issue shapes up, Zagzebski's proposed solution to the value problem presents itself as a satisfyingly subtle and differentiated proposal, albeit dependent on one's accepting a virtue-based analysis of knowledge. But I think there is a deep problem associated with her approach to the value question, and which is a generic problem with credit accounts. In short, they put the cart before the horse. Credit accounts purport to explain the fact that we value knowledge by pointing to the value of this or that form of epistemic creditworthiness—on Zagzebski's view the credit that transfers in all cases of knowledge to render it admirable (even where the content fails to render it desirable) is owing to the good epistemic motive that helps transform true belief into knowledge. But the idea that we value knowledge *because* we value good epistemic motive gets the order of explanation back to front. We do not value knowledge because we value good epistemic motive. Rather, we value good epistemic motive because we value the knowledge it tends to get us. Indeed such a motive only constitutes a good epistemic motive because it aims at knowledge or truth or some other suitably ultimate epistemic end. (The general point can be made equally well in terms of true belief rather than knowledge: we value good epistemic motive because of the truth it tends to get us. Either way, the point is that the value of good epistemic motive is most naturally to be explained by reference to the value of what it

⁶ Zagzebski (2003), 21.

⁷ Zagzebski (2003), 24.

gets us, and not the other way around.) So to suggest that we most fundamentally value knowledge because we value good epistemic motives, or whatever else is suggested as earning the relevant credit, is to put the cart before the horse.

The natural order of explanation marks a point of disanalogy between virtue epistemology and virtue ethics. In virtue ethics, it is a thoroughly plausible idea that the value of the various goods that virtues aim at cannot be specified independently from the values of the good motives animating the virtues. It is entirely plausible to say there is a non-vicious circularity in how we characterize these values—the virtuous agent is motivated towards the good, and the good cannot be specified independently from what motivates the virtuous person. But this becomes, at best, a far less plausible idea when transferred to the field of epistemological value. For it is all too easy to specify the value of truth, and thereby the knowledge that captures it for us, in purely practical terms without reference to our epistemic motives: we need plenty of true beliefs in order to successfully pursue our practical and other purposes in life. (We don't need *all* our beliefs to be true, of course; there are exceptions. We can generally afford to have a few false beliefs knocking around without any real consequence; in some circumstances, we might be pragmatically better served by some false simplifications of the truth; and sometimes we might personally need a fairly substantial false belief in order to be able to face another day. But the basic point stands.) Approached from this angle, the value problem presents itself as the question what good it does us to possess true beliefs specifically *as* knowledge, and so the question of knowledge's value now seems more adverbial: what is the greater value of possessing truths in the manner of knowledge?

I think an adverbial formulation captures the right way to approach the issue, but advocates of credit accounts do not naturally approach it in that way because they are committed to identifying the value of knowledge in something that does not reduce to the value of truth. That is one of the purported lessons of the critique of the reliabilist solution to the value problem: 'If the feature that converts true belief into knowledge is good just because of its conduciveness to truth, we are left without an explanation of why knowing *p* is better than merely truly believing *p*.'⁸ This can seem

⁸ Zagzebski (2003), 17.

to flow from the critique of reliabilism, but even disregarding the reservations about that critique to which Pritchard's objection gives rise, the lesson really only flows given a certain unwarranted presumption about how to frame the value question. The credit approach, and the value problem literature quite generally, is characterized by an unwarranted *Synchronic Presumption*, according to which the value question is conceived as a question about the comparative values of mere true belief and knowledge at a snapshot in time. At best, we are invited to compare a mere true belief that p and knowledge that p in a very short time frame. Accordingly, we tend to concentrate on what is of value in one or another moment's cognitive grasp of the directions to Larissa, rather than what epistemic transformations might occur as one trudges along the road, meeting other people along the way, passing or not passing various landmarks one had expected, and so on. Now, the Synchronic Presumption confines our philosophical attention to the present, and this has consequences for the lesson that Zagzebski draws from her critique of reliabilism—that the value of knowledge must be 'truth-independent'. The presumption causes a conflation of two quite different requirements of 'truth-independence', one narrow and one broad: (a) the narrow requirement that the value of knowledge be independent of the value that its constituent true belief already has in virtue of being true; and (b) the broad requirement that the value of knowledge be independent from the value of truth quite generally. It is a crucial motivation for credit accounts that the requirement of truth-independence is taken as the broad one in (b), for if the requirement were merely narrow as in (a), there would be far less motivation to cast the philosophical eye inward into the character of the agent in order to find some element, in itself supposedly independent of the value of truth, which contributes the distinctive value of knowledge. The Synchronic Presumption helps construct the value problem to suit certain styles of 'solution'.⁹

⁹ We see this phenomenon in more obvious form in the so-called tertiary value problem—the 'problem' of revealing the value of knowledge as *different in kind* from the value of truth (or from the value of whatever else may fall short of knowledge—see the definition given in Pritchard, this volume). The tertiary value problem really is a piece of philosophical artifice, corresponding to no natural philosophical intuition or question. It has surely come into being largely for the benefit of those with a 'solution' at the ready.

The conflation of (a) and (b) obscures the possibility that the value of knowledge (or one fundamental value of knowledge) *is* in fact reducible to the value of truth, even while remaining independent from the value of the truth of the constituent true belief. Thus it conceals the possibility that the value of knowledge consists in something about knowledge that helps us retain our true beliefs over time. Pursuing the way(s) in which knowing assists our general purchase on truth is, I think, the right way to approach the value question. The answer Socrates gives to Meno's question points us in this direction.

The Superior Resilience of Knowledge

The conclusion given in the *Meno* about the value problem is that knowledge is more valuable than true opinion because it is 'shackled'. Like the statues of Daedalus which were so life-like that people tied them down to stop them running away, states of knowledge are shackled so that we do not lose them. Now it is explicit in Plato's text that the metaphor of 'shackled' (or 'tethered') is intended by Socrates to signify awareness of reasons or evidence for the belief:

Socrates: If you own an original Daedalus, unshackled, it's not worth all that much—like a slave who keeps running away—because it doesn't stay put. But if you've got one that's shackled, it's very valuable. Because they're really lovely pieces of work. What am I getting at? My point is, it's the same with true opinions. True opinions, as long as they stay put, are a fine thing and do us a whole lot of good. Only, they tend not to stay put for very long. They're always scampering away from a person's soul. So they're not very valuable until you shackle them by figuring out what makes them true... And then, once they're shackled, they turn into knowledge, and become stable and fixed. So that's why knowledge is a more valuable thing than correct opinion, and that's how knowledge differs from a correct opinion: by a shackle'.¹⁰

¹⁰ Plato, *Protagoras* and *Meno*, 130.

If this is on the right track, and I think it is, then the value of knowledge will only reveal itself once we abandon the synchronic conception of the issue for a *diachronic* one. We have to conceive of epistemic subjects as placed in time in order to reveal the crucial difference: mere true beliefs are typically more vulnerable to being lost in the face of misleading counter-evidence. Reconsider the road to Larissa—and I shall try to be as epistemologically non-partisan as possible, assuming only what Socrates rightly assumes in his comment above, that knowledge typically involves arriving at one’s true belief on the basis of some suitable evidence or reasoning. The extra value in knowing the route as opposed to merely having a true opinion is that, over time, one is likely to come up against counter-evidence (you chat to a passer-by who says it’s the other way, you see a signpost that pranksters have turned to point the wrong way) and if you have some grasp of the evidence for your belief, as you typically will if you have knowledge, then you are in a better position to weight the new evidence. You are therefore less likely to abandon your true belief for a false one in the face of misleading evidence. The point is, possessing a true belief in the manner typical of knowledge shrinks the class of counter-evidence one will be misled by.

We want to possess truths because we need them to serve all our various purposes, but considered diachronically this entails that we value possessing them in a manner that is conducive to our retaining them over time in the face of misleading counter-evidence. Now one can instantly imagine various epistemically undesirable ways of doing this: sheer dogmatism will lead one to hold on to one’s beliefs, including true ones, in the face of any counter-evidence, including misleading counter-evidence. But, given our diachronic perspective, it is clear that dogmatism is epistemically undesirable because it is a thoroughly indiscriminate strategy over time. Sheer longevity is not what the value of knowledge consists in. Rather, I suggest we take Plato’s prompt that the value of knowledge resides in a tendency to survive the test of time in virtue of some kind of rational advantage. I suggest we coin a notion of ‘resilience’. Resilience is: *the tendency to survive misleading counter-evidence owing to the subject’s being in a position to weight it against positive evidence already possessed.*

Resilience names a typical feature of knowledge, not a necessary condition. It follows from the minimal assumption that knowers typically have a suitable grip on

reasons in favour of their belief, and that although this is not necessary for knowledge, it is none the less a central distinctive characteristic. Socrates explicitly has this characteristic in mind, for he says that mere true opinions are ‘not very valuable until you shackle them *by figuring out what makes them true...* And then, once they’re shackled, they turn into knowledge, and become stable and fixed’. Although Socrates’ comments are often interpreted as advocating a tripartite analysis of knowledge as justified true belief, there is no real commitment to any such thing taken as crucially distinct from most alternative modern analyses. It is historically more plausible to see modern analyses as variations, of a theoretically complex and highly inter-reactive sort, on the broad generalization rehearsed by Socrates, to the effect that knowing things typically involves believing them truly for a reason. This non-specific and minimal assumption about what typically distinguishes knowledge leaves my proposal non-partisan with respect to what might provide an adequate analysis of knowledge, and indeed to the question whether there could be any adequate analysis of knowledge.

Also taking his cue from the *Meno*, Tim Williamson has briefly made a similar suggestion in the context of his case for knowledge’s being a prime condition, and as such unanalysable. The context of his discussion is of course epistemologically partisan in its anti-analytical commitment, though I think something like his point can be made in the non-partisan spirit I am urging for responses to the value question. Williamson is not primarily engaged in establishing any particular answer to the value question as such, but rather in establishing the superiority of knowledge construed as a prime condition when it comes to predicting and explaining action. In the course of that framing argument, however, he draws the crucial conclusion about the value of knowledge, that ‘present knowledge is less vulnerable than mere present true belief to *rational* undermining by future evidence’, and that is indeed the point we should draw out of Socrates’s remarks.¹¹ However, his argument for this proceeds exclusively on the basis of two rather specific sorts of comparative case. Firstly, the case in which mere true belief is lost upon discovery that it was based on a false belief, whereas knowledge cannot be lost in that way because a true belief arrived at by way of a false

¹¹ Timothy Williamson, *Knowledge and Its Limits* (Oxford: Oxford University Press, 2000), 79.

lemma is not knowledge. And, secondly, the case in which mere true belief is lost upon discovery of misleading counter-evidence abundant in one's environment, whereas knowledge cannot be lost in that way because a true belief possessed in a context in which it might be defeated at any moment by counter-evidence is too unstable to constitute knowledge.

That Williamson relies so specifically on these two sorts of case to establish his conclusion, with which I agree, however renders his argument peculiarly vulnerable to the invocation of mirror image examples of circumstances in which knowledge is more readily lost than mere true belief. Jonathan Kvanvig has exploited this strategy to argue against Williamson, specifically constructing examples to mirror the case where there is misleading evidence in the environment. Kvanvig's examples illustrate how the advent of such evidence can cause one to lose precisely not true belief, but knowledge. For instance, he suggests that one's mathematical knowledge could be lost owing to a renowned mathematician's mistakenly asserting (in a suitably public domain, yet unbeknownst to one) something to undermine it. An example like this aims to remind us that our knowledge can come and go without our awareness, and so knowledge in general might come to seem just as elusive in the face of counter-evidence as is true belief. In similar vein, we could perhaps construct an example to mirror the false lemma case: A visitor to the U.K. forms the true belief that Gordon Brown is a powerful figure in British politics on the basis that Gordon Brown is Chancellor of the Exchequer. But then when (shortly after the visitor has left, and unbeknownst to her) Gordon Brown stops being Chancellor to become Prime Minister, her knowledge is lost—yet her true belief remains. Whatever one makes of such cases, I think the moral here is that if one makes the case for knowledge's greater rational persistence too much by way of specific cases, the argument risks descending into a competition over how many examples can be lined up on either side to influence our sense of what is rule and what is exception.

After giving his mirror examples to the misleading evidence scenario, Kvanvig concludes that Williamson's claim that knowledge displays superior rational persistence is at best contingently true. He regards this as an objection, evidently assuming that if knowledge does have a distinctive value, then that value will apply to all possible cases of knowledge without exception. He emphasizes his objection by

pointing out that one could be at a possible world in which most of our beliefs are fixed not by evidence at all, but rather pragmatically, in which case most of our true beliefs would not be knowledge but would be none the less robustly persistent in the face of misleading counter-evidence. (Of course, like dogmatically held beliefs, they might persist in the face of *any* counter-evidence, but Williamson would not be able to dismiss the pragmatic scenario by pointing to the fact that our beliefs would simply not be in ‘good order’, for they would be in good order pragmatically speaking.) Maybe so, but we should reject Kvanvig’s assumption that revealing the contingency of Williamson’s thesis amounts to an objection. On the contrary, Williamson’s anti-analytical epistemological position means he is explicitly committed to its being impossible to define knowledge, as distinct from mere true belief, in terms of knowledge’s greater rational persistence, and it follows that there will be exceptions to the rule that knowledge has greater rational persistence. For present purposes I aim to remain non-partisan on the question of the analysis of knowledge, but there is in any case simply no reason to expect a solution to the value problem to amount to an exceptionless claim about knowledge; it should be obvious to us from the start that it may simply be a generalization. (I shall return to this shortly.)

Accordingly, the resilience proposal wears its admission of exceptions on its sleeve. It is explicitly only a generalization to say that knowledge typically involves possession of evidence so that knowers are at an advantage when it comes to weighing in new counter-evidence. Clearly, we can allow that there are circumstances in which knowledge would lack resilience, or where its resilience would lack its usual value. For we can allow that there are exceptional circumstances in which knowledge is possessed without the usual evidential awareness, and exceptional circumstances in which retaining one’s true belief beyond the snapshot in time is simply of no interest. In such circumstances, we may readily admit, the knowledge in question would indeed lack the value it more normally possesses. Similarly, the resilience proposal also wears its more general contingency on its sleeve. It is manifestly a generalization confined to worlds significantly like this one. These worlds are, I take it, worlds in which most of our beliefs are fixed not pragmatically but more by evidence (even if some of them are partly *formed* as the result of pragmatic pressures, they none the less stand susceptible to evidential defeaters); and in which, most of the time, we achieve our true beliefs not by way of false lemmas and not in contexts where we are either

already surrounded by soon-to-be-observed existing misleading evidence, or soon-to-be-introduced-into-the-environment misleading evidence. These sound like exceptions to the rule and they surely are.

The cases on which Williamson so specifically bases his own argument, however, can help substantiate our more generally motivated resilience version of Socrates' point. They remind us that when we possess mere true beliefs, we may do so not only by complete fluke (as in my opening standard example of the true belief that there is a fire in the house—such happy flukes must almost never happen), but more often with some grip on the evidence, where that grip is inadequate—perhaps because it is flawed by a false lemma, or because it is rendered insufficiently stable by not-yet-observed misleading evidence. This supports the claim that it is typical of knowledge, as opposed to mere true belief, that one is in a better position to weigh in new counter-evidence, and so one is less likely to be misled. It is not that in most cases of mere true belief we have no grasp whatever of the evidential situation, but rather that, given we are falling short of knowledge, our grasp is bound to be inferior with respect to that crucial task of weighing new counter-evidence with existing evidence.

The value of resilience is of course reducible to the value of truth taken generally, for it is wholly derived from the value of sustaining true beliefs over time. But the resilience of knowledge is none the less a value over and above the value of the truth of any constituent true belief, which is the only value of truth made visible on the Synchronic Presumption at work in the value problem literature. The value of any item of knowledge is therefore not truth-independent in the broad sense given in (b); but it is independent from the pre-existing value of the truth of the constituent true belief, as is required in (a). By restricting our attention to the present, and so conflating (a) and (b), the Synchronic Presumption conceals the possibility that the value of knowledge consists, at least in part, in its superior resilience.

I said it should be obvious to us from the start that an account of the value of knowledge might take the form of a generalization about knowledge rather than a necessary condition. But we saw that Kvanvig presumes that Williamson must be aiming for a claim of necessity, and in this presumption he is in the good company of

most contributors to the value problem literature, not least because most are in the business of advocating their preferred analysis.¹² Why should it come as such a shock to entertain a thesis about the value of knowledge that is not a thesis about all possible cases of knowledge? The answer lies in a second presumption distorting the debate and limiting the responses we might make to the value problem: let us call it the *Analytical Presumption*. This is the presumption that the distinctive value of knowledge must be ready-contained in whatever warrant is said to convert mere true belief into knowledge. This methodological presumption stems from the analytical enterprise in epistemology, and whatever one may think about the wisdom of that enterprise vis-à-vis achieving an enlightening philosophical characterization of knowledge, it has certainly had a distorting effect on the present debate. The Analytical Presumption partly explains the attractiveness of credit accounts, for they locate the value of knowledge precisely in the allegedly ready-made form of the agent's epistemic credit that principally plays the warranting role. The Analytical Presumption is therefore a driving force behind the unfortunate cart-before-the-horse strategy of locating the value of knowledge in something whose own value can only be explained by reference to the prior value of ultimate epistemic ends such as knowledge. As I say, credit accounts can give satisfying solutions to the value problem as we have come to recognize it. But what we have come to recognize as the value problem has been substantially misshapen by the twin pressures of the Synchronic Presumption and the Analytical Presumption. The first presumption rules out the reducibility of the value of knowledge to the value of truth, and conceals the significance of the test of time; the second insists that the value of knowledge is to be found in some kind of warrant, and so conceals the possibility that a good response to the value question might be a generalization about what is distinctive of knowledge, and not a purported necessary condition.

A Different Diagnosis

¹² Needless to say there is absolutely nothing wrong in principle with advocating one's preferred position on any philosophical issue. My point is diagnostic, and only critical in so far as the partisanship has helped distort the value question.

Jason Baehr too has argued that the literature on the value of knowledge is, as it stands, on the wrong track.¹³ He observes that the literature is premised on the idea that there is a powerful and widespread pre-theoretical intuition to the effect that knowledge is more valuable than mere true belief, and he calls this the ‘guiding intuition’. He rightly observes that the guiding intuition is treated in the literature as placing a constraint on the analysis of knowledge, in the sense that any viable analysis must entail that knowledge is indeed more valuable than mere true belief. And he argues that the guiding intuition is not in good shape. Firstly, it would have to express an exceptionless generalization about knowledge; yet it is implausible that we have a real pre-theoretical intuition that expresses any such thing. Secondly, he argues it would have to be ‘formal’; that is, the guiding intuition would have to have no specific content beyond the fairly empty idea that knowledge is more valuable than mere true belief. It would have to be ‘formal’ in this sense, in order to make sense of the fact that there are so many contender vindications of it in the literature.

As my own arguments here make manifest, I am in complete agreement with the first point, that for the guiding intuition to provide a proper motivation for the literature as we know it, the intuition would have to be exceptionless, and that this is implausible. Baehr invokes the chief counter-examples of trivial knowledge and immoral knowledge. Recall Sosa’s example of counting the grains of sand. If we bother to do so, we may well achieve knowledge, but a piece of knowledge that clearly has no greater value than a trivial mere true belief with the same content. Such a triviality is not worth knowing any more than it is worth believing. Then, at the other end of the scale, there is immoral knowledge, such as knowledge of how to stir up ethnic hatred, or instigate genocide. Knowledge of these things has no greater value than the counterpart mere true beliefs, for they are so horrible that, again, such things are not worth knowing any more than they are worth believing. And so, argues Baehr (and contrary to Zagzebski, who, as we have seen claims admirability even for undesirable knowledge), knowledge is not *always* more valuable than true belief; and the guiding intuition is false. Thus one of the two chief motivations for the value

¹³ Jason Baehr, ‘Is There A Value Problem?’, in Adrian Haddock, Alan Millar, and Duncan Pritchard (eds.) *Epistemic Value* (Oxford: Oxford University Press, forthcoming).

problem literature as we know it—the guiding intuition *qua* strict universal—collapses. Agreed. Of course, one could always find ways of plausibly presenting our pre-theoretical intuition as in itself exceptionless. For instance, I think it is plausible to present it in a refined version that rules out the exceptions above, as expressing the idea: If there's value in believing it, then there's more value in knowing it. (An epistemic analogue to the proverb, 'If a job's worth doing, it's worth doing well.')

But this indeterminacy in how precisely to express our pre-theoretical intuition only serves to support Baehr's point, since our refined version would not provide a proper motivation for the literature as we know it. That literature aims to find the value of knowledge somewhere in a set of necessary and sufficient conditions for knowledge, and it follows that its guiding intuition must be an intuition about all cases of knowledge.

My point of disagreement comes only with the second argument: that the guiding intuition would have to be 'formal', that is, contentless beyond the mere claim that knowledge is more valuable than mere true belief. The thought here is that in order to motivate the array of philosophical appropriations of the guiding intuition, that intuition would have to say nothing substantive about why we value knowledge; yet it is implausible that we have any such empty pre-theoretical intuition. So the guiding intuition is found, on this score too, to be false. While I appreciate the argumentative aim here, I do not see that the guiding intuition would have to have no content beyond the sheer idea that knowledge is more valuable than mere true belief. It is commonplace for pre-theoretical intuitions of all sorts in philosophy to be in need of unpacking and/or clarifying. The puzzle of identifying what is implicit in a given intuition—and sometimes this can be the same thing as identifying its basic grounds—is just the sort of thing that different philosophical theories compete to solve. And I see no reason to regard the value of knowledge literature as straying from this standard model. While I can agree it is implausible that we have a pre-theoretical intuition which contains not even an inkling of why knowledge is more valuable than mere true belief, still it seems to me entirely plausible to say we have a pre-theoretical intuition whose inkling of substantive content is implicit, or confused, or in some other way ready for philosophical explicitation and development. In this respect, it seems to me, the guiding intuition stands.

While I agree, then, with Baehr's conclusion that there is something profoundly amiss in the value of knowledge literature, I prefer a different diagnosis. My diagnosis points to the significance of two presumptions: the Synchronic Presumption, which obscures the twin possibilities that the value of knowledge is reducible to the value of truth, and that it is to be revealed in knowledge's advantage in surviving a certain test of time; and the Analytical Presumption, which misleads us into thinking the value of knowledge must come ready-made in some kind of warrant, and, therefore, into thinking that if there is a positive answer to be found to the value question, then it must apply to all possible cases of knowledge. Both presumptions serve to obscure the fundamental value of knowledge, which, taking my cue from Plato's Socrates, I have been arguing for in terms of knowledge's superior resilience. More broadly, we might say that both presumptions obscure the significance of the test of time, but in slightly different ways. The Synchronic Presumption obscures it simply by confining our attention to the present snapshot in time; and the Analytical Presumption obscures it by directing the philosophical gaze to the retrospective matter of the aetiology of the true belief (was it formed by way of a reliable faculty/agent/good epistemic motive?) and so away from the prospective matter of how well it will survive misleading counter-evidence as time goes by.

No doubt resilience is not the only basic value of knowledge. Another presumption one might explore the extent of in this debate is that of individualism. If we expand our conception of the value question not only through time to embrace the diachronic perspective I have been urging but also out across social space, we may find that the only true beliefs that we may responsibly pass on to others by testimony constitute knowledge, and, correlatively, that the only true beliefs we should accept from others constitute knowledge. If so, the two-way sharability of truths is another basic value of knowledge.¹⁴ A key point I hope to have put across is that approaching

¹⁴ Edward Craig's *Knowledge and The State of Nature* can be read as an extended explanation of the value of knowledge in terms of the sharability of truths, though his argument is not geared explicitly to the value question, but rather to a practical explication of why we come, of necessity, to have the concept of knowledge at all. But one instantly sees how such an explication of why we have the concept might simultaneously constitute an explanation of its value. (See *Knowledge and the State of Nature: An Essay in Conceptual Synthesis* (Oxford: Clarendon Press, 1990). Martin Kusch has discussed Craig's genealogy as providing a social explanation of knowledge's value in 'Testimony and the Value of Knowledge' forthcoming in

the question of knowledge's value from an epistemologically partisan point of view has tended to carry unwarranted presumptions into how the issue is viewed, presumptions that actively obscure the value of resilience. While all approaches have their point, and are likely to reveal some layer of knowledge's value—most accounts are after all addressing, in some form or other, the evidential sensitivity that marks out knowers—there is clearly a role for different, and less partisan approaches to this question, so that we may achieve a fuller picture of the different, often inter-related, values of knowledge.¹⁵

Miranda Fricker
Birkbeck College

Haddock, Millar and Pritchard (eds.). Ward Jones's early paper on the value question also places Craig's work and the issue of testimony centre-stage; see 'Why Do We Value Knowledge?', *American Philosophical Quarterly*, 34: 4 (1997); 423-439.

¹⁵ I presented earlier versions of this talk to a Workshop on the Value of Knowledge at the University of Copenhagen, organized by Klemens Kappel; and at research seminars at the University of Glasgow and University of Bristol. I thank all those who took part for their questions and comments. I am also grateful to Alan Millar for a conversation that first prompted me to latch on to the idea that knowledge's value has something to do with the handling of counter-evidence; and to Jason Baehr for discussion and subsequent email exchanges that helped clarify the differences between our views.