

# Social Knowledge and Social Norms

Peter J. Graham

## 1. Introduction

Many philosophers in the twentieth century argued that knowledge was social, for belief was social, for thought was impossible without language. A weaker view was also popular: although perhaps thought does not depend on other minds, knowledge does, for knowledge requires the ability to justify one's belief to others according to social standards of justified belief. By the end of the twentieth century, however, both views fell on hard times. The current wisdom holds that neither belief nor knowledge is essentially social.

Nevertheless, a good deal (if not most) of what we know we know because we believe what other people tell us. The extent of our knowledge owes a great deal to what we learn from others. Epistemology would be far from complete if it did not thoroughly address testimony. Although discussed during the Modern Period, it was not until the late twentieth century that testimony came to the fore.<sup>1</sup> In this chapter I focus on testimony as a central source of knowledge.<sup>2</sup>

## 2. Justification versus knowledge (the anti-luck condition)

I begin with the distinction between justification and the anti-Gettier or 'anti-luck' condition on propositional knowledge. I will rely on this distinction to frame our discussion.

On the traditional analysis of propositional knowledge, a subject's knowledge that P is comprised of (1) S's belief that P, (2) the fact that P (so that S's belief that P is true), and (3) S's having a justification for believing that P. Gettier (1963) showed that there are cases of justified true belief that P that fall short

of knowledge. Here is an example from Bertrand Russell (1948): an individual may look at a clock and as a result truly believe, with justification, that the time is 2 p.m. But what if the clock is stopped and the individual only looked at just the 'right' time? Then the individual does not know it is 2 p.m. (despite having a justified true belief): you can't learn the time by looking at a stopped clock. The 'Gettier problem' led most epistemologists to conclude that propositional knowledge includes an additional 'anti-luck' condition.<sup>3</sup> Let's then assume that justification is one thing and the anti-luck condition is another, for you can meet one without meeting the other.

### 3. What is testimony?

This is a question not about the speech act per se – namely, what is testimony as one speech act among others? – but rather about the *process* that leads to so-called testimony-based beliefs. The question 'What is testimony?' (for our purposes) is the question 'What makes a belief a testimony-based belief?' Testimonial knowledge is then testimony-based belief that meets the conditions for knowledge: testimonial knowledge is justified true testimony-based belief that satisfies the fourth, anti-luck condition, just as perceptual knowledge is justified true perceptual-belief that satisfied the fourth, anti-luck condition.

Here are three easy examples of testimony-based belief. You want to know the time, so you ask a passer-by. '2 p.m.' she says. You believe her. You want to know the best time to visit Shanghai, so you do a web search. The top hits tell you that the second week of October has the best weather and the fewest tourists. You decide that you want to know a little chemistry, so you buy a textbook. When you read about the elements, you form a series of testimony-based beliefs.

These examples have the following three ingredients in common.

- (1) Someone's testimony that P is the *distal cause*. A belief is testimony-based because it is causally based or sustained (in part) on someone else's testimony, where 'someone's testimony that P' includes assertions, sayings, tellings, even assertive conversational implicatures.
- (2) A hearer's representation as of a speaker's testimony that P is (part of) the *proximal cause*. The hearer must represent the speaker's speech act as a telling or saying that P – the hearer must comprehend the speaker's testimony as a part of the process of forming a testimony-based belief. No comprehension *as of* a speaker asserting that P as a normal part of the formation of a testimony-based belief, no testimony-based belief that P.

- (3) The psychological transition from comprehension to belief involves *deference* or *epistemic dependence* to the epistemic authority of the speaker. The hearer is ‘turning to’ the speaker (to the passer-by, to informants on the internet, to chemistry as a science) for information (for knowledge), and deferring to, or depending upon, the speaker (the sender) for the facts the hearer wants to know. Testimony-based beliefs are ‘second-hand’ beliefs, beliefs formed through deference to, or dependence on, the word – the testimony – of another.<sup>4</sup>

Given our distinction between justification and the anti-luck condition, I divide the epistemology of testimony into two questions:

- (i) What is testimonial justification? When and why are testimony-based beliefs justified?
- (ii) What is testimonial knowledge? When a hearer forms a justified true testimony-based belief, when and why does the hearer satisfy the anti-luck condition?<sup>5</sup>

But before discussing those questions, we should first decide whether testimonial knowledge is even possible.

#### 4. The possibility of knowledge transmission

John Locke is frequently quoted as claiming that testimonial knowledge is not even possible:

For, I think, we may as rationally hope to see with other men’s eyes as to know by other men’s understandings. So much as we ourselves consider and comprehend of truth and reason, so much we possess of real and true knowledge. The floating of other men’s opinions in our brains makes us not one jot the more knowing, thought they happen to be true ... In the sciences, every one has so much as he really knows and comprehends. What he believes only, and takes upon trust, are but shreds; which, however well in the whole piece, make no considerable addition to his stock who gathers them. Such borrowed wealth, like fairy money, though it were gold in the hand from which he received it, will be but leaves and dust when it comes to use. (*Essay Concerning Human Understanding*, Bk I, ch. III, sec. 24)

People read this as denying the very possibility of testimonial knowledge, of the transmission of knowledge – gold – from one mind to another; all that results from deference and dependence is shreds and dust.

This strikes nearly everyone as incredible; we are deeply social creatures after all, living lives of massive mutual interdependence. Surely we learn (come to know) enormous amounts from relying on others. Locke's British intellectual heir, David Hume, famously remarked that 'there is no species of reasoning more common, more useful, and even necessary to human life, than that which is derived from the testimony of men, and the reports of eye-witnesses and spectators' (*Inquiry Concerning Human Understanding*, sec. 10, para. 5). As opposed to Locke and in line with Hume, common sense, it seems, asserts that knowledge from one mind to another can be, and often is, transmitted through testimony. C. A. J. Coady (1992: 389) agrees: 'If S knows that P ... then S can bring his listeners to know that P by telling them that P'. Elizabeth Fricker (1987: 57) claims: 'If S's belief is knowledge, then we may allow that title to H's belief too.' And Tyler Burge (1993: 477) asserts: 'If one has acquired one's belief from others in the normal way, and if the others know the proposition, then one acquires knowledge.' A cursory review of the literature on this question will result in a score of similar passages. Indeed, the agreement is so widespread on this issue that it must form a piece of our folk epistemology.

Although it goes without saying, these passages require qualification: the knowledge is transferred only if the hearer is also justified in forming the corresponding testimony-based belief. If someone knows that P and tells you that P, but you have every reason to believe that they are either lying or mistaken, then it is far from obvious, and probably false, that you would come to know that P by believing them.

Is the speaker's knowledge that P also necessary for a hearer's testimonial knowledge that P? The answer here, too, seems to be in the affirmative: a hearer can acquire testimonial knowledge that P only if the speaker has knowledge that P to transmit. If the speaker doesn't have the knowledge to transmit, how could the hearer, in depending on the speaker, acquire knowledge? Magic?

Commentators agree. Robert Audi (1997: 410) says: 'My testimony cannot give you testimonially grounded knowledge that P without my knowing that P'. Tyler Burge (1993: 486) writes: 'If a recipient depends upon interlocution for knowledge, the recipient's knowledge depends on the source's knowledge as well.' Angus Ross (1986: 62) asserts: 'Your telling me that P can only be said to provide me with knowledge if you know that P'. And Michael Welbourne (1986: 302) says: 'It is necessary, if there is to be a successful process of testimonial transmission, that the speaker have knowledge to communicate.'

Philosophical reflection, in line with common sense, reveals that (*pace* Locke) testimony transfers knowledge from one mind to another.<sup>6</sup>

## 5. Testimonial justification: reductionism versus anti-reductionism

This brings us back to our two questions: what is testimonial justification, and what is testimonial knowledge? In this section I discuss the first question; I discuss the second in the subsequent two sections.

Why are we justified, when we are, in believing another's testimony? There are two opposing perspectives on this question. The so-called anti-reductionist believes that we enjoy a default epistemic right to take the assertions and testimonies of others at face value. If someone tells you that P, you have a default defeasible epistemic right or entitlement to believe that P; testimony-based beliefs are default defeasibly justified. Tyler Burge (1993, 2013) calls this the 'Acceptance Principle':

A person is entitled to accept a proposition that is presented as true and that is intelligible to him or her, unless there are stronger reasons not to.

When Burge advances this principle, he is working within a broadly reliabilist view of justification (or 'warrant'), where entitlement (his word for justification or warrant that does not rely on having a reason or reasons) entails the reliability of the belief-forming competence in normal conditions when functioning normally (Burge 1993, 2013; Graham forthcoming). Very roughly, his idea is that, when you take someone to have told you that P, you enjoy a prima facie defeasible epistemic entitlement to believe what they tell you, for your ability to comprehend others reliably leads to true beliefs in normal conditions when functioning normally: relying on others (in normal conditions) is a good route to truth.

Thomas Reid advanced a similar view in the eighteenth century. In section 24 of chapter 6 of his *Inquiry into the Human Mind* (1764), he argued that there were two principles of human social psychology that tallied with one another – the principles of veracity and of credulity.

The wise and beneficent Author of Nature, who intended that we should be social creatures, and that we should receive the greatest and most important part of our knowledge by the information of others, hath, for these purposes, implanted in our natures two principles that tally ... The first ... is a propensity to speak truth ... so as to convey our real sentiments. This principle has a powerful operation, even in the greatest liars; for where they lie once, they speak truth a hundred times. Truth is always uppermost, and is the natural issue of the mind. It requires no art or training, no inducement or temptation, but only

that we yield to a natural impulse. Lying, on the contrary, is doing violence to our nature; and is never practised, even by the worst men, without some temptation. . . . Another original principle implanted in us by the Supreme Being, is a disposition to confide in the veracity of others, and to believe what they tell us . . . [T]he former . . . may be called the principle of veracity [and the latter] we shall . . . call this the principle of credulity. It is unlimited in children, until they meet with instances of deceit and falsehood: and it retains a very considerable degree of strength through life.

The result of these paired principles of human psychology is that true information, for the most part, flows reliably from one mind to another. Or so Reid believes.

Although Reid attributes this fortunate epistemic situation to the wisdom and good will of the Supreme Being, most contemporary followers of this broadly Reidian view would ground our social psychologies in our evolutionary and cultural history, cultivated in different ways by our varying social norms and traditions. I'll touch again on this issue.

Burge, however, ambitiously argues that the ground of the reliability of assertive communication lies not in our social psychologies, but rather in the very nature of reason. Famously, he argued that the Acceptance Principle is a priori true, for the reliability of assertive communication in normal conditions when functioning normally follows from the very nature of the expression of propositional content as a sign of reason, and the very nature of reason as a sign of truth. Regrettably, space does not allow for discussion of Burge's ambitious argument for this surprising conclusion.<sup>7</sup>

Without assuming that anti-reductionists are 'simple' reliabilists about epistemic justification (generally, they're not), let's assume that behind the spirit of anti-reductionism is the assumption that our default defeasible entitlement involves the reliability of assertive communication in normal conditions.

Not everyone agrees that the Acceptance Principle is true. Those who reject the Principle are often called reductionists. They deny that we can take for granted, in the absence of defeating evidence, that other people are trustworthy. Reductionists believe that, in order to be justified in believing what another tells you, you need independent, non-essentially testimony-based, positive reasons for believing that the person is sincere and competent on the occasion, either because *testimony* is generally trustworthy, or because *they* are generally trustworthy, or because they are apt to be trustworthy *on this occasion*. The reductionist believes another's report that P as such is epistemically neutral on the question whether P; on its own, it provides no support to believe that P.

Their main intuitive motivation seems to be this: you just can't trust other people to be as reliable as you need them to be. There is too much mendacity and honest error, too many sources for error in the chain of communication, to simply risk believing what other people tell you. Indeed, for all that you know, most testimony may be misleading. To avoid being duped, either intentionally or accidentally, you need to sort out the good from the bad, and that requires positive non-testimonial grounds for thinking that your source is sincere or reliable or both. Not doing so is a recipe for objectionable gullibility (Fricker 1994: 144–5).<sup>8</sup>

The structure of this debate between the reductionist and the anti-reductionist parallels an older debate over perceptual justification. The classical foundationalist held that beliefs formed through perception about the external world were not *prima facie* justified (experiences as such were neutral guides to external reality), but instead required for their justification background support from more 'basic' or 'fundamental' beliefs about the characteristic patterns of our own consciously known inner sensations and perceptual experiences. From beliefs about the patterns of inner experience, we would then infer the characteristics and existence of objects in the external world, presumably via an inference to the best explanation. The justification for beliefs about the external world then 'reduced' to the justification for beliefs about our own minds and what we can in turn infer about the external world from those. The classical foundationalist's opponent (often called 'moderate' foundationalism) insisted instead that perceptual beliefs about the external world are justified 'directly' by the perceptual experiences (perceptual representations) that cause them. There was then no need to 'reduce' perceptual justification to any other form of justification; we enjoy default entitlement to take perceptual experience at face value.

You should now see the parallel, and the point of the labels. Reductionists about testimony think that testimonial justification reduces to first-hand justification through perception, stored in memory, and extended through reasoning (for testimony as such is epistemically neutral; the fact that someone told you that P is no right on its own to believe that P), just as reductionists about perception think that perceptual justification reduces to other sources (for perceptual experience as such is epistemically neutral; the fact that it perceptually appears to you as if P is not right on its own to believe that P). Anti-reductionists about testimony think that testimonial justification doesn't so reduce, but instead stands on its own two feet, just as anti-reductionists about perception think that perceptual justification doesn't reduce to other more 'basic' sources of justification.<sup>9</sup>

What is the main motivation for anti-reductionism? Here is the standard argument. First, most of our testimony-based beliefs are justified. Second, as a matter of fact we lack the reductive non-testimonial reasons required to reductively justify the extent of our reliance on testimony. The full extent of the justification for our testimony-based beliefs thus does not ‘reduce’ to a non-testimonial basis (see, especially, Coady [1992]). The anti-reductionist sees the reductionist alternative as overly demanding.<sup>10</sup>

Although many people find this point against the reductionist compelling, it has its detractors.<sup>11</sup> Instead of pausing to evaluate these replies, let me rehearse instead its most plausible instance: childhood testimony. Children form justified testimony-based beliefs, even before the age of two. From two years onwards, they are deeply dependent for information on what other people tell them, especially on topics they cannot observe for themselves (Harris 2012; Harris and Lane 2014). Infants and young children have a strong and adaptive tendency to rely on testimony (Cole, Harris, and Koenig 2012). But children lack some of the conceptual resources to formulate meta-arguments about the reliability of their interlocutors. They also lack sufficient first-hand evidence to justify the premises in those arguments, even if they can formulate them. And, last, even if they have the evidence and can formulate the arguments, it is unlikely that they rely on that evidence and formulate those arguments; they seem to lack the executive ability. Thus, if children have justified testimony-based beliefs, but they cannot ‘reduce’ that justification to justification from other sources, then testimony-based beliefs do not require positive non-testimonial reasons for their justification.<sup>12</sup>

There have been two main responses to this argument in the literature. The first denies that children have justified testimony-based beliefs; since they lack the positive reasons, they lack justified beliefs (e.g. van Cleve 2006). The second denies that children lack the positive reasons – recent developmental psychology has shown us that children are smarter than we think they are. I simply reject the first reply. As for the second, too much time would be needed to review it. Suffice it to say that the recent evidence from developmental psychology, in my view, shows no such thing.<sup>13</sup>

Fricker (1995) offers a unique and telling reply. She grants the case for children but then goes coherentist for adults. That is, an Acceptance Principle is true for justified testimony-based beliefs when the hearer has little to go on by way of background support (the childhood case), but then coherentism is true when the hearer has a coherent set of beliefs to deploy to justify her reliance on

testimony (the adult case). By my lights, this is just to grant anti-reductionism. First, it grants the case where background belief is insufficient. Second, it grants the anti-reductionist argument against the distinctively *reductionist* requirement on background reasons. Third, the existence and relevance of coherent background beliefs is fully compatible with anti-reductionism (Graham 2006c). Fourth, without embracing coherentism, the anti-reductionist can even explain why so many of those background beliefs are justified, and thus relevant to new testimony-based beliefs: those background beliefs were themselves justified by testimony, not by further reductive background beliefs (Graham 2006b; Burge 2013). A main concern of Fricker's, it seems to me, is not the abstract issue of the structure of testimonial justification, but instead the feared laziness or epistemic irresponsibility that might result if a believer took the Acceptance Principle as a license to ignore relevant background beliefs, or to ignore evidence of untrustworthiness on the speaker's part. And that concern is shared equally by the anti-reductionist.

The reductionist thinks that the Acceptance Principle is too permissive; the anti-reductionist thinks that the alternative is just too demanding. Is there a way to settle this debate?<sup>14</sup> Let me offer three points in favour of anti-reductionism. First, as I just remarked, the Acceptance Principle is not a recommendation to ignore counterevidence: it's not a policy or a license to hand over all of your thinking to what other people tell you, willy-nilly. Second (and relatedly), the Acceptance Principle is fully compatible with hearer's possessing, developing, and deploying filters and countermeasures. Various forms of monitoring and epistemic vigilance are not only compatible with the Acceptance Principle (compare analogous filters and monitors for perceptual justification), but are often recommended by anti-reductionists (Goldberg and Henderson 2006; Henderson 2008; Graham 2010; Sperber et al. 2010). Third, reductionist anxiety is nearly always driven by armchair considerations. Although this is not objectionable as such, one might naturally wonder what the empirical evidence suggests. Surprisingly, empirical studies suggest that 'gullibility' isn't such a bad thing, especially if our goal is to promote truth and avoid error (Michaelian 2010; Shieber 2014; Ahlstrom-Vij 2016). Some initial computer modelling reaches a similar conclusion (Zollman 2015). I think these three points, individually and collectively, squarely place the burden of proof on the reductionist. This conclusion should not be surprising, once we recognize not only the *extent* of our reliance on other people for information, but the extent of our *nature* as social beings.

## 6. Counterexamples to transmission

I now turn to our question about testimonial knowledge: when and why does a hearer satisfy the anti-luck condition? You might think that the answer to this question is easy. Since the speaker knows and thereby satisfies the anti-luck condition, the hearer has only to satisfy the justification condition to come to know. The neat thing about testimonial knowledge is that the speaker (in possessing knowledge) does the work for the hearer. Problem solved. Hence, we don't need an independent treatment for the hearer, beyond an account of the hearer's justification. This seems to be the point behind the common-sense view that testimony *transfers* knowledge.

Unfortunately, the issue isn't so easily resolved. Even if the speaker has knowledge, that may not be enough for the hearer to satisfy the anti-luck condition. Here's an example of just that possibility (modified from Nozick [1981]):

HOSPITAL. A father knows his son is fine today, even though his son suffers serious health problems. The father's mother (the son's grandmother) is sick in the hospital. When the father visits, he tells her that her grandson is fine. But if his son were sick (or even dead), he would not tell her, so as not to upset her. Although the father knows his son his well, his mother does not learn from him, for he would easily tell her that her grandson is fine even when that is not so. Relying on his testimony about her grandson's well-being, she would easily form a false belief.<sup>15</sup>

The hearer forms a justified true testimony-based belief, but the hearer doesn't acquire knowledge, for in a sense the hearer's true belief is just luckily true. Although the speaker has knowledge to transmit, and the hearer (we assume) has every right to believe the speaker (and so forms a justified belief), something goes wrong.

Matters get even worse for the simple common-sense answer, for it may be possible for a hearer's justified testimony-based belief to satisfy the anti-luck condition, even though the speaker does not have knowledge to transmit, so it cannot be that the speaker's knowledge that secures the hearer's knowledge. Here is a variant of a much discussed case (Graham 2000b, 2006a, 2016b):

FOSSIL. A devout creationist teaches at a public school where she must teach a section on evolutionary theory. She does not believe a word of it, but is a dedicated and responsible teacher. She develops a near-expert understanding based on deep reading of books and articles on evolutionary science. She

even develops a deep understanding of fossils that parallels highly skilled scientifically trained expertise. On a field trip, she discovers a fossil that proves that ancient humans once lived in this area (itself a surprising discovery that no one previously knew). Although she does not believe it, when she tells this to her students, they believe her. Because of her commitment to teaching, her exposure to evolutionary science, and her mastery of fossils, she would not say what she did unless it were true. Relying on their teacher, the schoolchildren would not easily be mistaken when coming to believe what she tells them. The children come to know something that no one has ever previously known.

As I said, variants have been much discussed (Lackey 1999; Carter and Nickel 2014).<sup>16</sup> Although many people are willing to accept that knowledge transmission does not always succeed (as in HOSPITAL), people find FOSSIL much harder to accept. How can children acquire knowledge from relying on another, if the sender doesn't have knowledge to pass along? That's like borrowing money from someone who doesn't have any to lend. How can someone give you something that they themselves don't have? So, how can the speaker ensure that the hearer meets the anti-luck condition for knowledge, if the speaker doesn't have knowledge herself?

Digging in their heels against the intuitive force of the example, people have tried three different replies. First, people have argued that, even though it looks like the children acquire knowledge, they really don't, for on closer inspection the teacher has a flaw that prevents her from inducing knowledge in her pupils: she doesn't always tell them what she believes (when it comes to evolutionary science, she tells them what the science would say, and not what she believes). But this isn't very persuasive. It would mean that they never learn from their teacher (and that you can never learn anything from someone who isn't completely honest with you about everything). And, after all, isn't it intuitive that, because she relies *on the science*, the children are clearly positioned to acquire knowledge from her (Burge 2013)?

Second, people have argued that, although it looks like the children form testimony-based beliefs, they are really relying on their background beliefs about the reliability of their schoolteachers and their first-hand, independently acquired knowledge of this teacher's track-record for telling the truth. Their beliefs are more 'first-hand' than 'second-hand', and so they are not really deferring to their teacher or epistemically depending on her expertise as opposed to their own, and so they are not really forming *testimony-based* beliefs, and so their knowledge isn't *testimonial* knowledge after all (Audi 2006; Fricker 2006). But this isn't very persuasive. The children are young children; they are not nascent critical reasoners sorting out which teacher to trust and which to

ignore. They are schoolchildren, after all, there to learn from their teachers. If anyone learns from deferring to others, to depending on the expertise of their informants, small children do.

Third, people have argued that, even though the children acquire knowledge from the teacher and form genuinely testimony-based beliefs, they don't acquire *testimonial* knowledge. 'Testimonial knowledge', they say,

is essentially knowledge from another person's knowledge through testimony. If the speaker doesn't have knowledge to transmit, it is *a priori* analytically impossible for the hearer to acquire *testimonial* knowledge; testimonial knowledge is knowledge *secured from someone else's knowledge*. It's just a part of the very concept, or the very idea, of *testimonial* knowledge. It is what the phrase 'testimonial knowledge' *means*.

So, even though the children acquire knowledge through epistemic dependence on their teacher's assertion (they form a testimony-based belief that is knowledge), they don't acquire *testimonial* knowledge.

I've never found this reply very persuasive, but I'm happy to grant it. I shall not, that is, dispute about a word. The children then acquire *testimonial knowledge\** (testimony-based belief that is knowledge) but not *testimonial* knowledge (knowledge that *a priori* analytically entails knowledge from a speaker's knowledge). Although all *testimonial* knowledge is also *testimonial knowledge\**, not all *testimonial knowledge\** is *testimonial* knowledge.

But, having made this linguistic concession, I shall ignore it. For the interesting and worthwhile category of inquiry is not the narrower category of testimonial knowledge, but is instead the broader category of *testimonial knowledge\**. That, I believe, is the category we should strive to understand. I'll give a brief argument for this at the end of the next section.

We still need an answer to our second question: how and why do hearers satisfy the anti-luck condition when forming justified true testimony-based beliefs? The answer from reflection on common sense – that the speaker knows, and so the speaker, in satisfying the anti-luck condition, takes care of for the hearer – didn't work. We need to dig a little deeper.

## 7. A safe-basis account of testimonial knowledge

Maybe we should choose a proposed anti-luck condition, and see where that leads. Since Gettier published his short article in 1963, there has been no

shortage of attempts to articulate the correct ‘fourth’ or ‘anti-luck’ condition on propositional knowledge. The issue is still a matter of controversy, well over fifty years later. Even so, a handful of leading contenders have emerged: the ‘no-essential falsehoods’ account, the ‘defeasibility’ account, the Dretske-Nozick ‘sensitivity’ account (Dretske 1971; Nozick 1981), and the ‘safety’ account (proposed by Luper-Foy [1984] and Sosa [1999]; taken up by Pritchard [2005], among others). Although the points I am about to make are compatible with both the sensitivity and the safety accounts, for reasons that need not detain us I prefer to work with the safety account over the sensitivity one. And the other accounts, to my mind, tend to over-intellectualize propositional knowledge: they start at the wrong end of the spectrum. But that’s a debate for another day. In the rest of this section I will explain the safety account, and then put it to work in answering our question, to see where it leads.

According to the safety account, propositional knowledge that P is (justified, true) belief that P held on a *safe* basis. When is a true belief that P held on a safe basis? When holding the belief that P on that basis, you would not easily be mistaken. When would you not easily be mistaken? This is usually glossed in terms of possible worlds: a true belief that P is held on a safe basis just in case, in each nearby possible world, if one forms a belief within that world on the same basis as in the actual world, then one forms a true belief in that world.<sup>17</sup>

How does the safe-basis account apply to our cases? Take HOSPITAL. The father knows that his son is healthy; he believes on a safe-basis (perception, testimony, background knowledge, etc.), so he would not easily be mistaken. But his mother’s belief that her grandson is healthy (from her son’s testimony) is not formed on a safe basis, for there is a nearby world where she relies on her son’s testimony to the effect that her grandson is fine but she forms a false belief. That’s why she doesn’t acquire knowledge from him, even though he knows that his son (her grandson) is fine.<sup>18</sup>

Now take FOSSIL. The schoolteacher relies on a safe-basis (careful observation, a sophisticated understanding of the contents of evolutionary theory and the fossil record, etc.) to reach the conclusion – a conclusion she ‘accepts’ but, because of her personal convictions, does not believe – that ancient humans once lived in this area: she relies on a safe-basis to reach a cognitive state of ‘acceptance’ but not of belief (cf. Bratman [1992]). ‘Accepting’ that conclusion, she tells the children that ancient humans once lived in the area. Young schoolchildren being young schoolchildren, they believe her. Her ‘acceptance’ results from a safe-basis, and she wouldn’t say what she does unless it was formed on such a basis, and so her testimony, too, is a safe-basis for the

schoolchildren. When they believe her, they, too, form a belief on a safe-basis. Relying on her testimony, the schoolchildren would not easily be mistaken. That's why they come to know something, something that (as the example goes) no one has previously known.<sup>19</sup>

The safe-basis account also allows us to argue that ordinary cases (where a speaker transfers knowledge) and unusual cases (like FOSSIL) fall within the same epistemic kind. Imagine an ordinary case where knowledge transfers. Why did that happen? Because the speaker believed on a safe basis, which partly explains why her assertion is a safe basis for the hearer, which partly explains why the hearer forms a belief on a safe basis. What happened in the FOSSIL case? The speaker did not believe on a safe basis, but she did 'accept' on a safe basis, which then partly explains why she asserts on a safe basis, and so on. The underlying 'epistemic mechanics' is the same in both so-called *testimonial* knowledge cases and what I called *testimonial knowledge\** cases. But if the epistemic mechanics are the same, the epistemic kinds are the same. I conclude that the 'broader' category of *testimonial knowledge\** is the category that we want to understand if we want to understand the epistemology of testimony, especially the transmission of knowledge from one mind to another.

After all, knowledge per se does not transfer from one mind to another. Suppose that I have some perceptual knowledge. Then I have knowledge on a safe-basis, where the basis involves perception and perceptual experience. Then I tell you, and you learn from me. Did I transfer my perceptual experience to you? Do you now have exactly what I have? No, not at all. You have neither my *perceptual* knowledge nor my *perceptual* warrant. But you do have knowledge, *testimonial* knowledge. Why? Because you formed a belief on a safe basis, because the epistemic force (the safety) of my perceptual belief was transferred through testimony to your belief. Testimony transmits *epistemic force* – it transmits the safety of one basis (perception) to the safety of another basis (comprehension and belief) through the safety of a medium (assertion, telling). Testimony doesn't *transmit* my knowledge to you per se. That's why the motivation for the so-called narrower category '*testimonial* knowledge' (as knowledge essentially derived from someone else's knowledge) puts the wrong foot forward.<sup>20</sup> It is not the speaker's knowledge that secures knowledge for the hearer; it is the epistemic force supporting the speaker's belief that secures the epistemic force supporting the hearer's belief. The same story explains unusual cases like FOSSIL: same epistemic mechanism, same epistemic kind.

## 8. The reliability of testimony and social norms

But why, you might ask, is testimony reliable enough for justified uptake and knowledge? Consistent with everything else that we know about ourselves, what explains the reliability of testimony?

There is so much to say about this topic that I can barely begin to scratch the surface. I'll make three quick remarks, and then discuss one answer in a little more detail, just to take a stab at this important issue.

First, there's always the possibility, consistent with Reid, that it is a part of our God-given nature. Second, there's the parallel possibility, consistent with the growing literature on the evolution of cooperation, that we're born to be helpful, especially in communication. Young children are especially helpful at providing information, unprompted, to adults in apparent need (Tomasello 2009). For a similar reason, parents tell their children the truth out of parental concern. Third, it is obviously frequently the case that, to coordinate with others, we need to tell the truth (I'll need to tell you the correct time that my flight arrives if you are going to pick me up). Another answer, consistent with these three, is that human life is governed by socially shared prescriptions – social norms – for telling the truth and providing useful information. It's this last answer I'll discuss in a little more detail.

What's a social norm? According to Cristina Bicchieri (2006, 2014, 2017), a social norm is not merely a collective pattern of behaviour in a group. Everyone may want to stay warm when they go out, and so everyone may wear a coat during the winter, so we are all behaving the same way. But that's not a social norm; that is a social custom, a pattern of behaviour in a group where everyone acts the same way because we all have the same needs and we've all discovered the same solution. A social norm is not a trend or a fashion either, where people imitate a crowd or follow the leader to go along or to simply rely on what probably works. Nor is a social norm a convention, where people act the same way in a recurrent situation to coordinate their behaviour, such as driving on the right-hand side of the road in the United States.

According to Bicchieri, social norms are patterns of behaviour that are collectively approved or disapproved in a group or population, and enforced by informal sanctions, positive and negative. Social norms arise from (1) (second-order) beliefs we have that other people believe we ought – normatively and not just prudentially – to behave a certain way (these are normative social expectations), and (2) a conditional preference to behave a certain way, provided

we believe that other people believe we ought to believe that way (a preference to meet normative social expectations). The conditional preference exists because of the sanctions, often external but also internal. Positive sanctions include tangible rewards, praise, status, reputation, and so on. Negative sanctions for censured behaviour include punishments, shaming, ostracism, ridicule, and so on. Often just thinking that others would disapprove can stop us in our tracks. Likewise, just thinking that others would approve might lead us to conform (Pettit 1990; Graham 2015b).<sup>21</sup>

Social norms – our normative expectations and conditional preferences to conform – then strongly motivate compliance. ‘With social norms,’ Bicchieri (2017: 34–5) explains, ‘the normative influence is strong and plays a crucial role in driving compliance. It matters to us that most people in our reference network believe we ought to conform to a certain behavioural pattern. This point must be emphasized ... [T]he social pressure to conform, expressed in social expectation that one ought to conform, is a powerful motivator.’ And it is the pressure to conform, enforced by sanctions, that really does a good deal of the motivational work. For, left to ourselves, we often prefer not to conform. Indeed, many social norms emerge as solutions to social dilemmas, where narrow individual self-interest might lead us to behave in non-cooperative ways, such as prisoner’s dilemmas and public goods games. ‘With a [social] norm, there is often the temptation to transgress it – this is precisely why norms must be socially enforced’ (39).

However, sometimes we internalize the norm, so that we positively value the behaviour, and we might conform, regardless of social expectations. ‘Especially with norms that are well established, norm followers tend to value what the norm stands for. An external observer may be induced to think that, since people have a positive attitude toward a norm, they may obey it regardless of what others around them do’ (40). Even so, personal values – our personal normative attitudes – often fall short of motivating behaviour (Fishbein 1967; Eagly and Chaiken 1993), so that social expectations and conditional preferences often play a much stronger role than meets the eye.

What if truth-telling is a social norm? What if we believe that other people believe that we ought to tell the truth, and what if we also believe that there are rewards for conformity and costs for noncompliance? Then we’ll be motivated, in addition to any other motivations that might already be in place, to tell the truth. The social pressure to conform would then be a powerful motivator. We might even positively value telling the truth.

Is it? Yes, indeed. It's a textbook case of a social norm: open any textbook discussion of social norms, and *tell the truth* will surely be on the shortlist of examples. In her contribution to a recent introductory anthology of social science, Bicchieri (2014: 208) proclaims that '[e]ach group has its own norms, and some, like reciprocity or truth telling, [are] very general, spanning all groups'. When providing examples of social norms in his book *Social Action*, Seumas Miller (2001) lists refraining from violence, remaining faithful to one's spouse, avoiding incest, keeping promises, telling the truth. Bowles and Gintis (2011) emphasize a number of times the social norm that one ought to tell the truth. Philip Pettit (1990) treats truth-telling as the central case in his work on social norms. As I said, it's a textbook case.

Social norms are surely a part of the story – a very complex story – as to why testimony is as reliable as it is, not just in general but in varying domains and in different situations. Although maybe the Supreme Being didn't implant the principle of veracity in our hearts, it has a powerful operation in our lives, even so.<sup>22</sup>

## Notes

- 1 Testimony was discussed by Locke, Hume, Reid, and Kant (among others), but then was for the most part marginalized until the late twentieth century. Michael Welbourne (1986) was often a lone voice in the 1970s, later joined by Elizabeth Fricker, among others, in the 1980s. Then Coady's (1992) book and Burge's (1993) essay stirred considerable interest in the early 1990s, which led to a handful of dissertations completed by the end of the 1990s and then a flowering of interest that continues unabated, including the recent publication of two excellent textbooks (Gelfert 2014; Shieber 2015).
- 2 I won't pretend to be as exhaustive or even-handed as possible – I'll be opinionated, for sure – but I'll do my best in the space available. For other ways of orienting the literature, I recommend Fricker (2004), Lackey (2011), Adler (2012), Greco (2012), Gelfert (2014), Goldman and Blanchard (2015), and Shieber (2015). There are issues that I will sideline, including disagreement, epistemic injustice, group knowledge, and group testimony, whether the reliability of a process, or even the cognitive process itself, somehow interestingly 'extends' to the speaker (Goldberg 2010), among other issues. And the assurance view, regrettably, will receive only the barest mention. For the assurance view, see Ross (1986), Moran (2005), McMyler (2011), and Hinchman (2014); against it, see Lackey (2008), Schmitt (2010), and Owens (2017).
- 3 For an up-to-date discussion, see Hetherington (2016).

- 4 For discussion of testimony as a speech act, see Graham (1997, 2015a). For more discussion of testimony-based beliefs, see Graham (2015a, 2016).
- 5 Not every participant to the literature sharply distinguishes these two questions, sometimes resulting in unnecessary confusion. Robert Audi (1997), however, sharply distinguishes these two, arguing that a hearer can acquire testimonial knowledge without simultaneously acquiring a justified testimony-based belief. Audi relies on a general framework where knowledge is belief based on a reliable indicator, and justified belief is belief based on accessible justifiers that can be cited when justifying the belief.
- 6 For the record, I think that this representation of Locke is far from the truth. For one, in context this passage is targeting an almost blind deference to the so-called ‘authority of others’ on matters of theological and abstract philosophical affairs, exactly the kind of topic where Locke, as a spokesperson of the Enlightenment, believed that we can, and should, think for ourselves. For another, Locke also had a very different *conception* of knowledge than we have today. On his *conception* of knowledge, knowledge is the perceived agreement or disagreement among ideas (roughly, but not entirely, what we would consider knowledge of self-evident truths and what can be self-evidently deduced from self-evident truths). Given that conception of *knowledge*, you cannot just acquire knowledge from relying on someone who knows, for even though they have perceived the relevant agreements among ideas to reach a conclusion, it does not follow that by believing their testimony you will ipso facto perceive the relevant agreement among ideas yourself. Hume probably would have agreed with Locke, for they probably shared this conception of knowledge. So even the quotation from Hume is, for the opposite reason, also a little misleading. Furthermore, in line with the quotation from Hume, in Book IV of the *Essay* Locke includes testimony as a ‘ground of probability’: namely, a ground for reasonable belief – what in many cases *we* would count as knowledge. In sum, Locke probably did not mean what most people who cite this passage take him to mean (see Shieber [2009] for more discussion). Regardless, this representation of Locke is a useful foil for common sense, which is why I started with it.
- 7 See Graham (forthcoming) for critical exposition and discussion of Burge’s argumentation. There’s a tendency to equate Burge’s a priori defence of the Principle with the truth of the Principle, so that if the Principle isn’t a priori necessary then it isn’t even true (van Cleve 1996; Faulkner 2011). Burge assumes the Principle and then asks for its basis. The Principle, as we’ve just implied, might have more than one basis: God, evolution by natural selection, human psychology, social psychology, social structures, the nature of promising and the institutions of promising, and so on, or the very nature of reason itself. Burge opts for the last option. His account of its basis may fall short without the Principle falling short in any way.

- 8 For other arguments, see Fricker (1987, 2004, 2016) and Lackey (2008). Lackey thinks that she has a persuasive counterexample to anti-reductionism. Imagine stumbling across an apparent diary written apparently in English, where it is clear to you that it was written by an alien from another planet. You possess no positive reasons for what it says, and simultaneously you possess, she says, no defeating reasons not to believe what it says. But intuitively, she thinks, you are not justified in believing what it says. Generalizing, every testimony-based belief requires non-testimonial support. I don't agree. By my lights, a mature adult will be puzzled by the fact that it was an alien, and so is surely to have defeaters. Perrine (2014: 3236) argues along these lines: the adult in Lackey's case 'has a wealth of background knowledge that ought to make him skeptical in this case' (cf. Burge 2013). However, if the recipient were a young child who could read but had few thoughts, if any, about the consequences of the author's alien status, then we'd have a case of no reasons on either side. But then it seems fine to say that child has a *prima facie* justified testimony-based belief.
- 9 I provide more detailed formulations of some of these issues elsewhere (2006, forthcoming). For testimony aficionados, there are two possible ways to frame the reductionism/anti-reductionism debate. One is about justification (that is how I see it); the other is about knowledge (about converting true belief into knowledge). If you see it the latter way, then you see the reductionist as saying that a true testimony-based belief, provided that it is reductively justified (usually inductively by a track-record), is then knowledge: to really 'reduce' the hearer's *knowledge* to perception, memory, and induction, everything required *for knowledge* must 'reduce' to the hearer's track-record. Seen this way, the 'reductionist' rejects the common-sense idea that the hearer gets what she needs for knowledge from the speaker, and thus rejects testimonial knowledge as involving knowledge transfer of either knowledge or justification from the speaker. This 'reductionist' view is then easy to dismiss with Gettier counterexamples. Lackey (2008: 142–59) does so. Imagine that the speaker is right just by luck, but that the hearer has a pretty good track-record of evidence to justify reliance. Then intuitively the hearer, too, despite the track-record, has a Gettierized justified true belief. The hearer's track-record is then not enough to convert a true belief into knowledge. Lackey argues that a further, anti-Gettier, condition is required, involving the reliability of the speaker's testimony. Paralleling her formulation of reductionism, she formulates 'non-reductionism' as the view that a true testimony-based belief, provided that the hearer has no defeaters, is knowledge. This view also suffers the same fate: Gettier counterexamples are easy to formulate. She concludes that both reductionism (about knowledge) and anti-reductionism (about knowledge) are wrong, for the same reason. She then formulates a 'hybrid' theory, neither 'reductionist' nor 'non-reductionist', for testimonial knowledge. (By these lights, any theory that sees the speaker as playing some role in meeting the anti-luck condition is a

- 'hybrid' theory – which includes nearly everyone working on the topic, including our common-sense view that the hearer's knowledge depends on the speaker's knowledge.) Putting all of this to one side, Lackey ends up advancing a reductionist account of testimonial *justification* (she would call this the 'rationality dimension' of justification), for she adopts a 'positive reasons requirement' on testimonial knowledge, where 'for each report R, the positive reasons justifying R cannot ultimately be testimonially grounded, where this means that the justificatory or epistemic chain leading up to R cannot "bottom out" in testimony' (186). Although Lackey spends some time in her book arguing that adults have plenty of positive reasons supporting testimony, she spends little time addressing whether adults satisfy this reductive requirement, which is just the requirement at issue.
- 10 For other arguments, see Coady (1992) and Rysiew (2007). For some discussion of these arguments, see Fricker (1995) and Graham (2000c).
  - 11 For detractors, see Lyons (1997), Shogenji (2006), van Cleve (2006), Lackey (2008), and Kenyon (2013). Fricker (1995, forthcoming a) is well known for the distinction between 'local' and 'global' reductionism. She grants that testimonial justification does not 'globally' reduce (individuals do not possess a reductive justification for the proposition that testimony is generally reliable), but she insists that it 'locally' reduces (individuals possess reductive justification for propositions of the form that this speaker on this occasion is trustworthy on this topic). Many commentators, including myself, doubt that this move succeeds, for they doubt that individuals possess adequate *reductive* local justifications, never mind adequate reductive global justifications. See Insole (2000), Gelfert (2009), and Graham (2016a).
  - 12 Selective trust in special cases some of the time is one thing; reductive trust in every case is another altogether. I discuss this argument in considerable detail elsewhere (2016a).
  - 13 See Graham (2016a, sec. 5) for a review of the literature and references. See also Harris (2012), Cole, Harris, and Koenig (2012), and Harris and Lane (2014).
  - 14 Paul Faulkner (2011) claims to occupy a middle ground. He argues that the anti-reductionist is wrong – we don't enjoy a *prima facie* default right to testimony – but the reductionist is wrong, too – we don't need a reductive justification showing that testimony in general, or in the particular 'local' case, is reliable. Because of the 'problem of communication' – because hearers want the truth, but because speakers, as rational self-interested agents, want the freedom to lie when that serves their interests – hearers need a reason in each and every case (even small children) to justifiably believe the speaker's report; it would be epistemically irrational otherwise to believe testimony without such a reason. But if the reason is not a reductive argument showing the reliability of the speaker, what is it? The reason is the hearer's affective trust that the speaker will prove trustworthy, where that is the normative expectation that the speaker should prove trustworthy, for

the speaker and the hearer are both subject to, and have internalized, the social norm of trustworthiness (of truth telling). (More on social norms in the final section.) The speaker, then aware of the hearer's normative expectation, will prove trustworthy (will fulfil the expectation). The hearer then possesses a mini-argument rationalizing acceptance of the speaker's assertion: (1) I normatively expect truth telling; (2) the speaker knows this; (3) the speaker has also internalized the norm, and so will prove trustworthy; (4) so, the speaker is telling the truth; (5) so, it is rational to believe the speaker's assertion. Although I agree with Faulkner about the important role that social norms play in the overall epistemology of testimony, I do not agree they tell the whole story: they are one mechanism among others underwriting the reliability of testimony, not the only one. (More on this in the final section.) Also, I think that he has conflated two issues. One issue is that of whether hearers need to represent, in their psychology, a rationalizing argument justifying their reliance. Another issue is that of whether there is a problem of communication, whether the possibility of deceit shows that the anti-reductionist cannot be right. Once social norms of truth-telling enter the picture, the problem of communication goes away. So, the case against the anti-reductionist dissolves. Why then think that everyone (even small children), in each and every case, needs a mini-argument to rationalize trust? Well, if you were an internalist with a tendency to intellectualize (some would say hyper-intellectualize) the epistemology of testimony, then you'd go for that conclusion. But since I don't, I won't. Faulkner thinks that the problem of cooperation motivates an internalist result, but it is (I believe) his internalist convictions that are doing the work, not the problem of cooperation. The structural problem of communication is one thing; the case for internalism is another. For discussion of Faulkner in more detail, see Graham (2012, 2013).

- 15 Here is another: ASTROLOGIST. Mary sometimes believes that it is raining in her village because she pulls the drapes and looks outside, and sometimes because she consults an astrological table with the drapes closed. Today, Mary looks outside and sees that it is raining. But when she takes your phone call and tells you that it is raining, you do not come to know that it is, for she would just as easily tell you that it is raining when it is not. Relying on her testimony about the weather, you would easily be mistaken (Peacocke 1986). For discussion of the issue and additional counterexamples, see Graham (2000a, 2016b) and Lackey (2008).
- 16 The difference between FOSSIL and Lackey's (1999, 2008) SCHOOLTEACHER example is that FOSSIL involves testimonial knowledge of a proposition that no one has ever previously known. In Lackey's example, the schoolteacher is passing on knowledge of evolutionary theory, well-known throughout the science. FOSSIL seems to illustrate the possibility that testimony can 'generate' knowledge, namely,

be the very first source of a new piece of knowledge, something that typically only occurs in perception, introspection, a priori understanding, and reasoning from those three sources. Those who think that testimony cannot ‘generate’ knowledge allow such cases like Lackey’s (they only insist on knowledge in the chain of sources), hence the need to create a case that avoids the qualification. For further discussion of these and related cases, see Audi (2006), Fricker (2006, 2015, 2016), Graham (2006a, 2016b), Faulkner (2010), Burge (2013), Carter and Nickel (2014), Wright (2016), and Bachman and Graham (forthcoming).

- 17 ‘Basis’ safety is not to be confused with ‘belief’ safety. A basis is safe just in case it leads you to a true belief in the actual world and doesn’t lead you astray in any nearby possible world. A belief is safe just in case it is true in the actual world and true in all nearby worlds. The standard example to illustrate this difference involves flipping a coin to form the belief that  $14 + 23 = 37$ . If you believe because you flipped a coin, then you didn’t believe on a safe basis. In a nearby world the coin flip will come up the other way, and you’ll end up believing a falsehood. But the belief that  $14 + 23 = 37$  (on the other hand) is true in all nearby worlds where you believe it, because it is true in any world where you believe it, because (as a necessary truth) it is true in all worlds! Forming a belief on a coin flip is not a safe basis, but the belief in the mathematical proposition is a safe belief. The safety account of knowledge is a safe-basis account, not a safe-belief account: there are safe beliefs that fail to measure up to knowledge, but no belief is knowledge without being held on a safe basis.
- 18 Mary in ASTROLOGY works the same way. By looking out the window through the drapes, Mary forms the belief, on a safe basis, that it is raining outside. That’s why she knows that it is raining. But when she tells you that it is raining, you do not form a belief on a safe basis. She would easily tell you that it is raining outside when it is not; relying on her testimony about the weather would then easily lead you astray. That’s why you don’t acquire knowledge from her, even though she knows it is raining today.
- 19 For more discussion, detail, and defence of – and further references for – the safe-basis account of testimonial knowledge, see Graham (2016b).
- 20 For additional arguments for a further broadening of the epistemology of testimony, see Graham (2000a, 2015a, 2016b).
- 21 Among our social practices – our shared patterns of behaviour – there are many kinds: customs, descriptive norms, conventions, and social norms. Thus, it is not trivially true that some of our epistemic norms are social norms (or habits, or customs), just because we have a social practice of asking for and giving reasons. For that practice might be a custom, a descriptive norm, a convention, or a social norm. The case must be made.
- 22 For a longer discussion of some of these issues, see Graham (2015b). For critical engagement, see Fricker (forthcoming a). For more exploration of Bicchieri’s

framework for epistemic norms as social norms, see Henderson and Graham (forthcoming a,b). Faulkner's (2010, 2011) work is clearly relevant here, too (as discussed previously). The social norms account is an important alternative to both the constitutive norms account of the epistemology of testimony (e.g. see Goldberg [2015]; for critical discussion, see Johnson [2015]) and the assurance view, for each is sensitive to the normative dimension of testimony, but each locates it in other kinds of normativity, kinds that do not directly address the motivational basis for the reliability of testimony.

## References

- Adler, J. (2012), 'Epistemological Problems of Testimony', in E. N. Zalta (ed.), *Stanford Encyclopedia of Philosophy*, <https://plato.stanford.edu/entries/testimony-episprob>.
- Ahlstrom-Vij, K. (2016), 'Is There a Problem with Cognitive Outsourcing?', *Philosophical Issues*, 26: 7–24.
- Audi, R. (1997), 'The Place of Testimony in the Fabric of Justification and Knowledge', *American Philosophical Quarterly*, 34: 405–22.
- Audi, R. (2006), 'Testimony, Veracity, Credulity', in J. Lackey and E. Sosa (eds), *The Epistemology of Testimony*, Oxford: Clarendon Press.
- Bachman, Z. and Graham, P. J. (forthcoming), 'Counterexamples to Transmission', in M. Fricker, P. Graham, D. Henderson, N. Pederson, and J. Wyatt (eds), *The Routledge Handbook of Social Epistemology*, Routledge.
- Bicchieri, C. (2006), *The Grammar of Society: The Nature and Dynamics of Social Norms*, New York: Cambridge University Press.
- Bicchieri, C. (2014), 'Norms, Conventions, and the Power of Expectations', in N. Cartwright and E. Montuschi (eds), *The Philosophy of Social Science: A New Introduction*, Oxford: Oxford University Press.
- Bicchieri, C. (2017), *Norms in the Wild: How to Diagnose, Measure, and Change Social Norms*, New York: Oxford University Press.
- Bowles, S. and Gintis, H. (2011), *A Cooperative Species: Human Reciprocity and Its Evolution*, Princeton: Princeton University Press.
- Bratman, M. E. (1992), 'Practical Reasoning and Acceptance in a Context', *Mind*, 101: 1–16.
- Burge, T. (1993), 'Content Preservation', *The Philosophical Review*, 102: 457–88.
- Burge, T. (2013), 'Postscript: Content Preservation', in his *Cognition through Understanding: Philosophical Papers, Volume 3*, Oxford: Oxford University Press.
- Carter, J. A. and Nickel, P. J. (2014), 'On Testimony and Transmission', *Episteme*, 11: 145–55.
- Coady, C. A. J. (1992), *Testimony: A Philosophical Study*, Oxford: Clarendon Press.
- Cole, C. A., Harris, P. L., and Koenig, M. A. (2012), 'Entitled to Trust? Philosophical Frameworks and the Evidence from Children', *Analyse & Kritik*, 34: 195–216.

- Dretske, F. (1971), 'Conclusive Reasons', *Australasian Journal of Philosophy*, 49: 1–22.
- Eagly, A. H. and Chaiken, S. (1993), *The Psychology of Attitudes*, Orlando, FL: Harcourt Brace Jovanovich College.
- Faulkner, P. (2010), 'Norms of Trust', in A. Haddock, A. Millar, and D. Pritchard (eds), *Social Epistemology*, Oxford: Oxford University Press.
- Faulkner, P. (2011), *Knowledge on Trust*, Oxford: Oxford University Press.
- Fishbein, M. (1967), 'A Consideration of Beliefs and Their Role in Attitude Measurement', in M. Fishbein (ed.), *Readings in Attitude Theory and Measurement*, New York: Wiley.
- Fricker, E. (1987), 'The Epistemology of Testimony', *Aristotelian Society Supplementary Volume*, 61: 57–106.
- Fricker, E. (1994), 'Against Gullibility', in A. Chakrabarti and B. K. Matilal (eds), *Knowing from Words: Western and Indian Philosophical Analysis of Understanding and Testimony*, Dordrecht: Kluwer Academic.
- Fricker, E. (1995), 'Telling and Trusting: Reductionism and Anti-reductionism in the Epistemology of Testimony', *Mind*, 104: 393–411.
- Fricker, E. (2004), 'Testimony: Knowing through Being Told', in I. Niiniluoto, M. Sintonen, and J. Woleński (eds), *Handbook of Epistemology*, Dordrecht: Kluwer Academic.
- Fricker, E. (2006), 'Second-Hand Knowledge', *Philosophy and Phenomenological Research*, 73: 592–618.
- Fricker, E. (2015), 'How to Make Invidious Distinctions amongst Reliable Testifiers', *Episteme*, 12: 173–202.
- Fricker, E. (2016), 'Unreliable Testimony', in B. P. McLaughlin and H. Kornblith (eds), *Alvin Goldman and His Critics*, Oxford: Blackwell.
- Fricker, E. (2018), 'Inference to the Best Explanation and the Receipt of Testimony: Testimonial Reductionism Vindicated', in T. Poston and K. McCain (eds), *Best Explanations: New Essays on Inference to the Best Explanation*, Oxford: Oxford University Press.
- Gelfert, A. (2009), 'Indefensible Middle Ground for Local Reductionism about Testimony', *Ratio*, 22: 170–90.
- Gelfert, A. (2014), *A Critical Introduction to Testimony*, London: Bloomsbury.
- Gettier, E. (1963), 'Is Knowledge Justified True Belief?' *Analysis*, 23: 121–3.
- Goldberg, S. (2010), *Relying on Others: An Essay in Epistemology*, Oxford: Oxford University Press.
- Goldberg, S. (2015), *Assertion: The Philosophical Significance of Assertoric Speech*, Oxford: Oxford University Press.
- Goldberg, S. and Henderson, D. (2006), 'Monitoring and Anti-reductionism in the Epistemology of Testimony', *Philosophy and Phenomenological Research*, 72: 600–17.
- Goldman, A. and Blanchard, T. (2015), 'Social Epistemology', in E. N. Zalta (ed.), *Stanford Encyclopedia of Philosophy*: <https://plato.stanford.edu/entries/epistemology-social>.

- Goldman, A. I. (1986), *Epistemology and Cognition*, Cambridge, MA: Harvard University Press.
- Graham, P. J. (1997), 'What Is Testimony?' *The Philosophical Quarterly*, 47: 227–32.
- Graham, P. J. (2000a), 'Conveying Information', *Synthese*, 123: 365–92.
- Graham, P. J. (2000b), 'Transferring Knowledge', *Noûs*, 34: 132–52.
- Graham, P. J. (2000c), 'The Reliability of Testimony', *Philosophy and Phenomenological Research*, 61: 695–709.
- Graham, P. J. (2006a), 'Can Testimony Generate Knowledge?' *Philosophica*, 78: 105–27.
- Graham, P. J. (2006b), 'Liberal Fundamentalism and Its Rivals', in J. Lackey and E. Sosa (eds), *The Epistemology of Testimony*, Oxford: Clarendon Press.
- Graham, P. J. (2006c), 'Testimonial Justification: Inferential or Non-inferential?' *The Philosophical Quarterly*, 56: 84–95.
- Graham, P. J. (2010), 'Testimonial Entitlement and the Function of Comprehension', in A. Haddock, A. Millar, and D. Pritchard (eds), *Social Epistemology*, Oxford: Oxford University Press.
- Graham, P. J. (2012), 'Testimony, Trust and Social Norms', *Abstracta*, 6: 92–116.
- Graham, P. J. (2013), 'Review of Faulkner, *Knowledge on Trust*', *Notre Dame Philosophical Reviews*, <http://ndpr.nd.edu/news/knowledge-on-trust>.
- Graham, P. J. (2015a), 'Testimony as Speech Act, Testimony as Source', in C. Mi, M. Slote, and E. Sosa (eds), *Moral and Intellectual Virtues in Western and Chinese Philosophy*, New York: Routledge.
- Graham, P. J. (2015b), 'Epistemic Normativity and Social Norms', in J. Greco and D. Henderson (eds), *Epistemic Evaluation: Purposeful Epistemology*, Oxford: Oxford University Press.
- Graham, P. J. (2016a), 'Formulating Reductionism about Testimonial Warrant and the Challenge from Childhood Testimony', *Synthese*, online 16 June 2016: <http://dx.doi.org/10.1007/s11229-016-1140-y>.
- Graham, P. J. (2016b), 'Testimonial Knowledge: A Unified Account', *Philosophical Issues*, 28: 172–86.
- Graham, P. J. (forthcoming), 'Sincerity and the Reliability of Testimony: Burge on the A Priori Basis of Testimonial Warrant', in E. Michaelson and A. Stokke (eds), *Lying: Language, Knowledge and Ethics*, Oxford: Oxford University Press.
- Greco, J. (2012), 'Recent Work on Testimonial Knowledge', *American Philosophical Quarterly*, 49: 15–28.
- Harris, P. (2012), *Trusting What You Are Told: How Children Learn from Others*, Cambridge, MA: Harvard University Press.
- Harris, P. L. and Lane, J. D. (2014), 'Infants Understand How Testimony Works', *Topoi*, 33: 443–58.
- Henderson, D. (2008), 'Testimonial Belief and Epistemic Competence', *Noûs*, 42: 190–221.
- Henderson, D. and Graham, P. (forthcoming a), 'Epistemic Norms and the "Epistemic Game" They Regulate: The Basic Structured Epistemic Costs and Benefits', *American Philosophical Quarterly*.

- Henderson, D. and Graham, P. (forthcoming b), 'A Refined Account of the "Epistemic Game": Epistemic Norms, Temptations and Epistemic Cooperation', *American Philosophical Quarterly*.
- Hetherington, S. (2016), *Knowledge and the Gettier Problem*, Cambridge: Cambridge University Press.
- Hinchman, E. (2014), 'Assurance and Warrant', *Philosophers' Imprint*, 14: 1–58.
- Insole, C. J. (2000), 'Seeing Off the Local Threat to Irreducible Knowledge by Testimony', *The Philosophical Quarterly*, 50: 44–56.
- Johnson, C. R. (2015), 'Testimony and the Constitutive Norm of Assertion', *International Journal of Philosophical Studies*, 23: 356–75.
- Kelp, C. (2016), 'Assertion: A Function First Account', *Noûs*, online first: <http://dx.doi.org/10.1111/nous.12153>.
- Kenyon, T. (2013), 'The Informational Richness of Testimonial Contexts', *The Philosophical Quarterly*, 63: 58–80.
- Lackey, J. (1999), 'Testimonial Knowledge and Transmission', *The Philosophical Quarterly*, 49: 471–90.
- Lackey, J. (2011), 'Testimonial Knowledge', in S. Bernecker and D. Pritchard (eds), *The Routledge Companion to Epistemology*, New York: Routledge.
- Lyons, J. (1997), 'Testimony, Induction, and Folk Psychology', *Australasian Journal of Philosophy*, 75: 163–78.
- Luper-Foy, S. (1984), 'The Epistemic Predicament: Knowledge, Nozickian Tracking, and Skepticism', *Australasian Journal of Philosophy*, 62: 26–49.
- McDowell, J. (1994), 'Knowledge by Hearsay', in A. Chakrabarti and B. K. Matilal (eds), *Knowing from Words: Western and Indian Philosophical Analysis of Understanding and Testimony*, Dordrecht: Kluwer Academic.
- McMyler, B. (2011), *Testimony, Trust, and Authority*, New York: Oxford University Press.
- Mercier, H. and Sperber, D. (2011), 'Why Do Humans Reason? Arguments for an Argumentative Theory', *Behavioral and Brain Sciences*, 34: 57–74.
- Michaelian, K. (2010), 'In Defence of Gullibility: The Epistemology of Testimony and the Psychology of Deception Detection', *Synthese*, 176: 399–427.
- Miller, S. (2001), *Social Action: A Teleological Account*, Cambridge: Cambridge University Press.
- Moran, R. (2005), 'Getting Told and Being Believed', *Philosopher's Imprint*, 5: 1–29.
- Nozick, R. (1981), *Philosophical Explanations*, Cambridge, MA: Harvard University Press.
- Owens, D. (2017), 'Human Testimony', in his *Normativity and Control*, Oxford: Oxford University Press.
- Pritchard, D. (2005), *Epistemic Luck*, Oxford: Clarendon Press.
- Ross, A. (1986), 'Why Do We Believe What We Are Told?' *Ratio*, 1: 69–88.
- Russell, B. (1948). *Human Knowledge: Its Scope and Limits*, London: George Allen & Unwin.

- Rysiew, P. (2007), 'Beyond Words: Communication, Truthfulness, and Understanding', *Episteme*, 4: 285–304.
- Schmitt, F. (2010), 'The Assurance View of Testimony', in A. Haddock, A. Millar, and D. Pritchard (eds), *Social Epistemology*, Oxford: Oxford University Press.
- Shieber, J. (2009), 'Locke on Testimony: A Reexamination', *History of Philosophy Quarterly*, 26: 21–41.
- Shieber, J. (2014), 'Against Credibility', *Australasian Journal of Philosophy*, 90: 1–18.
- Shieber, J. (2015), *Testimony: A Philosophical Introduction*, New York: Routledge.
- Shogenji, T. (2006), 'A Defense of Reductionism about Testimonial Justification', *Noûs*, 40: 331–46.
- Sosa, E. (1999), 'How to Defeat Opposition to Moore', *Philosophical Perspectives*, 13: 141–53.
- Sperber, D. (2013), 'Speakers Are Honest Because Hearers Are Vigilant: Reply to Kourken Michaelian', *Episteme*, 10: 61–71.
- Sperber, D., Clement, F., Heintz, C., Mascaro, O., Mercier, H., Origg, G., and Wilson, D. (2010), 'Epistemic Vigilance', *Mind and Language*, 25: 359–93.
- Tomasello, M. (2009), *Why We Cooperate*, Cambridge: The MIT Press.
- van Cleve, J. (2006), 'Reid on the Credit of Human Testimony', in J. Lackey and E. Sosa (eds), *The Epistemology of Testimony*, Oxford: Clarendon Press.
- Welbourne, M. (1986), *The Community of Knowledge*, Aberdeen: Aberdeen University Press.
- Wright, S. (2016), 'The Transmission of Knowledge and Justification', *Synthese*, 193: 293–311.
- Zollman, K. J. S. (2015), 'Modeling the Social Consequences of Testimonial Norms', *Philosophical Studies*, 172: 2371–83.

