

Chapter 5: Testimony and Applied Epistemology

I would like to start this chapter with a very brief, very oversimplified bit of history. The reason for this, as I would like to tell it, is that the attitudes towards epistemology that I've been quarrelling with come from a highly contingent episode in the history of Western philosophy, and they are not as inevitable as they might seem.

It's probably fair to say that the vast majority of students who become familiar with epistemology through the Western canon will be largely introduced to the topic by the figures from the early modern European enlightenment. The usual protagonists are René Descartes, John Locke, David Hume, Immanuel Kant, and perhaps Isaac Newton and Galileo Galilei. Each of these figures made important contributions to epistemology that effectively shaped the field as we know it. However, they also all shared a common assumption in their thinking about the topic, which continues to shape the prejudices of those whom they have influenced. If we're not careful to examine this assumption, we run the risk of unconsciously adopting it ourselves. The common assumption that I'm speaking of is this: that knowledge is always an *individual* achievement.

For each of these philosophers, their model of the rational, enlightened inquirer is one who pursues developing their own worldview on their own terms. They gather all of their evidence on their own, they trust in their own senses and wits above everyone else's, they do not take any authorities for granted, and they always do their own research. They only accept something as true if they alone can verify it.

Think, for example, of the protagonist from Descartes's *Meditations*. Descartes presents a picture of a lone inquirer who sits beside his fireplace silently contemplating to himself. His isolation allows him to rebuild his worldview without any outside interference or distraction. It is within this solitary, isolated state that this figure sorts out what to believe from what not to believe. In the end, the only beliefs that he deems acceptable are the ones that he can clearly and distinctly perceive to be true, for *himself*.

I take it that this is emblematic of the intellectual virtues that were promoted during the European enlightenment: *conduct your own experiments, be skeptical of the authorities unless you can prove that they're credible, always research for yourself*. And given the intellectual climate of early modern Europe, there are a couple of obvious reasons as to why these values seemed so important at the time. The continent was still in the middle of the Protestant revolution and the scientific revolution. Before then, the Catholic Church had maintained total sovereignty over all types of intellectual activity. When it came to matters of mathematics, logic, physics, natural history, or religion, everyone before then had to defer to the authority of the church. All of this was upended when

Martin Luther exposed the religious authorities for demagoguery and Galileo proved that the earth revolves around the sun. Remember all of that fuss that Galileo caused? Apparently his adversaries refused to even look through his telescope. They found it preposterous that anything other than Aristotle's physics could be the final truth.

Like Galileo, Descartes was also a physicist with iconoclastic views. He also had his own theories that challenged the authority of Aristotle, and he faced significant resistance from the orthodox intellectual community. Bearing this in mind, consider again Descartes's method of doubt. Remember, Descartes recommends to his reader to purge their belief system by meditating on radical skepticism. Then, once your belief system has been destroyed, you rebuild again based on the evidence that you've acquired *for yourself*, freed from the interfering dogma of the dominating institutions. Could it be that Descartes had ulterior motives for promoting his method of doubt, besides the sheer love of truth and pursuit of knowledge? Here is a letter that Descartes wrote to his friend Mersenne:

I will say to you, just between us, that these six Meditations contain all the foundations of my Physics. But, please, you must not say so; for those who favor Aristotle would perhaps have more difficulty in approving them; and I hope that those who will read them will unwittingly become accustomed to my principles and will recognize the truth, before they notice that my principles destroy those of Aristotle.

When the historical Descartes invites his readers to suspend judgment on everything that they think they know, he is tacitly hoping to unearth their dogmatic confidence in Aristotle's physics. He's sowing the seeds so they'll become more open towards his own scientific theories. Given his intellectual climate, this is a very clever move. But who would've thought that the foundational text of modern philosophy had such a sneaky motive?

I would now like to suggest that something like this motive can explain the epistemological values that were the hallmarks of enlightenment philosophy. It is important to understand that many of the canonical figures from this period were essentially reactionaries. They were revolting against the dominant religious institutions that had established hegemony over the philosophical and scientific communities in Europe. Given this backdrop, *it made sense* for them to advocate for inquiry as a solitary enterprise.

This reactionary sentiment comes out most clearly when it comes to the topic of testimony. For many early modern philosophers, the entire topic is noticeably absent from any of their epistemological discussions. For the few who actually wrote about it—Locke and Hume—their take is extraordinarily skeptical. The general tenor of their discussion on testimony is that we cannot gain knowledge from the testimony of others (alone) because we can never be sure that other people are telling the truth—they could be lying to us or they could be mistaken. If we ever gain knowledge from testimony, it is only because we put in our own work to make sure that our informants are reliable sources,

that we've checked their track-record of telling the truth, and so on. Either way, knowledge is only possible if it's ultimately the result of our individual efforts.

I think that this degree of skepticism towards testimony is too extreme; but again, it is understandable given the historical climate. In fact, I will be making my case for the first of these claims in this chapter. I will argue that we can attain knowledge from testimony and that we need not understand this knowledge as a result of individual efforts. But before we go into the theory, it is worth pausing to remark on how different our situation is from that of the early modern philosophers.

Needless to say, our intellectual climate is not very close to Descartes's or Hume's. For one, our scientific disciplines are no longer in the grip of an autocratic religious institution. There is no organized religion dictating the permissible results of physics, biology, natural history, or medicine according to some pre-established holy text. Not only that, but it's hard to ignore the remarkable successes that our sciences have delivered since that time. We've been able to harness the energies of the atomic world, extend our lives by controlling the microscopic world, and predict all manner of things with striking accuracy. Certainly, there are many facets of our scientific institutions that deserve a level of trust that wasn't yet earned in Descartes's time.

Our scientific theories now wield knowledge that wasn't available to previous generations. Insofar as these theories contain this knowledge, they thus command our assent. Nonetheless, we mustn't be too sanguine about it either. There are other forces in the modern world that can distort our view of this progress.

One such distorting force is the patent plethora of anti-scientific conspiracy theories that are making their rounds through public space. You know the ones that I'm talking about: climate change denialism, the anti-vaxxer movement, the conspiracists about COVID-19, young earth creationism, and the rest. These are the people who openly oppose the dominant scientific establishment. They think that the credentialed experts do not deserve the authority that has been given to them.

There's a certain irony involved in the rhetoric of these people. Often you'll find that they're now the ones who espouse enlightenment-era epistemic values, echoing Descartes and Hume. *Don't trust the authorities; research things for yourself.* It's ironic, because, as I said, it was precisely the scientific luminaries who espoused these values as a reaction to the church. In modern times, those same values have become coupled with an anti-science agenda.

Another force of distortion is the cacophony of false or misleading news stories and studies that are propagated by various outlets. For pretty much any crazy suspicion you might have, it's not difficult to "do your own research" to confirm what you suspect. If you think that some scientific result is a mistake, or some news story is fabricated, you can easily discover the rabbit hole of amateurs, posturing as heterodox experts, who purport to have proven your suspicion. The problem is that most people lack the relevant background and tools that are required to properly evaluate evidence in the areas that they lack expertise. Without that background, it's easy to get fooled by fallacious

reasoning, whereas a genuine expert can spot the fallacies. It's also easy to become too attached to the idiosyncratic views of a contrarian thinker, whereas the genuine experts see the full range of viewpoints in their field.

Altogether, we live in hazardous times for the independent thinker. So what can we, as philosophers, do about this? Well, the first thing to say is that Descartes's model of the isolated inquirer is a myth. The number of topics that we can make non-testimonial, first-hand, authoritative judgments on, compared to the rest of what we think we know, is exceedingly small. At most, they just include the *a priori* subjects that we study in school, mundane facts about the locations we inhabit and have visited, and the events of our friends' lives that we were present for. But this amount of knowledge pales in comparison to what we need in order to be informed citizens of modern society. It excludes practically all knowledge of science, history, politics, and current affairs. Nearly all of the time, we cannot possibly double-check the experts without relying on the further testimony of other experts. As a result, we must accept that there's a division of epistemic labour. Knowledge must be a social phenomenon; it cannot be up to the individual.

Here is the plan for this chapter. First, I'm going to look at the arguments for skepticism about testimony and individualism about testimonial knowledge. In contrast to these positions, I will be arguing that testimony is a basic source of knowledge, and that the conditions for gaining testimonial knowledge depend more on the community than the individual. I will say up front that my theoretical principles will largely echo what has already been developed in previous chapters. The view that I favour will retain a strong element of *externalism*. Since the context is testimony, the externalist element is especially likely to invite the worry that my view isn't suitable for giving practical advice. So, for the second half of this chapter, I would finally like to address this worry. That is, I want to finally spell out how I envision the relation between foundational epistemology and practical concerns about what to believe.

1 The division of epistemic labour vs. individualism

Let me begin by being more explicit about the contrast between two ways of thinking about testimonial knowledge. The contrast is roughly whether we hold that the onus for testimonial knowledge is on the individual or on their community at large. On one way of seeing things, knowledge is an individual achievement; on the other, it is a communal good. We can make these rough ideas a bit more precise by introducing some official definitions. Here is the substance of the two views:

The individualistic view. If S knows that p on the basis of someone else's testimony, then S must adduce reasons, which are ultimately non-testimonial, for believing that their informant is telling the truth.

The division of epistemic labour (DOEL) view. If an informant knows that p and asserts that p, and S believes that p on the basis of their assertion, then S thereby knows that p.¹

(For those who are familiar with the literature on testimony, what I am calling the “division of epistemic labour view” is more often called “non-reductionism” or “the direct view.” Its foremost modern defender is Jennifer Lackey. The “individualistic” view covers two views that are on offer in the contemporary literature: reductionism and skepticism. Reductionism claims that we can get testimony knowledge by adducing the appropriate reasons in favour of the testimony; skepticism says that *if* testimonial knowledge is possible, there must be strong enough reasons available to the subject to support its truth, but since there aren’t strong enough reasons, we cannot (in fact) gain such knowledge.)

Notice how the labels are appropriate for each view. For the individualistic view, it really is *up to the individual* to procure reasons or evidence that their informant is telling the truth, and these reasons cannot ultimately circle back to testimony. So if they are to learn anything from testimony, they must rely on *their* personal experiences, *their* reason, *their* prior commitments, and so on, in order to build their case for believing their informants. Again, the onus is on the individual.

Contrast this with the DOEL view. For them, knowledge is not simply up to the individual: it is a social enterprise. It is something that, if you possess it, then you are also able to share it. Others can gain knowledge of what you know on behalf of your generosity.

It might be fair to say that the individualistic view likens other people to encrypted stores of information; you can gain knowledge from them, but only if you decipher the right codes. The DOEL view, on the other hand, sees other people as free-flowing conduits of information. We can (so to speak) extend our cognitive reach beyond our first-personal experiences by using the testimony of others. When an informant tells me of an event that they have witnessed first-hand, it is as if they are opening a channel that relates my epistemic situation to facts that they have witnessed. (Just as I can use a telescope to see the

¹It is common to articulate the direct view of testimonial knowledge with the inclusion of a ‘no defeater’ clause. That is, one can know on the basis of knowledgeable testimony *provided that there isn’t any funny business that spoils the knowledge*. For example, if an informant knows that p, and they assert in the same breath both *that p is true* and they assert *that p is false*, it would seem highly counterintuitive to credit me with knowledge if I choose to accept only their first assertion. The common response on behalf of the direct view is to say that the second assertion acts as a defeater so that I cannot rationally accept only the first (on the basis of testimony alone), and hence I cannot gain knowledge that p in this case. But contrary to common sense, I’m not going to include a no-defeater clause to the DOEL view. Instead my preference is to lean into the externalism and then reiterate the line that I took in chapter 2, section 5.1.

This is, without doubt, the worst thing that I’ve said in this entire project: that one can gain testimonial knowledge from an informant that contradicts themselves. But it isn’t strictly necessary to say this, given my general framework. I could try to include a ‘no defeater’ clause into DOEL and then leave the rest of my overall picture as it is. The reason why I have chosen not to do this is simply because I’m doubtful that the notion of a defeater can be worked out in a satisfactory way. So instead I will bite this one unintuitive bullet.

rings of Saturn, I can use the testimony of another person to “see” the facts about events that I wasn’t physically present for.) It’s as if your informants can serve as your eyes and ears to the facts that you do not witness first-hand. Of course, the DOEL view doesn’t ignore the fact that these channels can often be corrupted when other people fail to tell the truth, or assert something they don’t know. Their only claim is the conditional one: *if* the testifier knows what they assert, *then* knowledge can be attained through their testimony.

It is worth observing that the DOEL view has a curious implication about the status of testimonial knowledge. Since it doesn’t require the subject to infer the truth of the testified proposition from prior premises about the reliability of the informant, it follows that testimonial knowledge can be *foundational* (in the sense of chapter 1). If an informant of mine witnessed first-hand that p and thereby knows that p , and later asserts to me that p , then, according to the DOEL view, I can come to *know* that p on the basis of their testimony, without the help of inference from other beliefs of mine. This might sound like a peculiar consequence. It might sound weird to call a testimony-based belief an item of *foundational knowledge*. But if this sounds weird to you, it is important to remember the definition we gave for foundational knowledge: foundational knowledge is simply knowledge gained without the use of inference from other beliefs. So if testimonial knowledge doesn’t require inference, then it qualifies as foundational. But this admission does not carry with it any of the other connotations that “foundational knowledge” has in the cartesian tradition. It does not imply that testimonial knowledge is *certain*, or *impossible to doubt*, or *a priori*, or *immune to rational revision*. (Although, *as knowledge*, testimonial knowledge must be *true*. So on that score, it cannot be revised without losing a true belief.)

It should be apparent that the enlightenment figures that I mentioned in the beginning section are more closely allied with the individualistic view of testimony. As a result of their out-sized influence in philosophy, it is this view of testimony that has been historically popular in the West. But the DOEL view has also had its share of defenders. Within early modern European philosophy, it was advocated by Thomas Reid and other proponents of the Church. Going back even further, the view was advanced in the Nyāya Sūtras in India, thousands of years before the issue was given any serious attention in the West.

2 In defence of the DOEL view

As I’ve already indicated, my own preference is for the division of epistemic labour view. I hold that if the conditions are right—that is, if the informant knows what they assert—and the subject believes them, then that’s all it takes for the subject to gain knowledge. No further ratiocination is required on behalf of the subject.

My reasons for taking this view are largely theoretical. It is the picture of testimonial knowledge that best coheres with the general epistemological framework that I have developed in chapters 1-3. To be clear, it isn’t strictly *implied*

by the foundationalist, externalist, ‘one-world’ understanding of knowledge that I have been advocating. But it does share the same motivations. So given my previous commitments, it makes sense for me to take this stance towards testimony.

In chapters 1 and 2 I developed a theoretical framework for knowledge that was both *foundationalist* and *externalist*. According to my main proposal, a belief counts as foundational knowledge when it is the product of a channel that relates the subject’s cognition to reality. Now, in previous chapters, we primarily attended to perception as one such channel; but there are others. In my view, testimony is one of them. Oftentimes the knowledgeable relations that we bear to the facts are mediated through the knowledge and testimony of other people.

In fact, if you take the DOEL view, then there is a close analogy between perceiving that something is the case and learning that something is the case from a knowledgeable informant. In one situation, I perceive that it is raining outside by looking out my window. My perception has thus put me into direct contact with the facts about the weather in my vicinity, through the use of my eyes. (Of course, in order for this to happen, my perception must be working properly. But that is just a necessary condition for perception-based knowledge; it need not figure into the reasons I adduce for my belief.) Since my belief is based on an information channel that has put me into contact with the facts (i.e. perceiving), I thus obtain foundational knowledge. In the other situation, my roommate comes into my apartment after being outside and tells me that it is raining. Although I haven’t seen the weather myself, my roommate is acting as my eyes and ears. They have perceived the state of the weather, and so they know the facts about the weather outside. By testifying about the weather, they open up an information channel for me to get into contact with those facts, through the use of their knowledgeable testimony. (Of course, in order for this to happen, the testimonial channel must be working properly; they must know what they’re asserting. But that is just a necessary condition for testimony-based knowledge. In order for me to attain knowledge through them, it isn’t necessary for me to form any beliefs about their knowledgeability; they just have to be knowledgeable.) Since my belief is a product of this link between me and the facts, my belief counts as foundational knowledge.

One of the key pillars of the theory of knowledge developed in previous chapters was its commitment to externalism. This is the claim that a belief’s status as knowledge isn’t solely a function of how things appear to the subject, or how things seem from their point of view. It is possible for there to be two subjects (e.g. me and my brain-in-a-vat counterpart) that share all of the same appearances “from the inside”, and yet one has knowledge that the other doesn’t. This is because knowledge partly depends on relations borne by the subject to their external environment, which hold as an objective matter of fact and need not be distinguishable by the appearances. Well, something very much like this idea is upheld by the DOEL view of testimony. This view also permits there to be two subjects who have each taken the same measures on their end to attain knowledge through testimony, and yet one succeeds and the other fails.

Suppose that you and I are new to a city and are each looking for directions to the city centre. One stranger (a knowledgeable one) tells me that it is to the East. Another stranger (an ignorant one) tells you that it is to the West. Both you and I are equally trusting of our respective strangers and have done equally as much (or as little) to ensure their reliability. Nonetheless, I succeed in getting knowledge (because my stranger is knowledgeable) and you do not. (Although you aren't to be blamed for being duped.) The reason is simply that my knowledgeable informant related me to the facts whereas your informant didn't.

The externalist understanding of knowledge essentially sees knowledge as the product of the subject's cognitive efforts *and* the relations they bear to their external environment. Well, the DOEL view is very much in agreement with this way of understanding knowledge. The subject's cognitive situation, considered in isolation, is the wrong place to look to determine whether they possess knowledge from testimony. Instead, one must look to their broader social context. If the subject belongs to a community that possesses a lot of knowledge and passes this knowledge around, then the individual has a lot of knowledge available at their fingertips. If not, then not.

As you can see, the division of epistemic labour view of testimony fits very well within the general understanding of knowledge that I've already laid out and defended. So my *main* argument in its favour is really just to gesture at the general framework developed so far. That framework, recall, is a combination of foundationalism and externalism that understands foundational knowledge as based on real informational links between the subject and the world. But still, there is more to say than just that. There are also the reasons that I have given for this framework, which we can import specifically into the testimonial setting. So what were those reasons, again? In short, the main motivations in chapters 1-3 have to do with its handling of various types of skepticism. My foundationalist, externalist, one-world picture of knowledge explains where various skeptical patterns of reasoning go wrong, in ways that are unavailable to rival pictures of epistemology. So in order to display the real virtue of the DOEL view, we should consider the arguments for skepticism towards testimony.

3 Skepticism and testimony

Now, what reasons are there for being skeptical towards testimony? At first it might seem as if there's an abundance of skeptical problems for us to consider. For example, a skeptic might argue that we can't know that our informants are knowledgeable or reliable and we can't distinguish the real experts from the frauds. But on closer inspection, we find that these skeptical arguments are not all that different from the ones we've encountered before.

Take, for instance, the skeptic that wants to ask, "how do you know that your informant is knowledgeable?". The response to this skeptic, from the DOEL perspective, is that it is simply not a requirement on testimonial knowledge that the subject must produce an answer to this question. We can gain

knowledge from testimony simply by trusting knowledgeable informants. For this to work, the informants must actually *be* knowledgeable, but the subject who listens to them is not required to know this about them. (This response should be familiar from §2.4.1 and §2.4.2, albeit making the necessary changes for the testimonial setting.) Again, this doesn't preclude the subject from also developing knowledge about the knowledgeability of their informants. Indeed, developing a body of knowledge about who is trustworthy and who isn't is a *great thing*; and it is a practical necessity of social life. But this knowledge isn't a necessary condition on the mere transmission of testimonial knowledge.²

We will get back to the point about practical necessities later. For the time being, it is important to look at the skeptical argument that stems from considerations about sensitivity to the truth. Here, the skeptic will argue that you cannot gain knowledge from testimony since it is possible for the informant to be ignorant, or lying to you, without you realizing it. Since you cannot rule this possibility out, you cannot know that what your informant says is true.

This argument is highly reminiscent of the sensitivity argument discussed in §2.4.3. Let's remind ourselves of how that argument went. Throughout chapter 2, we were looking at highly mundane claims to knowledge: specifically, the knowledge that I have hands. The skeptic aimed to contest that knowledge by raising the possibility of a radical skeptical scenario in which my belief would be false and yet things would appear to be the same. In that chapter, their favoured skeptical scenario was the one where I'm a handless brain-in-a-vat that is fed hallucinatory experiences. The argument proceeds by claiming that, since I cannot rule this possibility out, it follows that I don't really know that I have hands. And the reason that I cannot rule out this possibility is (roughly) because things would seem to be the same "from the inside" whether or not I was a BIV. My belief that I'm not a BIV isn't sensitive to the truth. Just to repeat, the general pattern of the argument goes like this. So you think that you know that P? Well let me present to you a scenario where Q, such that if Q is true then P is false, and if Q is true then things will still *appear to you* as if P is true and Q is false. Since you cannot rule out Q, it follows that you cannot know that P.

Now that our memory is refreshed, I would like to convince you of two things. First, I would like to make the case that these skeptical arguments are not merely academic exercises. Despite the impression that it is solely within the interest of theoretical philosophy, the study of these brain-in-a-vat-style skeptical arguments actually have a great deal of practical importance outside of the philosophy classroom. That is because the *pattern* of skeptical reasoning can actually crop up quite frequently in ordinary reasoning and public discourse. Secondly, I would like to once again pinpoint where these arguments go wrong. Doing so will require me to echo the points that I made in chapter 2, but this

²Perhaps contrary to first appearances, this claim is highly plausible. Don't think of the paradigm of testimonial knowledge as deciphering the facts about current events from the news. Rather, think of a child learning their first share of mundane facts from their parents. Clearly the child isn't required to have complex knowledge about the reliability of their parents in order to absorb their first bits of elementary knowledge.

time rephrased in terms of the DOEL view of testimonial knowledge. If my approach is right, then this will also complete my motivations for preferring the DOEL view over the individualistic view.

It is a curious fact about the brain-in-a-vat-style skeptical arguments that most people, upon first considering them, find their premises to be intuitively compelling. Most people are tempted to view them as sound when they see these arguments in isolation. And yet, most people do not *act* as if they do not know that they have hands. This shows something deep and important about our psychological proclivities: it shows that we have a deep tendency to be irrational. I think that Michael Huemer puts this point best:

The fact that skeptical arguments seem plausible to us... suggests that we hold a set of very stringent criteria for justified belief—criteria so strict, in fact, that they can be used to rule out any proposition whatsoever from being considered justified. But the fact that we reject skepticism and accept common sense beliefs indicates that, at the same time, we hold a much looser set of criteria for justified belief—criteria that allow lots of propositions to be considered justified. Here’s something that might happen: I come upon claim A, which I happen to like. So I apply my loose standards of justification, and find A to be justified, whereupon I accept it. Then, some time later, I come upon claim B, which I don’t like. So I apply my strict (skeptical) standards of justification, and find B to be unjustified, but my inconsistent standards enable me to believe whatever I like... My experience is that the human capacity for self-deception is both vast and subtle. It enables us to seize upon any available tools for maintaining the beliefs that we prefer, while avoiding full consciousness of its own operation. In fact, it takes a concerted, conscious effort not to engage in this otherwise automatic faculty.³

Huemer is clearly right about what these skeptical arguments reveal about our psychology. If we aren’t careful and self-critical, we can easily find ourselves employing skeptical reasoning in preferential ways to protect our cherished beliefs. This alone gives us excellent—and not merely academic—reason to study the skeptical arguments. It helps us to alert ourselves for when they occur in the wild.

To add to Huemer’s point, it is important to see that this skeptical pattern of reasoning often does occur in the wild. And in no place is this more common than the beliefs that we gain from testimony. Consider any important and highly politicized belief that you think that you know. I, for one, think that I know that climate change is real, that the experts are telling the truth, and that I can know this from their expert testimony. Yet it’s not hard to imagine a climate change skeptic arguing, “So you think you know that climate change is real? Well, what if the climate scientists are all part of a conspiracy where

³Michael Huemer, *Skepticism and the Veil of Perception*. Huemer speaks of justification in this passage whereas I’d prefer to replace talk of justification with talk of knowledge.

they're deceiving you for their financial gain. You don't know that they're not!." Although they wouldn't make their own argument explicit, here is how it goes:

- 1 If you know that climate change is real, then you must know that the climate scientists aren't paid off to deceive you.
- 2 You do not *know* that the climate change scientists aren't being paid off to deceive you.
- 3 You do not know that climate change is real.

Here is their argument for premise 2:

- i Whether we live in a world where the climate change scientists are telling the truth or deceiving you, *things will appear the same from your point of view*.
- ii If things would *appear the same*, then they are the same with respect to knowledge.
- iii If the climate scientists were deceiving you for nefarious purposes, then you couldn't know that they weren't.
- iv Therefore you don't know that they aren't deceiving you.

Here's another example of the same sort of thing. You might think that you know that Supreme Court Justice Brett Kavanaugh assaulted Christine Blasey Ford. Well, you can imagine (because it happened often) a Kavanaugh apologist arguing that *you can't know* because *let me describe a scenario where it's not true*, such that *things would seem the same to you from your point of view*. And here's another. You can imagine an election where some despot claims victory even though the initial ballot numbers show that he lost. Afterwards, his supporters argue that *you can't know that vote count wasn't rigged against him*. The desired effect of raising this skeptical point is to prevent you from protesting the rule of their despot. Once you start to look for it, the brain-in-a-vat-style skeptical arguments start cropping up everywhere. You think you know that P? Well *what if Q* (where *if Q then P*, and *if Q then things would seem as if P*). Therefore you don't know that P.

Not only is this style of reasoning prevalent, but it can be put to great mischief. Because, as Huemer observes, we all have a psychological tendency to sympathize with such arguments. And it's not hard to see how bad actors can exploit these tendencies of ours to promote their selfish ends. Those who have an interest in polluting the planet only need to raise skeptical doubts about climate change in order to lull people into compliance. Powerful people who have an interest in avoiding accountability for their transgressions only need to raise skeptical doubts about their actions. The brain-in-a-vat-style skeptical arguments can be weaponized to silence their opponents.

In his book, *How Propaganda Works*, Jason Stanley points out that those who promote conspiracy theories oftentimes don't even believe the theories that

they promote. They raise them as hypotheticals that they want *other people* to take seriously. The motive is to place the burden on their opponents to *disprove* their conspiracy theories. (Think of how the conspiracy theorist will often hedge their assertions: “I hear a lot of people are saying that...”, “what if it turned out that...”, “we should really look into this.”) The effect is that the public sphere gets so flooded with garbage that a general skeptical attitudes prevails and it becomes harder for truth to flourish. And when people are left in the dark, it becomes harder to hold bad actors accountable.

Sometimes the propagandists are open that this is what they’re doing. In 2002, the Republican strategist Frank Luntz wrote a memo to President George W. Bush concerning the framing of the issue of climate change. He admits to Bush that “the scientific debate is closing against us”, and so advises him to defend the administration’s policies by promoting skepticism; “Should the public come to believe that the scientific issues are settled, their views about global warming will change accordingly. Therefore, you need to continue to make the lack of scientific certainty a primary issue.”⁴ (Luntz has since recanted about his approach to this issue, but the damage done by this strategy lives on.) The point is, you should beware the people who dwell too much on hypotheticals for the purpose of promoting skepticism in the public sphere. It tends to coincide with nefarious politics.

Okay, let’s get back to the theory. Is this pattern of skeptical reasoning sound? According to my own general principles, it is not. The DOEL view gets off board precisely at subpremise (ii). As a form of externalism, we claim that testimonial knowledge is not solely a function of how things appear to the audience. It isn’t incumbent upon the audience to reconstruct any sort of argument as to why they can trust the appearances, or anything like that. Rather, as long as the informant really knows what they assert, the audience can gain knowledge just by believing them.

To make this objection more precise, we need an explicit formal reconstruction of the skeptic’s pattern of reasoning. Suppose that an informant asserts to me that P and I come to believe them. Here is the general formula for how the skeptic will challenge my knowledge.

- 4 If I know that P, then I know that my informant isn’t deceiving me.
- 5 If I know that my informant isn’t deceiving me, then if they *were* deceiving me, then I would know it.
- 6 But I wouldn’t know that my informant was deceiving me if they were.
- 7 So I don’t know that my informant isn’t deceiving me.
- 8 Therefore, I do not know that P.

A reader familiar with the material from chapter 2 will recognize that this argument relies on the “simplified sensitivity principle.” The idea is that knowledge

⁴Burkeman, Oliver (March 4, 2003). “Memo exposes Bush’s new green strategy”. *The Guardian*. London.

must be sensitive; if you know that P , then you shouldn't be easily tricked into believing P if it were false. Many of the arguments for skepticism towards testimonial knowledge appeal to something like the sensitivity principle.

We have already found the first step to disarming such arguments. The “simplified” sensitivity principle which underlies premise (5) is false (see §2.4.3). The only kernel of truth that it contains is that the *basis* on which we form our knowledge must be sensitive to the truth.⁵ This was called the “basis sensitivity principle” and it says that *if you know that P on the basis B , then if P were false you wouldn't believe that P on basis B* . In other words, your basis for knowledge cannot easily trick you into believing falsely.

Since I am willing to accept this principle, the focal point of the skeptical argument will now become whether our testimonial beliefs are founded on a sensitive basis. If they are not, then testimony cannot deliver knowledge and the skeptic wins. If there is a sensitive basis for testimonial knowledge, then this kind of skeptical argument can be defused. We defuse it in the same way we did in §2.4.3: we argue that testimony can deliver knowledge, and that when it does we can run the Moore-shift and *know* that our informants aren't deceiving us. And so we come to the question: what is the basis for testimonial knowledge?

4 The transmission of knowledge

The question of bases is more finicky when we're considering testimony than when we're considering other pathways to knowledge, such as perception. When it comes to perception, we're dealing with a factive state. This means that when S *perceives that p* , it follows that p must be a fact. If p weren't a fact, then S could not perceive that p ; at most, S could only *apparently perceive that p* . So when it comes to perception—that is, *genuine perception*—its sensitivity is pretty much guaranteed. But obviously testimony is not like that. A proposition being asserted is no guarantee of its truth (*obviously*). Testimony, *as such*, is not a truth-sensitive basis for belief. So in order to vindicate our knowledge through testimony, we clearly cannot identify the basis for this knowledge as testimony *as such*.

What can we say instead? Well, here is what the individualistic view of testimony will say. They will say that the basis for testimonial knowledge is limited to what's apparent to the audience. If someone asserts to you that p , then here is what you have to go by: the fact that they said that p (or appear to have said that p) and your clues as to their trustworthiness (or apparent trustworthiness, from their appearance, their past track record, etc.). These are the things that are allegedly supposed to support your knowledge through testimony.

The problem with their view is that the elements that comprise their basis are not sensitive to the truth. This is easy to see. Take two situations where an apparently trustworthy person tells you something and in each case you

⁵At least, I think that this principle is true. But this is controversial—as discussed in chapter 2.

believe them. In one case they're telling you what they know, in the other one they're telling you something false, and by all appearances they seem equally trustworthy in each case.

The upshot for the individualistic view is that it is in tension with the basis sensitivity principle, provided that the basis for knowledge is what they say it is. So if this is the view that you favour, you have two options: either you succumb to skepticism or you must reject the basis sensitivity principle. Just as before, there are many "internalist" philosophers who would rather take the second option, but I don't envy their choice.

At this stage, it is important to pay attention to the differences between perception and testimony. Here's an obvious one. Unlike testimony, perception is capable of generating *new* knowledge. When I perceive something first-hand, I can thereby gain knowledge of it that's possible for nobody to have had before. Testimony, on the other hand, does not generate new knowledge. In order for me to obtain knowledge through testimony, my informant must possess the knowledge before me.

For this reason, there is something quite similar between testimony and *memory*. Like testimony, memory does not produce any new knowledge that didn't exist before. Rather, it *preserves* knowledge through time. At some previous time, I gained some new knowledge through one of the channels that generate knowledge (e.g. perception). Later on, I conjure up this knowledge again by remembering what I had previously known.

The process of storing information in one's memory and bringing it to mind later is not, in itself, a sensitive basis for knowledge.⁶ Sometimes it produces false memories, and sometimes the original information that was stored was false. Nonetheless, it is still eminently plausible that memory can *preserve* knowledge through time. Let's say that yesterday I perceive that it was raining outside. Today I remember what I perceived. It seems deeply right to say that my memory has served to preserve the knowledge that I had on this past basis (perception). I know that it was raining yesterday *because* I perceived it in the past, and my memory preserved this knowledge. So now we raise the all important question: what is the *basis* for my current knowledge? The answer comes from after the "*because*" in the previous sentence. My basis for knowing that it was raining yesterday is *past perception preserved through memory*. Identifying the basis for knowledge in this way ensures that the basis is sensitive to the truth.

Now let's take the situation back to testimony. Just as memory can preserve knowledge from a past time-slice of a person to a future time-slice of the same person, testimony can do something analogous between distinct people. Let's say that my roommate perceived that it was raining a few minutes ago and then comes inside and asserts to me that it is raining, and I believe them. Es-

⁶I have switched to talking about the *processes* of information storage in memory rather than the *state* of having remembered because it is plausible to think that *S remembers p* entails both *p is true* and *S knows that p*. So the state of having remembered *is* sensitive. But it isn't the basis of new knowledge; it's the state one has when one conjures up old knowledge.

entially what's going on is my roommate is gathering knowledge through their perception, and then preserving and transmitting this perceptual knowledge to me. So, then, what is the basis for my knowledge? Answer: *my roommate's past perception transmitted through testimony*.

I would like to now suggest that this is how we should think of the basis for testimonial knowledge. Testimonial knowledge (when it is actually knowledge) is always based on *another agent's original basis for their knowledge and it gets transmitted through testimony*. If the knowledge was originally based on perception, then my testimonial knowledge is based (in part) on *the original witness' perception*; if the original knowledge was *a priori*, then my testimonial knowledge is based (in part) on *the original subject's a priori reasoning*; and so on.

If I'm right about this basis, then we do, in fact, have a truth-sensitive basis for testimonial knowledge when it is the genuine item. That is because the *original* basis for this knowledge (the testifier's basis) must be sensitive to the truth (according to my previous commitments), and the basis for testimonial knowledge includes *theirs*. So when I base my belief on my roommate's testimony, I'm basing my belief (in part) on *their* perception, and this basis wouldn't be available if what they said was false. Thus the "basis sensitivity principle" is in fact satisfied.

It would be nice if I could now declare to have dissolved the style of skeptical argument that appeals to sensitivity. And in fact, *I can*—at least to the letter. In its most explicit formulation, the argument relies on the basis for testimonial knowledge being insensitive. But we have found that we can identify a basis that *is* sensitive. So in print, the argument is unsound.

And yet, to many readers I'm sure this victory over skepticism will feel like a cheat. It will feel like a cheat because the basis I'm identifying is, in the first instant, *somebody else's*. As the audience to somebody else's testimony, it seems strange to say that their basis can be part of my own. This basis for my knowledge would be removed from me in both time and space. So how can it serve as a basis for *my knowledge*?

What we are confronting here is the deep commitment to *externalism* that is the hallmark of the division of epistemic labour theory. To those steeped in the individualist approach, this externalism can be jarring. When it comes to confronting skeptical arguments, this externalism is a feature, not a bug. Externalism offers us a way out of the skeptical argument. But still, there are profound worries that this will also render the view useless when it comes to more practical matters. This is the final issue which must be addressed.

5 Information bubbles and the norms of belief

Here are the lessons that we have learned so far (assuming, of course, that my view is right). First, *if* one's informant knows that *p* and asserts that *p*, then it is possible to know that *p* just by believing their assertion. Moreover, this knowledge does not rest solely on how things appear to the recipient. Rather,

it depends on the informant's original basis for their knowledge. Secondly, if the informant doesn't know what they assert, then the audience cannot gain knowledge by believing their assertion. Thus, we can gain knowledge from testimony if and only if our informant knows what they're talking about. This might seem vaguely truistic, but it's not entirely. It isn't a mere platitude precisely because it rejects wholesale skepticism towards testimony.

It is now time to address a related, but apparently different matter: the question of what we should believe. When we're inquiring, gathering evidence, and deciding what to believe, what are our obligations? What are the epistemic norms or rules that we ought to follow? It might seem as if we have hardly touched on this question so far. I have said when a belief counts as knowledge, but I haven't said anything about when you should believe a testifier and when not to. It might seem as if my view doesn't say anything about who you can trust and who you cannot trust.

But contrary to first appearances, my view does have a take on our epistemic obligations. Let's assume, going forward, that we should believe something if and only if believing it would constitute knowledge.⁷ Hence, our epistemic obligation is to obtain knowledge when possible. In that case, my view does say something about who you should trust: it says that you should trust people in all and only those occasions in which they assert their knowledge. If someone genuinely knows what they say, then you should believe them. But if their testimony is based on ignorance, then you shouldn't believe them.

Some folks object to this idea because they claim that it can't be faithfully followed. We are oftentimes not in a position to know whether someone is speaking from knowledge or speaking from ignorance. But when that's the case, then how do we know how to fulfill our epistemic obligation? Not only that, but if we're not in a position to faithfully follow this alleged epistemic obligation, then how can it really be an obligation? If we don't always know which choices it recommends, then isn't there something arbitrary about it? One person can unwittingly follow the obligation and another can unwittingly violate it, and the one person would be lucky and the other one would be unlucky. But isn't that unfair?

There are a number of assumptions being made in this series of objections. To bring them all out and dramatize them, let's think about the concept of an information bubble. Each of us are, by now, all well aware of information bubbles. They are the sort of thing that social media has become very good at creating: they are information environments where their inhabitants only ever receive information from one side of some point of controversy. They are "bubbles" because those who are within them are insulated from opposing points of view. (This doesn't mean that the opposing point of view is never presented in the bubble. It's just that, whenever it is presented, it is always shown in such a negative light that it never gets a chance of being taken seriously.)

⁷This is contentious, but I'm not going to defend it here. The idea that knowledge is the aim of belief, as well as our basic epistemic obligation, is one of the guiding assumptions of this project. I haven't given arguments in its favour, but I've said a great deal to defend it against the charge that it's a bad choice for a goal because it's unattainable.

Let's imagine two hypothetical people, who we will call Jack and Jill. Imagine that they each grew up in different areas of the United States. Jack grew up as the son of two climatologists, and has been surrounded with information about climate science for all of his life. Nobody seriously doubts the reality of climate change within his community. Perhaps Jack has been made aware of the claims of the skeptics, but he has also been taught rebuttals to these claims which everyone informs him are definitive. The people in his community all seem clearly trustworthy to Jack and he has no reason to doubt their expertise or sincerity. So on the basis of all of their testimony, he believes in the reality of climate change.

Jill, on the other hand, grew up in a community of climate change skeptics. Her family's life's work is to dispute the claims of climatologists. Nobody that Jill is acquainted with actually believes in climate change, and she is constantly surrounded by a wealth of information that purports to prove that climate change is a hoax. She is aware of the mainstream opinions about climate science, but the members of her community have provided her with a list of so-called 'facts' that debunk each item of mainstream opinion. Everyone in her community *seems* to be knowledgeable about what they're saying; they seem to be genuine experts and they seem sincere. So from Jill's point of view, she sees no compelling reason to doubt them. She thus takes them at their word and forms the belief that climate change is unreal.

Both Jack and Jill are living in their own respective information bubbles. Within each of their contexts, their views about climate change will be fairly insular. Jack doesn't really have much of a chance of being exposed to anything that will change his mind, and neither does Jill.

So between Jack and Jill, who is doing the right thing, epistemically speaking? We can fill in the details in such a way so that, from their respective points of view, it will seem as though the reasons to trust in the members of their own communities will appear to be equally strong to each of them. Just as Jack's community appears credible to him, Jill's community appears credible to her. Given this parallel, it's quite natural to think that their epistemic obligations should also be parallel. That is, if Jack should believe in climate change given the information he has from his community, then Jill should also disbelieve in climate change. It might appear as though Jill has no stronger reason to abstain from believing than Jack does. (Remember, Jill has never been presented with evidence for climate change in a credible light. Whenever she's been shown the evidence for climate change, it is always followed by counter-evidence from people whom she trusts.) They thus seem to be in ostensibly analogous situations.

But from the perspective of the theory that I'm developing, this couldn't be further from the truth. *Normatively* speaking, their situations couldn't be further apart. One of our characters, Jack, has a wealth of *knowledge* available to him through his community. The other character, Jill, isn't surrounded with such knowledge. Unfortunately for her, she's surrounded by ignorant informants. So again, what should each of our characters believe? *Answer:* Jack *ought* to believe the members of his community about climate change, because doing so will result in him achieving knowledge. In doing so, he'd be doing the

epistemically right thing. Jill, on other hand, *shouldn't* believe her community. That is because, by believing what she's told, she's increasing her ignorance. Even in her situation, she would be doing the epistemically wrong thing to believe what has been told to her, all of the apparent evidence notwithstanding.

Yet, it's not hard to feel that Jill has been the victim of bad luck here. It's not her fault that she was raised in a community of deceptive and ignorant people. We might think that by listening to her trusted community members when they appear to be credible, she's 'making the best of a bad situation.' She may be *trying* to be epistemically responsible, and trying to do everything right on her end, but through no fault of her own, she's failing because of circumstances that are not in her immediate control.

Some philosophers take this line to show that knowledge can't really be the standard for the real obligations for our beliefs. Rather, it must be something that is more 'within our control.' After all, if we can't *blame* Jill for believing in climate change, then doesn't that show that she hasn't really done something wrong? If she's faultless, then doesn't that mean she hasn't failed her obligation?

There is both a simple answer and a complicated answer to this worry. I am going to leave the complicated answer to a footnote and provide only the simple one.⁸ The proper thing to say is that we must distinguish between *fulfilling an obligation* and *having an excuse*.⁹ Sometimes it happens that people fail their obligations even though it wasn't really *their fault*, and hence we shouldn't heap blame upon them. But that *only* means that we shouldn't blame them; it doesn't mean that they did the right thing. Being *not at fault* is simply not the same as *doing the right thing*.

Here's a mundane example to make this point obvious. In my country, the law says that we're not allowed to drive over 30 km/h in a school zone. This is a legal obligation that we're bound to by law. Now suppose that I'm driving in a car with an inaccurate speedometer, so that it shows that I'm driving 30 km/h when I'm actually driving 40 km/h. Also suppose that I have no other way of knowing how fast I'm going (40 km/h doesn't feel any different from 30 km/h

⁸The complicated answer is best put to the philosopher who says that our epistemic obligations should be such that we should always be in a position to know when we're following them. According to their thinking, what you should believe should depend solely on what you can be perfectly aware of. (Otherwise, there will be cases where you ought to believe something, fail to believe it, and also fail to realise that you're failing your obligation.) The reply to this comes from Timothy Williamson's anti-luminosity argument discussed in the conclusion of chapter 3. Basically, using fairly mild assumptions, we can prove that there are *no non-trivial* states that are perfectly transparent to a subject (states such that they'd be in a position to know that they obtain whenever they obtain). Recall my proof about the appearances: there are no such things as "appearances" that always have the properties that they appear to have. Once we accept this result, we must also accept that there cannot be epistemic obligations that are "perfectly luminous" either. There can't be a norm that says that you're obligated to believe something if and only if N, and you're always perfectly aware whether or not N. The desire to find norms and obligations with that magical epistemic property is a pipe dream. The sooner we realise this, the sooner we can accept that our real obligations won't always make it apparent to us what we should do. In that event, we need to make a distinction between obligation-failure and blameworthiness as I do here, and we no longer have a reason to fault the view of this chapter.

⁹Srinivasan (2015) does a great job of explaining the points made in this section.

to me) and that I have no idea that the speedometer is broken. Did I in fact break the law? Answer: *yes of course I did*. I was going over the speed limit, after all. I still have to pay the repercussions if I get caught. “My speedometer was broken” is not a legally valid defence. The law states that *we ought not to go over 30 km/h*; it doesn’t say *you ought not to let your speedometer go over 30 km/h*.

Having said that, the fact that I broke the law does not, in this case, show that I have an irresponsible predilection to ignoring the law. After all, I thought that I was following my legal obligation, even though I wasn’t. It was just a case of bad luck. For this reason, you shouldn’t regard me as an apt target for scorn or blame for going over 30 km/h. It would be inappropriate for anyone to hold this against me. But still, just because I shouldn’t be blamed doesn’t mean that I followed the law.

Although this example is fairly simple, the lesson generalizes to all other sorts of obligations. *Seeming to do the right thing* should not be confused with *doing the right thing*. Conversely, a person can do the wrong thing while it seems to them that they’re following their obligations. In our earlier example, it might seem to Jill that she ought to believe the testimony of her community about climate change. It might appear to her that by believing her trusted informants, that she would be gaining knowledge, and hence doing the epistemically right thing. She might even have belief-forming habits that are ordinarily typical of good epistemic agents. But none of this would imply that her beliefs are any *good*. In my view, a belief can never be good if it is false.¹⁰ A false belief will always be an epistemic failure, even when it is produced by the most careful and diligent mind. Again, this doesn’t mean that we must blame Jill for her bad belief; it just means that she hasn’t believed what she ought to. The point is, we shouldn’t confuse her blamelessness as a sort of epistemic success.

There is thus a relevant respect in which belonging to an epistemic community, or information bubble, where your peers are unreliable is like driving in a car with an inaccurate speedometer. From the inside, it might appear that you’re doing the responsible thing (collecting knowledge; driving by the speed limit). But your real obligation isn’t just to do what is *apparently* responsible; it is to do what is *actually* right.

All of this raises a very interesting question for our own situations. We know that information bubbles are prevalent, and so each of us likely belongs to many of them. So given that you *do* belong to several information bubbles, what should you do about it? Are you obligated to broaden your horizons and seek out possible counter-evidence to the insular narratives that you’ve been provided? Or is it best to trust the traditions that you’ve been brought up with? The answer to this, when given in broad outline, is bound to be underwhelming.

If you’re in an information bubble that passes around knowledge, then your epistemic obligations are easy. All you have to do is believe this knowledge and you’d be doing the right thing. There is little or no obligation to seek out

¹⁰One can also have true beliefs that are failures of epistemic obligation, when those beliefs aren’t knowledge.

possible counter-evidence. On the other hand, if you're in a bubble that passes around ignorance, then your obligations are demanding. In that case, the best thing that you can do is question the dominant narrative and find your way out of the bubble. For although you may not always be blameworthy for trusting your peers, it would still be a failure of what you should believe. The reason that you should question them is because what they say is *wrong*.

Therefore, which obligations you have and whether they are demanding will depend on what kind of peer group you are actually in—whether they are knowledgeable or not. Of course, all of this begs the higher-order question of *how do you know which kind of bubble you're in?* I don't raise this question to suggest that you need to have this knowledge in order to learn from your peer group if they are reliable. But still, wouldn't it be good to have some strategies for knowing whom to trust?

6 Armchair philosophy and applied epistemology

Admittedly, all of this has the air of platitude. (“How do I get rich?”; “Buy low and sell high”, a useless stock market advisor would say.) It certainly may not seem as though this is the kind of lesson that can directly bear on our practical strategies for belief formation. But if this is the kind of worry that bothers you, then I have one final point to make. In order to make this point, we need to unravel the various aims of epistemology to show why this isn't a fault of the externalists theses (e.g. DOEL) *per se*. Once we do this, we can then go on and identify the sense in which externalism is a practical doctrine.

Let's consider an analogy to ethical inquiry. Whereas epistemology is the branch of philosophy that covers what we ought to *believe*, ethics is the branch of philosophy that covers what we ought to *do*. Yet, there are various levels of questions that we can ask about ethics. Generally speaking, the questions have organized themselves into three distinct levels: practical ethics, normative ethics, and metaethics. Here is a brief rundown of their distinct aims:

- I Practical ethics is the branch of ethics that is concerned with what we should do in everyday circumstances. The sort of questions that qualify as practical ethics should be familiar: is it permissible to legalize recreational drugs?, is it permissible to eat meat? Etc.
- II Normative ethics is the branch of ethical inquiry that is no longer chiefly concerned with what we should do in specific circumstances, but rather with identifying the most basic ethical principles, and how the system of ethical principles is structured. The paradigmatic dispute in normative ethics is the age-old controversy between consequentialism and deontology. (Ultimately speaking, should we do what has the best outcome for people overall or take care to protect individual rights?) This is a debate over identifying the most basic and fundamental principle of ethics. In

addition, there are further debates over the structure of ethical principles. For instance, inquiry into normative ethics will concern the abstract relationships between the notions of *value*, *obligation*, and *rights*.

- III Roughly speaking, metaethics is the investigation into the *nature* of the ethical properties and principles. The guiding questions are: *what is the nature of goodness?*, and *what is the nature of ethical obligation?* Sometimes metaethicists are concerned with identifying these properties (reducing them to other properties), and sometimes they are concerned with identifying the states that these properties are grounded in. For instance, they want to know whether ethical properties are *natural*, or *non-natural properties*. They want to know whether they are determined by societal conventions, or human flourishing, or God's commands. Lastly, they would want to know whether ethical properties are real, or mere fictions.

I want to now make a point that *hopefully* shouldn't be too controversial: that the three tiers of ethical inquiry can—and should—operate with a fairly high degree of independence from one another. What I mean is that we shouldn't expect a given answer in a 'higher' tier to provide any straightforward solution to the questions of 'lower' tier. Whether or not *goodness* is a natural or non-natural property will not provide any straightforward answer to the question of whether consequentialism or deontology is correct.¹¹ It certainly would offer no hint as to whether it's okay to eat meat.

There isn't even a straightforward route from the answers to normative ethical questions to the answers to the questions of practical ethics. Certainly, whether consequentialism or deontology is correct will have some effect on the proper assessment of practical ethical questions, but it will leave much undecided. It will determine what *sorts of reasons* are appropriate for ethically assessing a practical situation, but it won't tell us what the all-things-considered assessments of particular situations should be. That is because each practical ethical problem has its own peculiarities that are not straightforwardly addressed by the fundamental ethical principles; a lot of additional reasoning is required in order to apply the fundamental principles to a practical situation. And typically, what makes an ethical dilemma difficult is not going to be a matter of what's fundamentally valuable, but rather how to assign the correct values to the specifics of the situation at hand.

Imagine a philosophy hermit who spends his entire life on a mountain top, isolated from the rest of society, meditating on the nature of ethics. After years of isolation, he achieves ethical clarity. He discovers, with absolute certainty, that goodness is a natural property and that consequentialism is true. He then descends from his mountain to share his wisdom with the rest of society: that

¹¹One might think that a naturalist metaethical outlook might favour consequentialism since the most obvious way to identify an ethical property with a natural property is to take *goodness* to be *promotion of sentient flourishing*. That might be right. I'm not claiming that the three tiers of inquiry are completely autonomous. Indeed, I don't think they are. My only point now is to warn the reader against expecting straightforward solutions to one tier of inquiry from an answer in another.

we should promote the greatest flourishing for the greatest amount of beings and reduce the most suffering for the most amount of beings. Thanks to the hermit's impeccable armchair reasoning, this principle is now *known*.

How likely is it that this hermit will be able to offer perceptive and nuanced solutions to the contemporary problems of society? Answer: *not very likely*. The hermit's isolation from society has made him woefully uninformed about the complexities of each issue—of the conflicting interests of the relevant parties, the histories that gave rise to their conflicts, etc. Without this worldly knowledge the hermit's insight into normative ethics will make him relatively unequipped to give much practical advice.

At this juncture, it is useful to distinguish *normative ethical principles* from *practical decision procedures*. Normative ethical principles tell us what is *fundamentally good*, or *what our fundamental obligation is*. A *practical decision procedure* is a set of 'rules of thumb' or *heuristics* for deciding how to act in a given situation. For instance, consequentialism is the normative ethical principle that tells us that we ought to do the action that maximizes good consequences. Accepting this view of obligation does not straightforwardly tell us how to decide what to do in a given practical decision. That is because we will not typically be in a position to know what all of the consequences of our actions will be, or how to assess their value, and so we still need some heuristics to decide what to do given our limited information. These heuristics will typically have to take into account a lot of social facts that cannot be learned from the armchair (or the top of a mountain). Consequentialism does not offer any straightforward answer as to what those heuristics should be. It can weed out certain *bad* heuristics that have no place given consequentialism, but it does not determine what the best heuristic will be. (It's even consistent for a committed utilitarian to resort to deontic reasoning when they have to make quick decisions in real life, because they think that treating others as an end is typically a good way to maximize utility, and they can't foresee all of the potential consequences of their actions.) The limitation of the hermit's exercise in armchair philosophy is that he only knows that consequentialism is true, but he's in no position to offer any good advice as to the practical decision procedures we should adopt given the complexities of the society we live in.

With that said, it's important to see that armchair philosophy isn't completely useless either. The reason for this is because there may be (and are) many members of society who are tempted into bad ideas in practical ethics because they are in the grips of a false normative or metaethical theory. For instance, there may be ethical egoists who claim, as a matter of normative principle, that we should only ever look out for our own good. As a result, they promote public policies that selfishly favour their own ends at the expense of everyone else. Or worse, there may be ethical nihilists who claim that there's no such thing as right or wrong, and so they selfishly pursue their own interests while dismissing all ethical reasons to do the contrary. Even worse than this, there may be whole groups of bad actors who weaponize radically skeptical or nihilistic views on the normative or metaethical level in order to selfishly promote oppressive ethical practices. (Consider those religious zealots who claim

that there would be no morality without God’s commands in the form of their holy text, and then use these beliefs to justify a practice of oppressing women and LGBTQ+ people. Or consider the ethno-nationalists who use a sort of ‘social darwinism’ or cultural relativism to justify their racist ethical practices.) As I see it, the most important practical value of the hermit’s mountaintop reflections is that he can serve to debunk these bad ideologies. (Granted, if these bad actors are irrational, then his refutations won’t change their minds. But I assume that even rational people can sometimes be tempted by arguments for false normative or metaethical theories, and the philosopher’s role is to save them from temptation.)

I claimed that the three tiers of ethical investigation are “semi-autonomous.” I can now say what I mean by that. Tier N of inquiry is semi-autonomous from tier N+1 if and only if an answer from Tier N+1 may be able to weed out and debunk some radically *bad* answers at Tier N (in particular, those that are motivated by radically skeptical or nihilistic tier N+1 ideas), but it will still leave it underdetermined as to what’s the best answer to a tier N question.

Now let’s return to epistemology. Recall that in §1.1 and §2.5.2, I outlined three tiers of epistemological inquiry. Here they are again.

E.I *Applied epistemology*. This branch of inquiry is concerned with the practical questions about what we ought to believe given the specific evidential situations we live in. For instance, should I believe the reports of *Fox News* or *The Atlantic*? Should I trust the traditionally accredited environmental scientists with what they have to say about global warming, or should I believe the studies produced by heterodox climate change deniers? When a woman accuses a powerful man of sexual misconduct, and he denies it, and there’s only one person’s word against another, whom should we believe? There are also more challenging questions about the proper methods of scientific inquiry (witness the controversies over the replication crisis in psychology and the social sciences).

E.II *Normative epistemology*. Just as in ethics, there’s a tier of inquiry into the fundamental normative principles of epistemology. These are the basic norms that govern what we ought to believe, and dictate how *knowledge*, *justification*, and *doxastic obligation* are structurally related to each other. For instance, some philosophers claim that we ought to believe what is *practically useful* because knowledge is unattainable; I have claimed—and defended—that we ought to pursue *knowledge* because it *is* attainable (chapters 1 and 2). Some philosophers claim that we ought to structure our beliefs into a coherent web of justification; I have claimed that knowledge has a foundationalist structure (chapter 1). All of these inquiries and debates sit at the same tier as what’s analogous to normative ethics in ethical inquiry. That is, they remain semi-neutral to the metaphysics of epistemic properties, and they do not offer straightforward practical advice as to what to believe in real-life situations.¹²

¹²Although as I concluded in §2.5.2, there are occasions in which meta-epistemological

E.III *Meta-epistemology*. This is the investigation into the *metaphysics* of the various epistemic properties—namely, knowledge, justification, and epistemic obligation. Epistemologists interested in this inquiry want to know what the *nature* of these properties are (the epistemological literature in the late 20th century was replete with attempts to analyze these properties into more basic properties). Moreover, one interested in the metaphysics of these properties will also be interested in the sort of states of affairs that determine them. In particular, one might be interested in whether they supervene on internal states of the subject, or not. Hence, the *internalism vs. externalism* issue is best seen under the heading of ‘meta-epistemology’ or ‘the metaphysics of epistemic properties.’

Having laid out these divisions with the analogy to ethics in mind, I think the lessons straightforwardly carry over.

Lesson 1: practical decision procedures for belief are underdetermined by meta-epistemology and normative epistemology.

So who should you believe, *Fox News* or *The Atlantic*? Answer: certainly not Fox News. But mind you, I didn’t reach this answer by attending only to the norm that says *believe that which is knowledge*. Presumably, the audience of Fox News will say that they agree with that norm, at least to the letter. Rather, I reached this conclusion because I also happen to know that Fox News fares pretty poorly whenever it is fact checked by independent third-party media watch dogs. This knowledge doesn’t come from reflections on philosophy; it comes from being informed about the reliability of the various media outlets. Of course, the advocates of Fox News might disagree with me about this assessment because they’re also suspicious of the third-party media watch dogs. They think that they’re biased and unreliable too. (As I said, they may believe themselves to be adhering to the knowledge norm of belief to the letter, they’re just wrong about the non-philosophical facts about which news sources are reliable.)

For the most part, *practical problems call for practical solutions*. That’s practically a truism. But it’s a truism that we need to keep in mind when we’re assessing the tiers of epistemological theorizing that are a step or two removed from the practical problems that concern us.

There’s a certain objection to the DOEL view of testimonial knowledge and other externalist accounts like it that claims that they miss the point of epistemology because they don’t provide useful guides as to what to believe. My response to this point is that *of course* externalism is going to underdetermine the answers to most practical decisions as to what to believe. It was never intended to answer those. (Nor was internalism! Internalism merely says that justification supervenes on the internal. Does that help you decide whether to trust *Fox News* or *The Atlantic* without further information? Of course not.) Externalism is a thesis that concerns the *metaphysics* of epistemological

doctrines (i.e. externalism) can aid us in defending a normative epistemological conclusion (that knowledge is attainable).

properties. To complain that it fails to provide solutions to applied epistemology is akin to complaining that metaethical naturalism (by itself) fails to deliver an answer as to whether we should eat meat. A person who levels this complaint is seriously misunderstanding the different aims of the different tiers of inquiry.

Lesson 2: Externalist theories are practical insofar as they can debunk harmful overly-skeptical ideas.

Am I conceding that externalism is a practically useless doctrine? No. Not at all. My only claim is that externalism will *underdetermine* the solutions to practical problems without the help of additional societal, scientific (i.e. non-philosophical!) facts as to what are the reliable sources of testimony and scientific knowledge. But that's exactly what one should expect from any doctrine whose aim is to answer questions about the metaphysics of epistemological properties. Philosophy isn't the answer to all problems, and nor should you expect it to be.

But remember what I said about the practical utility of foundational doctrines in ethics—what good the hermit can do. He may not be able to solve practical problems, but he can warn us against *bad answers that are rooted in bad foundational doctrines*. This is exactly what I claim to be practical utility of externalism.

Earlier we observed how bad epistemology can be put to great mischief. Overly-skeptical patterns of reason are routinely weaponized to protect the misdeeds of powerful people, flood the public sphere with propaganda, and subdue people into inaction. Moreover, each of us is vulnerable to such patterns of reasoning, because we each have an irrational tendency to selectively sympathize with some of these skeptical arguments when it suits us.

Now let's return to the parable of the hermit. Imagine now an epistemological hermit who spends their entire life on a different mountaintop studying Tier-2 and Tier-3 epistemological questions, in relative ignorance of the facts of society. Reflections on skepticism, the nature of knowledge, the norms of belief, etc., convinces them, *with impeccable reasoning and absolute certainty*, that (i) knowledge and justification are dependant on environmental factors, that (ii) knowledge is the norm of belief, (iii) and that the brain-in-a-vat-style skeptical arguments fail because they falsely assume a far too 'internalist' criteria for our evidence. (In other words, they reach the conclusions of chapters 1 & 2.) They then descend from their mountaintop into society. What wisdom would they have to offer us?

For lack of any practical knowledge of our society, they wouldn't be able to tell us straight-away whether to trust *Fox News* or *The Atlantic* (they've never even heard of *Fox News* or *The Atlantic*). They wouldn't even be able to tell us whether to trust the climate scientists or the climate change deniers—at least, not until they learn some more stuff about climate science, the history of denialism, and all these other facts that can't be learned from armchair philosophy. But just because they can't help us with *that*, it doesn't follow that they're useless. Indeed, they would be more perceptive than the rest of us of

the dangers and fallacies of skeptical patterns of reasoning. They see the bad actors wield skeptical arguments to try to silence their opponents, and many of us regular folk get duped. Armchair philosophy's utility for society is that it can show us, in a philosophically satisfying way, how to resist these dangerous patterns of thinking.

7 Conclusion

There has been a long history of philosophers who have been tempted by the following web of ideas. (No philosopher exhibits this web better than Descartes.) First, they wanted to achieve knowledge with certainty, and they further wanted a certain internal mark that would guarantee when it's achieved. To this end, they would attend to the world of appearances as the starting place for epistemology. Starting in this sphere was supposed to be safe and certain; and then perhaps, from there, they could articulate the rules for extending their knowledge outward on the basis of how things appear. The hope was that these rules would be impeccable guides to extending one's knowledge.

Having laid out this groundwork, these epistemologists would arrive at several striking conclusions. For one, it would follow that we couldn't base our knowledge on anything unless it was shareable between ourselves and our brain-in-a-vat counterparts. Since they wanted an internal guarantee that their knowledge was well-founded, they wouldn't trust anything if it was possible for the appearances to trick them. For another striking conclusion, it would follow that we couldn't learn from other people by taking their testimony at their word. Instead, we would have to treat the testimony of others with a default stance of agnosticism, until we could back up their credibility with evidence produced as an individual. In practice, this means barely trusting anyone at all (if ever).

Far from being an antidote to skepticism, this whole approach to epistemology ends us being the skeptic's partner in crime. By turning inward to look for certainty, one misses out on the wealth of knowledge that can be received from outside. By insisting that knowledge must have internal marks, one loses sight of the knowledge that's actually attainable.

At various points throughout this project, I have advocated to replace this picture with an externalist antidote to skepticism. Broadly speaking, this means that we can attain knowledge when our beliefs are the products of channels that relate our cognition to the facts. These channels include perception in the most basic case, and then memory and testimony as modes of knowledge preservation. This externalism also means that knowledge isn't solely a function of internal 'markings'. It is not solely based on how things appear, and in the testimony case, it is not based on the audience's assessment of their informant.

Both historically and in my own presentation, the philosophical purpose of these externalist theses was never so much concerned with delivering practical lessons that we could carry with us outside of the philosophy classroom. The main purpose of these ideas was always remote and theoretical: it was primarily to combat radical philosophical skepticism—i.e. the claims that we can't know

anything at all or anything about the external world. And although that is not a valueless goal, it is one that only a philosopher would care about.

But in this chapter we have reached a lesson that I think is rather profound, and reaches far beyond theoretical philosophy.¹³ As I've been at pains to stress, skeptical reasoning is pervasive. So any theoretical tools that we have to grapple with this reasoning has application in real life circumstances. And these are especially handy tools to have when skeptical reasoning is put to ill ends.

So here is my final word on the practical utility of externalism. Its opposite—internalism—has a tendency to cultivate a skeptical frame of mind that promotes suspension of judgment. In itself, that is not necessarily a bad thing. But it can become a *very bad thing* when it leads one to miss out on important items of knowledge—especially when the items of knowledge that one misses are vital to making practical decisions. My externalist epistemology, on the other hand, sets this right. On the abstract normative level (tier II), it shows how knowledge is possible and is hence an acceptable aim for belief. We thus avoid the stultifying effects of internalism and skepticism. This alone does not recommend any specific belief-forming strategies that could be used in the day-to-day quandaries we face about what to believe. Thus it does not necessarily lead to being too hesitant or too gullible. But this is as it should be. As I said, it's not solely the job of normative epistemology to guide you in your beliefs; that's the job of being an informed citizen. With that being said, these theories do give us the means to shield off *bad ideas* that are rooted in *bad epistemological principles*. And that, I think, is the most we can ask for, as far as philosophy is concerned.

¹³All of the credit to Jonathan Jenkins Ichikawa, who makes this lesson a theme for his undergraduate philosophy courses. (I was a teaching assistant for one of them.)