

Epistemology

Draft 1

For Rossberg, M., *Cambridge Handbook of Analytic Philosophy*

Mona Simion

COGITO Epistemology Research Centre, University Of Glasgow

On the face of it, we care a lot about knowledge. As a society, we invest a lot of time and energy in the development of institutions whose aim it is to accumulate or distribute knowledge: universities, schools, libraries, and the internet are among the most prominent of these. On an individual level, we send our children to school so that they can acquire knowledge about a wide range of topics. Some of us go to considerable financial lengths in order to make this possible. It is hardly surprising, then, that the study of knowledge has historically received a great deal of attention in philosophy. A lot of effort has been made to get clear on what exactly is involved in knowing and having justified beliefs, and how we come to know from our senses and the word of others. Finally, the issue of whether we should care that much about knowledge to begin with has also been subject to thorough scrutiny.

This article briefly surveys recent work in analytic epistemology, or the theory of knowledge. The discussion is organised around the main contemporary debates in the field on issues concerning the value, nature, and extent of knowledge and epistemic justification.

1. The Nature of Knowledge
 - 1.1 Gettier
 - 1.2 Modal Accounts
 - 1.3 Virtue-Theoretic Accounts
2. The Value of Knowledge
 - 2.1 Value Scepticism
 - 2.2 The Final Value of Knowledge
 - 2.3 Knowledge and Stubborn Belief
3. Scepticism and the Extent of Knowledge
 - 3.1 Rejecting Closure
 - 3.2 Mooreanism
 - 3.3 Contextualism
4. The Nature of Epistemic Justification
 - 4.1 Internalism vs. Externalism
 - 4.2 Formal Epistemology
 - 4.3 Knowledge-First Epistemology
5. Social Epistemology
 - 5.1 Testimony

5.2 Disagreement

5.3 Group Belief

1. The Nature of Knowledge

What is knowledge? Most epistemologists would agree that knowing that something is the case implies that one believes it to be the case, their belief is true, and justified – i.e., held for a good reason, or based on solid evidence, or on reliable methods of inquiry. Most would also agree, however, that merely having a justified true belief (JTB) is not sufficient for knowledge. What more is needed? For more than half a century, a lot of philosophical ink has been spilled on trying to answer this question.

1.1 Gettier

In a famous paper (1963), Edmund Gettier puts forth a couple of counterexamples to the analysis of knowledge as JTB. Consider, for instance, the case of Rod: Rod drives past a field, looks out the window, and sees what looks exactly like a sheep. Rod thereby justifiably comes to believe that there is a sheep in the field: after all, perception is a paradigmatic source of epistemic justification. As a matter of fact, however, what Rod is looking at is not a sheep, but a sheep-looking dog. However, luckily, there is, indeed a sheep in the field, behind the hill, outwith Rod's visual field. Luckily, then, Rod's belief is also true. Rod has a justified true belief that there is a sheep in the field but, intuitively, he doesn't know that there is a sheep in the field: after all, he's looking at a dog. The JTB analysis of knowledge fails down the sufficiency direction.

One early proposal (Goldman 1967) to address this problem supplements the JTB analysis with a causal condition on knowledge (henceforth JTB+C). According to this account one knows that p iff one justifiably and truly believes that p and one's belief that p is appropriately caused by the fact that p . It is easy to see that this view gets the case of Rod right: after all, Rod's belief is not caused by the fact that there is a sheep in the field – since the sheep is obscured from view – but rather by the fact that there is a sheep-looking dog in the field. In this, Rob's belief fails to qualify as knowledge by the lights of the causal analysis.

Unfortunately, it turns out that JTB+C is not strong enough to account for a different variety of Gettier cases that have become known in the literature as Fake Barns-style cases (Goldman 1976). Suppose there is a county with the following feature: the landscape is peppered with barn-facades that from the road look exactly like barns. Suppose further that Henry is driving through Fake Barn County, looks out the window, sees a structure that looks exactly like a barn, and forms the (justified) belief that there's a barn in front of him. Now suppose further that Henry happens to be looking at the one and only real barn in the county – thereby, his belief is also true. Should he have looked out the window just a few seconds earlier or later, though, Henry would have formed a false belief. Intuitively, Henry's belief is luckily true and thus fails to qualify as knowledge. Note, though, that JTB+C

misses this prediction: condition C is met in this case. Henry's belief is directly caused by the fact that there's a barn in front of him.

1.2 Modal Accounts

A large number of analyses of knowledge have been proposed in the literature since 1963, in an attempt to escape Gettier-type cases. In what follows, I will not survey this intricate history but rather focus on most recent and widely endorsed such theories of knowledge.

One thing several philosophers have noticed is that what seems to constitute the knowledge-precluding problem in Gettier-type cases has to do with a particular variety of epistemic luck (Pritchard 2005), that intervenes between the belief's being true and the reason for which it is held. Modal accounts of knowledge aim to address this issue: according to these accounts, one's belief is lucky in the relevant, knowledge-precluding sense if and only if one could have easily believed falsely rather than truly.

One way to unpack this thought was defended by Robert Nozick (1981) and is known as the sensitivity condition on knowledge: a subject S's belief is sensitive iff in the closest possible world (possible situation) in which p is false, S would not continue to believe that p. It's easy to see that the sensitivity condition will help with Gettier cases like our toy case of Rod above: had there been no sheep behind the hill, Rod would have continued to believe that there is a sheep in the field. In this, Rod's belief fails to meet the sensitivity condition, which, in turn, explains why he lacks knowledge.

At first glance, sensitivity also predicts that Henry doesn't know in Fake Barns: after all, had there been a mere barn façade in front of him, he would have still formed the belief that there's a barn.

Unfortunately, we can easily modify Fake Barn cases as to create trouble for sensitivity: suppose that barn facades are always green, but genuine barns are always red (Kripke 2011): Henry's belief that he sees a *red* barn will be sensitive, even though his belief that he sees a *barn* will not. Since it is highly counterintuitive that Henry can know that there is a red barn in front of him, while failing to know that there is a barn in front of him, sensitivity is back in trouble.

The sensitivity account is also too strong: if knowledge implies sensitive belief, it would seem, most of our garden-variety induction-based beliefs will fail to qualify as knowledge (Sosa xxx). Consider: Sam drops a trash bag down the garbage chute of his apartment building, and forms the belief that the bag will fall to the basement garbage room. His grounds for so believing are inductive: it is possible that the bag will be snagged in the chute, but extremely unlikely. As it happens, the bag is not snagged in the chute and his belief is true. Intuitively, Sam knows that the bag is now in the basement; however, his belief is not sensitive: in the closest world where the bag gets snagged in the chute, Sam continues to believe it's in the basement.

Note, also, that denying knowledge to Sam would amount to generalised scepticism about inductive knowledge – which is a large amount of the knowledge we take ourselves to have. Our justification to believe that the sun will rise tomorrow is inductive. We don't want

out account of knowledge to predict that we don't know that the sun will rise tomorrow.

An alternative way to unpack the modal condition on knowledge that avoids this counterintuitive result is the safety condition, put forth by Ernest Sosa (1999): on this view, S knows that p iff, in all nearby worlds (i.e. situations similar to the actual one) where S believes that p , p is not false. Note that safety correctly predicts that Sam knows that the garbage bag is in the basement, and that we all know that the sun will rise tomorrow, and other such everyday matters of fact: after all, the worlds at which these things fail to occur are clearly fairly distant from our own.

Safety also correctly predicts that Rod does not know that there is a sheep in the field – since the world at which he falsely believes it is nearby, i.e. a world at which all that needs to change is that there is no sheep behind the hill – and that Henry does not know that there is a barn in front of him, since he could have easily been looking at a façade and falsely believed it was a barn. All in all, safety seems like a solid response to the Gettier problem.

However, several cases have been put forth in the literature to shed doubt on the necessity of safety for knowledge. Here is, for one, a case from Chris Kelp (2009): Russ wakes up in the morning, comes to down the stairs, has a look at the clock, sees that it reads 8.22 and on that basis forms a belief that it's 8.22. Russ's belief is justified and true. But now suppose Russ's arch-nemesis has an interest that Russ forms a belief (no matter whether true or not) that it's 8.22 by looking at the clock when he comes down the stairs. Russ's arch-nemesis is prepared to do whatever it may take in order to ensure that Russ acquires a belief that it's 8.22 by looking at the clock when he comes down the stairs. However, Russ's arch-nemesis is also lazy. He will act only if Russ does not come down the stairs at 8.22 of his own accord. Since Russell does come down the stairs at 8.22, Russell's arch-nemesis remains inactive.

Intuitively, Russ knows: he forms a belief that it's 8.22 based on reading a highly reliable clock, and it is, in fact, 8.22. However, Russell's belief that it's 8.22 is not safe: At all nearby possible worlds at which he comes down a minute earlier or later his arch-nemesis steps on the scene and sets the clock to 8.22.

1.3 Virtue-Theoretic Accounts

According to virtue-theoretic views, knowledge is a cognitive achievement: it is true belief generated via epistemic competence. On Sosa's AAA view (2007), for instance, S knows iff S's belief is accurate, adroit (sourced in competence) and apt (accurate because adroit).

It is easy to see that Rod's belief that there is a sheep in the field does not satisfy the aptness condition: it is true, and sourced in its perceptual competence, but it is not true *because* sourced in its perceptual competence, but rather because of luck.

The problem that these accounts encounter is one pertaining to the strength of the competence condition. How salient should competence be in the explanation of the fact that a true belief was formed? One possible answer is the strong answer: quite salient. This way to go explains, for instance, why Henry doesn't know in Fake Barns: Henry's perceptual competence is involved in his forming the

belief that there's a barn in front of him, but it does not carry the highest explanatory salience with regard to Henry's belief being true. Rather, luckily having looked out the window as the precise moment when the only real barn was in view does.

The problem with the strong version of the view, however, is that it is thought by many to get cases of testimonial knowledge wrong (Lackey 2007). After all, it would seem as though when a subject learns from the testimony of others, what saliently explains the fact that they acquire a true belief is the epistemic competence of the testifier, rather than S's own. If so, a strong competence condition will predict we cannot get knowledge from testimony. This is highly problematic: in virtue of our physical and psychological limitations, most of the things we know we believe based on other people's - teachers, parents, colleagues etc - say-so.

Alternatively, one can go with a weaker competence condition for knowledge: after all, it is plausible that some degree of epistemic competence is, as a matter of fact, involved in receiving testimony: you wouldn't just believe anybody, or anything they say. As such, maybe the right account is one that asks for the believer's cognitive competence to be involved in the formation of the true belief at stake, but not necessarily in the most salient manner.

The problem with this way to go, however, is that it fails to accommodate Fake Barns cases: after all, Henry's perceptual competence is clearly involved in his coming to believe that there is a barn in front of him. In this, virtue-theoretic accounts are facing a strength dilemma (Pritchard 2009) for the competence condition: they either misdiagnose Fake Barn cases (the weak version) or testimonial knowledge cases (the strong version).

For more work on the analysis of knowledge, see Ichikawa and Steup 2018).

2. The Value of Knowledge

Why is knowledge valuable? In particular, what is it about knowledge – rather than lesser epistemic states, such as, for instance, luckily true belief – that warrants all the special attention it has historically received both within and outwith philosophy?

Here is a naïve answer to start with: knowledge is instrumentally practically valuable. It's valuable because it is conducive to practical goods, such as resources, and power.

Plausible as it may seem, this answer leads straight into a philosophical puzzle. To see how, let's rehearse what we've just seen in a more precise fashion. We've said that, intuitively, knowledge is more valuable than merely luckily true belief:

- (1) $K > TB$: The value of knowledge is higher than the value of mere true belief.

$K > TB$ is a very plausible claim. To see this compare, a doctor who diagnoses you as having a cold based on thorough medical investigations, with one who merely tosses a coin and happens to get it right; which doctor would you want to hear from?

Second, we said that the following also holds:

- (2) VKP: The value of knowledge amounts to its practical instrumental value.

VKP is just the naïve assumption that we started with.

Now, here is the problem: it looks as though luckily true belief is just as valuable for practical purposes as knowledge is. Consider the following cases: in the first one, I am thirsty, I know that there is milk in the fridge, and based on this knowledge I open the fridge, find the milk, and quench my thirst. In the second case, I have no clue as to whether there's milk in the fridge; I toss a coin, however, and decide to believe that there is milk in the fridge should the coin land heads. The coin does land heads, I form the relevant belief, open the fridge and quench my thirst. In sum, it looks as though the following holds:

- (3) PVK=PVTB: Mere true belief is just as practically instrumentally valuable as knowledge.

(1), (2), and (3) generate a puzzle because they are inconsistent: one of these claims needs to go. The puzzle has become known in the literature as 'The Meno Puzzle' – after Plato's dialogue 'The Meno' where the problem is first discussed.

Unsurprisingly, solutions to the puzzle have focused on denying one of claims (1)-(3). In what follows, I will look at these solutions in turn.

2.1 Value Scepticism

Sceptics about the value of knowledge (Pritchard 2010, Kvanvig 2003) deny (1). According to these philosophers, our historical concern with knowledge is unwarranted. Rather, we should be focusing on more complex epistemic states, such as, for instance, understanding. Pritchard's argument runs as follows: he notes that, in order to account for our historical concern with knowledge, we don't merely need to explain why it is somewhat more valuable than merely luckily true belief. What we need is an account that predicts that knowledge is, in some way or another, distinctively valuable, in that it is not on a value continuum with mere true belief. In turn, Pritchard takes this requirement to imply that what we need is an account that predicts that knowledge has a different *kind* of value than mere true belief. Since, however, it is implausible that this should be so, Pritchard takes this to warrant scepticism about the distinctive value of knowledge.

The problem with this take on the issue is that, on closer inspection, it remains unmotivated: while Pritchard may well be right that our special concern with knowledge demands an explanation that predicts that knowledge and mere true belief are not on the same value continuum, this need not imply that their value is of different kinds. We know from Mill's (1963) value theory that values can display superiority relations. Weak superiority relations obtain when some amount of good x is better than any amount of good y. Parfit (1984) gives the example of one hundred years of a really good life being superior to a drab eternity, lived on bad music and potatoes. Strong

superiority relations obtain when any amount of good *x* is better than any amount of good *y*: any amount of love from one's child, for instance, is strongly superior to any quantity of good coffee. It is open, then, to the defender of the distinctive value of knowledge to argue that knowledge and mere true belief instantiate the same kind of value, but nevertheless do not find themselves on a value continuum in virtue of (weak or strong) superiority relations. Indeed, Kelp & Simion (2016) make this point: according to this account, knowledge is weakly superior in value to mere true belief in virtue of the fact that some amount of knowledge is conducive to a flourishing life, in a way in which no amount of merely luckily true belief is.

2.2 The Final Value of Knowledge

Another route towards escaping the puzzle is via denying (2) – the claim that the value of knowledge consists in its practical instrumental value. Virtue epistemologists take this route: according to John Greco (2010), for instance, knowledge is finally rather than instrumentally valuable. Final value is value for its own sake. Greco thinks knowledge is finally valuable because it's a species of achievement – i.e., cognitive achievement – and because he takes it that achievements are valuable in their own sake, independently of what there are conducive to.

The main problem with this account – even if we accept, with Greco, that knowledge is a species of achievement – is that it's not clear that all achievements display final value – i.e., value over and above the value of the outcome achieved. First, notice that 'bad' achievements don't seem to be particularly valuable; indeed, they seem, if anything, worse than the corresponding lucky successes. Take, for instance, a murderer who, through his excellent shooting skills, manages to hit the victim from the very first shot. Compare, now, with someone who shoots someone in a hunting accident. Clearly, the former state of affairs is worse than the latter. This suggests that 'bad' achievements bring negative value to states of affairs.

Alternatively, Greco might want to restrict his account of achievements to positive achievements only; after all, this is the proper genus of knowledge. Even so, though, the problem is that it's not clear, even for positive achievements, that they are always more valuable than the corresponding lucky successes. If they are not, it is implausible that they instantiate final value on top of the instrumental value towards the end achieved. To see this, think about the value of natural beauty, or untouched forests: we value these things much more than their humanly-fabricated counterparts, although the former came about by mere luck, while the latter are achievements through serious effort.

2.3 Knowledge and Stubborn Belief

Finally, some philosophers have denied claim (3) – that knowledge and true belief are equally practically instrumentally valuable. Plato himself chose this route in *The Meno*, and Tim Williamson (2000) is a notable contemporary champion of the view. This account, knowledge is practically more valuable than mere true belief because it is more stable, less easily lost, and thus more likely to guide action

to the achievement of one's practical goals. Plato gives the example of someone walking from Athens to Larissa whilst knowing the way there, in comparison to someone who merely luckily has a true belief about the correct route. According to Plato, the latter fellow is likely to lose his belief at the first turn that seems to go the wrong way. In contrast, the knower will continue ahead, since he knows that's the right way to Larissa.

The problem this account faces, however, comes from stubborn, dogmatic beliefs: these can often be as resilient as knowledge, if not even more so. One thing we don't want our account of the value of knowledge to predict, however, is that knowledge is just as valuable as dogmatic belief.

For more work on the value of knowledge, see (Pritchard, Turri and Carter 2018).

3. Scepticism and the Extent of Knowledge

Consider the following three claims:

- (1) I don't know that I am not a brain in a vat connected to a computer that is feeding me all my experiences (BIV).
- (2) If I don't know that I am not a BIV, then I don't know that I have hands.
- (3) I know that I have hands.

Claims (1), (2) and (3) are individually highly plausible but jointly inconsistent, and thereby give rise to a paradox: something has to go. This paradox has become known in the literature as the sceptical paradox.

Support for the plausibility of (1) comes from the subjective indistinguishability of sceptical scenarios from everyday situations: by stipulation, the BIV has no way to tell that their experiences are not mapping on to facts in the world anymore after envattment. From where they stand, they wake up in the morning just like they always have, go to work etc. How could one then know that one is not a brain in vat?

Support for (2) comes from the closure principle for knowledge: according to this principle, if one knows that p, and one knows that p entails q, then one can thereby come to know that q. The closure principle is not only extremely attractive, but also underlies our very capacity to extend our knowledge. To go back to our sceptical argument, note that knowing, for instance, that I have hands, and knowing that the fact that I have hands implies that I am not a brain in a vat, puts me in a position to come to know that I am not a brain in a vat. Since, however, by (1) I don't know that I am not a brain in a vat, it seems to follow that I don't know that I have hands, or any ordinary thing I take myself to know. But that contradicts (3).

It's important to note that the paradox that is problematic even if there are no sceptics out there, since it shows that there is an inconsistency in our thinking about knowledge. As such, solving the sceptical paradox is epistemologically important independently of the popularity (or lack thereof) of scepticism.

3.1 Rejecting Closure

One way to escape the sceptical paradox is to reject the closure principle for knowledge, i.e. to hold that the principle, while very plausible at first glance, does not hold in full generality. Defenders of sensitivity accounts of knowledge traditionally take this route against sceptical arguments (e.g. Dretske 1970). Suppose you know that *p* and that *p* entails *q*. Suppose also that you believe that *q*. If you know that *p*, your belief that *p* is sensitive: were *p* false, you wouldn't believe *p*. Suppose your belief that *q* is sensitive as well. Now, note, however, that the closest worlds at which *p* is false may differ from the closest worlds at which *q* is false. The difference may be that the closest worlds at which *p* is false are normal and you don't believe that *p*, while the closest worlds at which *q* is false are worlds of elaborate deception and so you still believe *q*. As such, sensitivity implies that the closure principle fails. My belief that I have hands is sensitive, since the closest world at which I have no hands is one where I can tell the difference (say, they were cut in an accident). In contrast, my belief that I am not a BIV is not sensitive, since at the closest world at which it is false I am a BIV, and thereby cannot tell the difference. Closure fails.

There are two main problems for this solution to the sceptical paradox: first, as we have already seen, there is very good reason to believe the sensitivity account of knowledge is incorrect. Second, sensitivity overgenerates closure failure. Consider a case in which you know that there is a great dane in the garden and infer that there is a dog in the garden. Intuitively, you can come to know that there is a dog in the garden based on this inference. But now say that behind the great dane there is a wolf; had there not been a dog in the garden, you would have seen the wolf and believed that it's a dog. According to sensitivity, and against intuition, your belief that there is a dog is not sensitive, whilst your belief that there is a great dane is sensitive.

3.2 Mooreanism

G.E. Moore (1959) and his contemporary followers notably attempt to escape the sceptical paradox by rejecting (1) – i.e., the claim that we don't know that we're not brains in vats.

Here is Moore's famous proof of the existence of the external world:

- M1. I have a hand.
- M2. If I have a hand, then I am not a handless BIV
- M3' I am not a handless BIV.

On the face of it, I know M1 and M2, and the argument is non-circular and valid: by providing the argument I prove the conclusion. So why is it that we all get the feeling that something suspicious is going on in Moore's argument?

One reason that might be is because, one might think, the sceptic has provided us with independent argument against M1 – i.e., the sceptical argument. If so, assuming M1 – i.e., assuming a claim that

the sceptic has already argued to be false – is begging the question against the sceptic.

Note, however, that the answer cannot be as easy as that: after all, while the sceptic's premises may well be highly intuitive, so is Moore's M1. If so, intuitive plausibility is not enough for the sceptic to get an upper hand over Moore. More is needed here to settle the debate. See e.g. (Wright 1985) and (Pryor 2000) for more on this topic.

3.3 Contextualism

We have seen that the Moorean strategy for escaping the sceptical paradox was to deny claim (1) – the claim that we don't know that we are not BIVs. Closure sceptics, in contrast, denied (2) – i.e. the entailment from our lack of knowledge of the denial of the sceptical hypothesis to lack of knowledge of ordinary facts. Sceptics, of course, deny (3) – i.e. our claim to know all the things we take ourselves to know.

Interestingly, these strategies do not exhaust all options explored on the market for escaping the puzzle: epistemic contextualism denies the very incompatibility of (1)-(3).

According to contextualism (e.g. DeRose 2009), the meaning of 'knows' varies with contextual features, such as practical stakes and tabled error possibilities: we are happy, for instance, to say that I know I will teach epistemology next year in normal circumstances, but reluctant to do so if, for instance, someone's life is under threat if we're wrong. Contextualists think that this is because the threshold for being in a strong enough epistemic position for knowledge is higher in some contexts – like high stakes contexts, for instance – than in others.

Similarly, according to contextualists, 'knows' means something different in (1) and (3). 'I know I have hands' is true in ordinary contexts of everyday life, since the threshold for being in a strong enough epistemic position for knowledge is meetably low, and we easily surpass it. At the same time, and compatibly, 'I don't know I am not a BIV' is true in contexts operative in the epistemology classroom during discussions of scepticism: the threshold is driven way up by the context, and it is impossible for us to surpass it.

As such, according to contextualism, the sceptical paradox is an illusion: anti-sceptics & sceptics both get it right, but in different contexts. In ordinary contexts 'I know I have hands' is true and so is 'I know I am not a BIV.' In sceptical contexts 'I don't know I am not a BIV' is true as is 'I don't know I have hands.' If so, there is no sceptical threat for our ordinary claims to knowledge.

There are two main problems with this solution to the sceptical paradox: first, many believe contextualism about 'knows' remains unmotivated by the linguistic data brought in its support (Simion 2021, Rysiew 2021).

Second, the viability of the contextualist solution to the sceptical paradox itself has been called into question. Several people (e.g. have worried that a semantic thesis about 'knows' is irrelevant to the epistemological question raised by the sceptic: after all, whether it is appropriate to attribute knowledge or not can be thought to be irrelevant to whether knowledge is instantiated or not (Sosa 2004). Anor excellent argument against the contextualist solution to

scepticism, coming from Baron Reed (2010), shows how the sceptic can make an easy comeback, by simply replacing 'knowledge' with a property that, even by contextualist lights, is context invariant, such as 'strength of epistemic position' (SEP) throughout. The sceptical paradox then becomes: (1) My SEP with respect to the proposition that p ('I have hands') entails non-BIV (I am not a brain in a vat) is excellent. (2) If my SEP with respect to the proposition that p entails Non-BIV is excellent, then my SEP for p cannot be much higher than my SEP for Non-BIV. (3) My SEP with respect to Non-BIV is poor. See (Rysiew 2021) for more discussions of contextualism.

4. The Nature of Epistemic Justification

After more than half a century of attempting to analyse the nature, value, and extent of our knowledge, it is fair to say that most recent literature features a justification turn – in that the focus moved towards trying to offer a satisfactory account of the nature of epistemic justification.

4.1 Internalism vs. Externalism

What does it mean for a belief to be justified? According to epistemic internalism, justification supervenes on features internal to the cogniser's mind (Conee and Feldman 2001); according to a popular version of the view, for instance, one is justified to believe that p iff it seems to one that p is the case. If it seems to me like there's a table in front of me, I am justified to believe that there's a table in front of me. To see the motivation for the view, think again about the Brian in a Vat scenario. To the envatted you it seems as though they are waking up in the morning, having a cup of coffee, going to work etc. Would the envatted you be justified in forming the corresponding beliefs? Intuitively, the answer is 'yes': they would be just as justified as the normal, full-bodied you: after all, they can't tell the difference.

The main problem for this view comes from cases of badly formed beliefs – or beliefs with dubious etiologies. Consider, for instance, cases of wishful thinking: say that, in virtue of wishing really hard that your partner loved you, it seems to you as though they're acting in a very loving manner no matter what they do. Intuitively, you're not justified in your belief. Similarly, think of cases of implicit bias: take a teacher who is firmly convinced that girls in his class are really bad at maths, based on sexist bias. Again, intuitively, no epistemic justification is present in these cases, although the relevant seemings are.

Cases like these have driven the vast majority of contemporary epistemologists to embrace epistemic externalism: on these views, justification does not merely depend on features internal to the mind; rather external features – pertaining, for instance, to the reliability of the process (Goldman 1979) or ability involved in belief formation (Sosa 2007) – call the shots. Reliability is a worldly feature: whether my perception gets it right most of the time or not is not a feature internal to my mind.

It is easy to see that externalism will, at first glance, score well when it comes to explaining away cases of wishful thinking and biased believing, since wishes and biases are not reliable ways to form beliefs.

There are two difficulties traditionally facing classical reliabilist externalist accounts of justification: the first concerns the reliabilist's sufficiency claim, the second sheds doubts on the necessity direction.

The worry for the sufficiency of reliable formation for epistemic justification is sourced in the phenomenon of accidental reliability (Bonjour and Sosa 2003). Consider a case in which whenever you wear something red in Berlin you come to believe, based on this, that it will be raining in Spain. That's a strange way to form beliefs: you have, intuitively, no justification whatsoever to believe it's raining in Spain when you form your beliefs in this fashion. Now consider, however, a strange situation in which, completely unbeknownst to you, your wearing red in Berlin is, as a matter of fact, strongly correlated with rain in Spain. Does this change the justificatory status of your belief? Intuitively, the answer is 'no'. According to reliabilism, however, you will be justified to believe it is raining in Spain insofar as the correlation obtains, independently of whether you are aware of it or not.

The main difficulty for the necessity claim lies with accommodating the intuition of justification in brain-in-a-vat-type cases (Cohen 1984): after all, the vast majority of beliefs formed by the envatted you are false, and thereby their belief-formation processes are highly unreliable. (More subtle versions of externalism deal better with this case: according to normal worlds reliabilism, for instance, justification requires for the belief formation process at stake to be reliable in normal worlds – i.e. worlds that are similar enough to the actual world. The envatted you will qualify as justified on this view, since the envatment world is clearly a distant one.) See (Pappas 2017) for more work on the internalism/externalism debate.

4.2 Formal Epistemology

The classic debate between internalism and externalism concerns the justification of full, outright beliefs – the kind of beliefs that are involved in knowing. Sometimes, however, we don't have enough justification to form an outright belief, although we do have some evidence in support of the proposition at stake. In these cases, many think we are justified to form a (lower or higher) degree of belief, or a credence. The question that arises, then, in parallel to the question of what justifies full beliefs, is the question of credence justification. For the most part, the issue is addressed in formal epistemology, i.e. the branch of epistemology that uses formal tools, from logic and mathematics, to put forth norms about rationality and justification.

Parallel to internalism and externalism about full belief justification, there are two main, competing views in the epistemology of degrees of belief: Subjective and Objective Bayesianism.

According to Subjective Bayesianism, one's credence that p is justified iff it is correct given the prior probability one assigns to p obtaining, updated in accordance with probability axioms. Subjective Bayesianism does not impose any restrictions on the way in which we assign the probabilities we start with; in this, Subjective Bayesians are the formal cousins of epistemic internalists about full belief justification. In contrast, Objective Bayesian accounts constrain our

prior probability assignments: on this view, there exists a correct probability assignment given a certain body of information available to the cognizer, and justified credences conform to these correct probability assignments. See (Weisberg 2021) for more on this topic.

4.3 Knowledge-First Epistemology

In light of more than half a decade of unsuccessful attempts to analyse knowledge in terms of epistemic justification (plus belief, truth, and some anti-Gettier condition), knowledge-first epistemology (first proposed and developed in Williamson 2000) proposes a radical methodological u-turn. These views commonly take knowledge to be a mental state in its own right, and employ it as a primitive in analysing justification.

On a simple knowledge-first account, of the kind proposed in Williamson 2000, one justifiably believes that p iff one knows that p (henceforth $JB=K$). The account escapes the Gettier problem altogether, since it denies that JTB is present in these cases to begin with: since Gettierized victims do not know, they do not justifiably believe either. The view also escapes the accidental reliability worry for traditional reliabilist externalisms: since one does not know based on accidentally reliable processes, one cannot be justified in holding the corresponding beliefs either.

The main problem for the account is constituted by brain-in-a-vat-type cases, where knowledge is absent, and thereby $JB=K$ predicts justification is absent as well (Brown 2018). Champions of this view commonly appeal to error theory to accommodate these cases: according to Williamson (2020) and Littlejohn (2020), our intuition of epistemic permissibility in these cases tracks epistemic blamelessness rather than epistemic justification. Justification, the thought goes, obtains when we comply with the epistemic norm of belief. Blamelessness, in contrast, is instantiated when we are in breach of the norm for no fault of our own. Since we are not good at distinguishing between intuitions of justification proper and intuitions of mere blamelessness, and since we think the brain-in-a-vat is doing nothing wrong, we tend to mistakenly attribute justification to them.

Other knowledge-first epistemologists are not satisfied by this error-theoretic explanation, and propose less demanding accounts of justification. There are two main versions of this view on the market: evaluative knowledge-first accounts (e.g. Bird 2007, Ichikawa 2014) claim that justified belief is would-be-knowledge: it displays all the epistemic evaluative properties a knowledgeable belief would.

Prescriptive knowledge-first views (Miracchi 2015, Kelp 2016, Simion 2019) are knowledge-first varieties of reliabilism, in that they focus their analysis on the properties of the belief-formation process, rather than those of the belief itself. On knowledge-first virtue-theoretic views, for instance, one is justified iff one forms their belief via an ability to know. Similarly, on proper functionalist accounts, a belief is justified just in case it is formed via a process that has the function of generating knowledge.

In contrast to the simple knowledge-first account of justification, these more complex views manage to predict that the brain in a vat is justified, since their beliefs have would-be-knowledge

properties/are formed in a knowledge-generating manner. At the same time, these views also rip the general benefits of the knowledge-first framework, in that they don't encounter Gettier problems, or accidental reliability problems, in virtue of their treating knowledge as primitive in analysing justification. For more discussion on knowledge-first views of justification, see (Silva 2020).

5. Social Epistemology

We are highly social creatures, dependent on each other for flourishing in all walks of life. Our epistemic endeavours make no exception: due to our physical, geographical, and psychological limitations, most of the knowledge we have is sourced in social interactions. We must inescapably rely on the intellectual labour of others, from those we know and trust well, to those whose epistemic credentials we take for granted online. Given the staggering extent of our epistemic dependence—one that recent technologies have only served to amplify—social epistemology is amongst the most thriving research areas in contemporary philosophy: if we are to successfully navigate the informational thickets of the modern world, replete with both easy-access information *and* misinformation, we need to know what sources of testimony we can trust, how to respond to cases of disagreement, and how to come to know as a group.

5.1 Testimony

According to reductionist views (e.g. Fricker 1987, Faulkner 1998, Lyons 1997) in the epistemology of testimony, testimonial knowledge requires that the hearer have independent positive reasons to believe what they are being told. This sort of view is called 'reductionism' about testimony, because it 'reduces' the justificational force of testimony to the combined justificational forces of perception, memory, and inductive inference.

There are 'global' and 'local' incarnations of this view: According to global reductionism, the hearer is *prima facie* justified in believing based on testimony if and only if she has independent positive reason to believe that testimony is generally a reliable source. The problem with this view is that it is too demanding: the observational basis of ordinary epistemic agents is much too small and limited to allow an induction to the general reliability of testimony.

Local reductionisms require the hearer to have independent positive reasons to believe that the particular testifier involved in the target testimonial exchange is reliable, or reliable on the topic at issue. The problem with local views, is that the hearer needs to do the background epistemic work of checking the credentials of the testifier over and over again, for every speaker and, indeed, for every topic. Since we are physically and psychologically limited creatures, and since we rely on each other for most of our knowledge, this view leaves a lot of testimonially generated epistemic value unexplained. We can, for instance, easily gain knowledge from strangers in the street about the right way to the train station, and other such everyday trivial matters, without having any clue as to how reliable they might be on

the subject matter.

Generally speaking, then, the problem for reductionism is that it lays to heavy a burden on the shoulders of individual cognizer to be empirically plausible, thus leaving a vast array of testimonial knowledge unexplained.

In contrast, according to anti-reductionism, testimonial justification is *sui generis* and, most of the times, fairly easy to come by. There are two main varieties of anti-reductionism defended in the literature: according to social anti-reductionism (Graham 2006, Goldberg 2010, Greco 2020, Simion 2020), some of the justificatory work is done by the social context: by the social norms present at the relevant context, or by the social roles of the parties in the testimonial exchange. Depending on the social context in which the testimonial exchange takes place, the hearer needs to do more or less justificatory work: in contexts with truth-telling norms, for instance, the hearer can just go ahead and believe what they are being told.

There is one main problem with these views: there are epistemically good and epistemically bad social norms. If our epistemology indiscriminately allows social factors to encroach on epistemic normativity, it threatens to license various forms of: (1) epistemic injustice, whereby socially marginalised knowers are not given due credibility. For instance, in societies where women are socially deemed to be intellectually inferior to men, a social-norms-based approach will have the result that one cannot justifiably believe women's testimony. (2) socially-generated epistemic errors: if particular social norms licence trust in unreliable testifiers (due, for instance, to their privileged social status, or popularity), this account risks licensing beliefs based on unreliable sources. What we seem to need is a principled way to distinguish epistemically good from epistemically bad social norms. But if that is the case, the social does not come first: epistemic value does.

A-priori anti-reductionism vindicates the thought that epistemic value comes first, in that they start analysis with what they take the value in question to be, and attempt to derive testimonial justification on a priori grounds. Tyler Burge (1993), for instance, takes it that intelligible presentations-as-true indicate generation by a source endowed with rational abilities, and that rational abilities have the function of generating true contents. Further, according to Burge, if something X has the function of phi-ing, then X reliably phi-s in normal conditions. If so, it will follow that rational abilities reliably generate true content in normal conditions, and, in turn, that intelligible presentations-as-true indicate generation by a source that reliably generates true content in normal conditions. From this, Burge derives the justification claim: one is *prima facie* entitled to believe based on intelligible presentations-as-true.

The most pressing problem with Burge's derivation is function-theoretic, and it concerns reason's dual function: theoretical and practical. Contra Burge, there are many examples (in nature) of functional devices that have two main functions, and in which, when in conflict, one function takes precedence over the other (e.g. our sexual/excretory organs). If conflicting functions are not an exception in nature, however, the practical function of reason may override the theoretical function when they come in conflict. If that is so, it need

not be the case that reason reliably generates true assertions, rather than practically advantageous assertions. See (Leonard 2021) for more discussion and topics in the epistemology of testimony.

5.2 Disagreement

The prevailing methodological approach in the epistemology of disagreement is the study of idealised, perfect-peer disagreement; cases where two perfect peers (roughly, agents in possession of the same evidence) disagree about whether a particular fact obtains. There are two main views on what one should (epistemically) do in the face of peer disagreement: ‘stead-fast’ views (e.g. van Inwagen 1996) hold that the individual cognizer has the right to hold on to her beliefs. The main problem for these views, unsurprisingly, is that one can think of many cases in which holding steadfast seems utterly irrational: think, for instance, about a situation in which, having had dinner at the restaurant, you and your friends all calculate how much everyone’s share of the bill is, and come at different results (Christensen 2007). Intuitively, the reasonable thing to do in this situation is to suspend belief and run the calculus again, rather than dogmatically stick to one’s guns: after all, we all make mistakes.

In contrast, according to ‘conciliatationist’ views (e.g. Elga 2007), when faced with perfect-peer disagreement, the rational thing to do is always to weaken one’s confidence in one’s belief. Conciliatationism easily accommodates the restaurant case. However, the view deals less well with cases of easy a priori knowledge: say, for instance, that my neighbour, whom I believe to be a peer, tells me that he just bought a round square table. Should I weaken my belief that there are no such things as round square tables? Intuitively, the answer is ‘no.’ For more work on disagreement, see (Frances and Matheson 2019).

5.3 Group Belief

What is it for a group – such as a team, or a corporation – to believe something? According to summativism, group belief reduces to the belief of (some) of the group members: a group believes that p iff (some of its members believe that p (e.g. Lackey 2020).

The problem with this account is that it looks as though a group can know a fact even when none of its members form the corresponding belief. The classical case is that of a jury who reaches the right verdict, even though all of its members are deeply racially biased, and thus cannot get themselves to form the corresponding belief.

Non-summativist views about group belief take group believers to be analogous to individual believers, in that they can host beliefs of their own, independently of whether any of their members host the belief in question. Non-summativisms come in two main shapes: acceptance views take group knowledge to not imply belief, but rather a weaker attitude, i.e. joint acceptance (Gilbert 1987). It is easy to see that these views will easily accommodate the case of the jury above. However, the acceptance view has problems explaining widely spread social biases, where conscious acceptance is missing.

Distributive non-summativists (e.g. Bird 2014) mainly conceive of group belief on a collaborative model, where the group members all contribute to the formation of the group belief, but need not actually host it: think of a scientific team, for instance, whereby all scientists, of different specialism, plug their results into a supercomputer that outputs the final proposition *p* based on this input. The thought is that the group thereby knows that *p*, but none of the individual scientists do.

One important difficulty for these views is answering the question as to which contributors to the final output count as being proper part of the believing group. Not all contributions will qualify: the mailperson delivering the correspondence to the group of scientists will not intuitively be part of the group hosting the belief that *p* (where *p* is a complicated scientific claim). As such, distributivism owes us a satisfactory account of group individuation.

See (Goldman and O'Connor 2021) for more developments in social epistemology.

References

- Bird, A. (2007). Justified Judging. *Philosophy and Phenomenological Research*, 74: 81-110.
- Bird, A. (2014). When is there a group that knows? Scientific knowledge as social knowledge. In J. Lackey (Ed.), *Essays in Collective Epistemology*, pp. 42–63. Oxford: Oxford University Press.
- BonJour, L. and Sosa, E. (2003). *Epistemic Justification: Internalism Vs. Externalism, Foundations Vs. Virtues*, Oxford: Blackwell Publishing.
- Brown, J. (2018). *Fallibilism*. Oxford: Oxford University Press.
- Burge, T. (1993) Content Preservation, *The Philosophical Review*, 102(4): 457–488. doi:10.2307/2185680
- Cohen, S. (1984). Justification and Truth, *Philosophical Studies*, 46(3): 279–295. doi:10.1007/BF00372907
- Conee, E. and Feldman, R. (2001). Internalism Defended. *American Philosophical Quarterly*, 38(1): 1–18.
- DeRose, K. (2009) *The Case for Contextualism: Knowledge, Skepticism, and Context: Vol. 1*, Oxford: Clarendon Press.
- Dretske, F. (1970). Epistemic Operators, *The Journal of Philosophy*, 67(24): 1007–1023. Reprinted in Dretske, 2000, pp. 30–47.
- Elga, A. (2007). Reflection and Disagreement, *Noûs*, 41: 478–502.
- Faulkner, P. (1998). David Hume's Reductionist Epistemology of Testimony", *Pacific Philosophical Quarterly*, 79(4): 302–313.
- Frances, B. and Matheson, J. (2019). Disagreement, *The Stanford Encyclopedia of Philosophy* (Winter 2019 Edition), Edward N. Zalta (ed.), [URL = <https://plato.stanford.edu/archives/win2019/entries/disagreement/>](https://plato.stanford.edu/archives/win2019/entries/disagreement/).
- Fricker, E. (1987). The Epistemology of Testimony, *Aristotelian Society Supplementary Volume*, 61: 57–83. doi:10.1093/aristoteliansupp/61.1.57
- Gettier, E. (1963). Is Justified True Belief Knowledge?, *Analysis*, 23(6): 121–123. doi:10.2307/3326922

- Gilbert, M. (1987). Modelling collective belief. *Synthese* 73, 185–204.
- Goldberg, S. (2010). *Relying on Others: An Essay in Epistemology*, Oxford: Oxford University Press.
- Goldman, A. (1967). A Causal Theory of Knowing, *The Journal of Philosophy*, 64(12): 357–372. doi:10.2307/2024268
- Goldman, A. (1976). Discrimination and Perceptual Knowledge, *The Journal of Philosophy*, 73(20): 771–791. doi:10.2307/2025679
- Goldman, A. (1979). What is Justified Belief? in *Justification and Knowledge*, George S. Pappas (ed.), Dordrecht: D. Reidel: 1–25.
- Goldman, A. and O'Connor, C. (2021). Social Epistemology, *The Stanford Encyclopedia of Philosophy* (Spring 2021 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/spr2021/entries/epistemology-social/>>.
- Graham, P. (2006). Liberal Fundamentalism and Its Rivals, Lackey and Sosa 2006: 93–115.
- Greco, J. (2010). *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge: Cambridge University Press.
- Greco, J. The Transmission of Knowledge. Cambridge: Cambridge University Press.
- Ichikawa, J.J. 2014. 'Justification is Potential Knowledge.' *Canadian Journal of Philosophy*, 44: 184-206.
- Ichikawa, Jonathan Jenkins and Matthias Steup, "The Analysis of Knowledge", *The Stanford Encyclopedia of Philosophy* (Summer 2018 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/sum2018/entries/knowledge-analysis/>>.
- Kelp C. (2009), "Knowledge and Safety," *Journal of Philosophical Research* 34: 21–31.
- Kelp, C. and Simion, M. (2016). The Tertiary Value Problem and the Superiority of Knowledge. *American Philosophical Quarterly* 53, 397-409.
- Kelp. C. (2016). 'Justified Belief: Knowledge First-Style. *Philosophy and Phenomenological Research* 93: 79-100.
- Kripke, S. (2011). Nozick on Knowledge, in *Philosophical Troubles: Collected Papers, Volume 1*, Oxford: Oxford University Press, pp. 162–224.
- Kvanvig, J. (2003). *The Value of Knowledge and the Pursuit of Understanding*, Cambridge: Cambridge University Press.
- Lackey, J. (2020). *The Epistemology of Groups*. Oxford: Oxford University Press.
- Lackey, J. (2007). Why we don't deserve credit for everything we know, *Synthese*, 158(3): 345–361.
- Leonard, Nick, "Epistemological Problems of Testimony", *The Stanford Encyclopedia of Philosophy* (Summer 2021 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/sum2021/entries/testimony-episprob/>>.
- Littlejohn, C. (2020). Plea for Epistemic Excuses. In F. Dorsch and J. Dutant (eds.), *The New Evil Demon Problem*. Oxford: Oxford University Press.

- Mill, J. S. (1963). *Essays on Ethics, Religion and Society*. In J. Robson (Ed.), *The Collected Works of John Stuart Mill, vol. 10*. Toronto: University of Toronto Press.
- Miracchi, L. (2015). Competence to Know. *Philosophical Studies*, 172: 29-56.
- Moore, G. E. (1959). Proof of an External World, and "Certainty," in *Philosophical Papers*, London: George Allen & Unwin, Ltd.
- Nozick, R. (1981). *Philosophical Explanations*, Cambridge, MA: Harvard University Press.
- Pappas, G. (2017). Internalist vs. Externalist Conceptions of Epistemic Justification, *The Stanford Encyclopedia of Philosophy* (Fall 2017 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/fall2017/entries/justep-intext/>>.
- Parfit, D. (1984). *Reasons and Persons*. Oxford: Oxford University Press.
- Plato. (1956). The Meno. In W. Guthrie (Ed.), *Protagoras and Meno* (W. Guthrie, Trans.). London: Penguin.
- Pritchard, D. (2005). *Epistemic Luck*, New York: Oxford University Press.
- Pritchard, D. (2009). Apt Performance and Epistemic Value, *Philosophical Studies*, 143(3): 407-416.
- Pritchard, D.H., Millar, A. and Haddock, A. (2010). *The Nature and Value of Knowledge: Three Investigations*. Oxford: Oxford University Press.
- Pritchard, D., Turri, J. and Carter, J.A. (2018). The Value of Knowledge, *The Stanford Encyclopedia of Philosophy* (Spring 2018 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/spr2018/entries/knowledge-value/>>.
- Pryor, J. (2000). The Skeptic and the Dogmatist, *Nous*, 34: 517-549.
- Reed, B. (2010). A Defense of Stable Invariantism. *Noûs*, Vol. 44, No. 2, pp. 224-244.
- Rysiew, P. (2021). Epistemic Contextualism, *The Stanford Encyclopedia of Philosophy* (Spring 2021 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/spr2021/entries/contextualism-epistemology/>>.
- Silva, P (2020). Knowledge-first Theories of Justification. Internet Encyclopaedia of Philosophy, <https://iep.utm.edu/kft-just/>
- Simion, M. 2021. *Shifty Speech and Independent Thought*. Oxford: Oxford University Press.
- Simion, M. Testimonial Contractarianism. *Nous*. Online First.
- Simion, M. 2019. Knowledge-First Functionalism.' *Philosophical Issues* 29 (1): 254-267.
- Sosa, E. (1999). How to Defeat Opposition to Moore, *Philosophical Perspectives* 13, 141-154.
- Sosa, E. (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge* (Volume I), New York: Oxford University Press.
- Sosa, E. (2004). Replies, in Greco (ed.) *Ernest Sosa and His Critics*, Cambridge, MA: Blackwell: 275-325.
- van Inwagen, P. (1996). It is Wrong, Always, Everywhere, and for Anyone, to Believe Anything, Upon Insufficient Evidence, in J. Jordan

- and D. Howard-Snyder (eds.), *Faith, Freedom, and Rationality*, Hanham, MD: Rowman and Littlefield, pp. 137–154.
- Weisberg, J. (2021). Formal Epistemology, *The Stanford Encyclopedia of Philosophy* (Spring 2021 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/spr2021/entries/formal-epistemology/>.
- Williamson, T. (2000). *Knowledge and Its Limits*, Oxford: Oxford University Press.
- Williamson, T. (2020). Justifications, Excuses, and Sceptical Scenarios. In J. Dutant and F. Dorsch, (eds.), *The New Evil Demon*. Oxford: Oxford University Press. Archived in Phil
- Wright, C. (1985). Facts and Certainty, *Proceedings of the British Academy*, 71: 429–472.