

Testimony

They told you so

In the realm of knowledge, many of our prized possessions come to us second-hand. We rely on others for our grasp of everything from the geography of distant places to mundane facts about the lives of our friends. If we couldn't use others as sources, we would lose our grip on topics as diverse as ancient history (except what we could discover through our own personal archaeological expeditions) and celebrity weddings (unless we start getting invited). Testimony evidently expands our horizons: the challenge is in explaining exactly how (and how far). Does listening to other people—or reading what they have written—supply us with knowledge in a unique or distinctive way? Do we need special reasons to trust people in order to gain knowledge from them? What should we think about resources like Wikipedia, where most articles have multiple and anonymous authors?

At one extreme, some philosophers have argued that testimony never actually provides knowledge (John Locke will be our star example of this position). At the other end of the spectrum, some philosophers argue that testimony not only provides knowledge, but does so in a distinctive way. In this view, testimony is a special channel for receiving knowledge, a channel with the same basic status as sensory perception and reasoning (this type of position was embraced in classical Indian philosophy, and is now popular in Anglo-American theory as well). By exploring both extremes, as well as leading middle-ground views, we can identify the factors that are broadly agreed to matter most to how we absorb what people say.

No way to know

When does testimony supply knowledge? Some philosophers say: 'Never.' To see why philosophers might be sceptical about testimonial knowledge, even if they aren't sceptical about other kinds of knowledge,

it first helps to clarify what we mean by ‘testimony’. In an act of testimony, someone tells you something—through speech, gestures, or writing—and the content of what they are telling you plays a special role in what you get out of the exchange. Even sceptics about testimonial knowledge can agree that ordinary perceptual knowledge can be generated by the event of hearing or reading what someone says. For example, imagine either seeing that someone has written ‘I have neat handwriting’ on a slip of paper, or hearing someone saying ‘I have a hoarse voice.’ If you can indeed see that the writing is neat or hear that the voice is hoarse, you come to know the truth of what is said or written. But your knowledge here is perceptual, rather than testimonial, because the content of what is written or said plays no special role in what you learn: the sentence ‘Smith got the job’ would work just as well to convey the beauty of the handwriting or the roughness of the voice. If you believe something on the basis of my testimony, you understand what I am saying, and take my word for it.

For John Locke, there was a sharp contrast between perceptual knowledge (for example, the knowledge that a voice you are now hearing is hoarse) and whatever it is we get through testimony (for example, the news that Smith got the job). The key difference is certainty, which for Locke is a necessary condition for knowledge. Because perception can make you immediately certain of something, as certain as you are intuitively that red is not black, you can gain knowledge perceptually. What we get from testimony, Locke says, is at best highly probable, as opposed to certain. During an English winter, if you see a man walking across an icy lake, then you know that this man is crossing the lake. If someone else tells you that she has seen a man walking across the lake, then as long as your informant is trustworthy and what she says fits your own past observations, it is rational to consider the report to be very likely to be true, but you will not actually know that it is.

A key role is played by your own background experiences. Locke tells the story of the King of Siam hearing from a Dutch ambassador that water in Holland becomes solid enough in winter to support the weight of a man, or even an elephant (if you could coax an elephant to Holland in the winter). The king is said to have replied, ‘Hitherto I have believed the strange things you have told me, because I look upon you as a sober fair man, but now I am sure you lie.’ Locke is sympathetic to the doubting king: given all the king’s past experiences in the tropics, it is

entirely reasonable of him to find it more probable that the ambassador is lying than that water ever naturally becomes solid.

Even if testimony is never quite certain, it can still be more or less likely to be true, and Locke counsels maintaining a level of confidence in testimony that reflects the strength of your evidence. He has a complex formula for determining the perfectly rational degree of confidence in testimony: after first weighing how well it fits with your own experience, you must take into account the following six factors:

1. The number of witnesses
2. Their integrity
3. Their skill
4. The purpose they have in supplying their report
5. The internal consistency of what is conveyed, and the circumstances of your hearing it
6. Whether there is any contrary testimony

While he thinks testimony doesn't transmit knowledge, Locke doesn't think we should generally resist what others tell us: he says that the reasonable person will assent to testimony. When it fits with our own observations and scores highly on his six-point checklist, what we get from testimony is for practical purposes very like knowledge ('we receive it as easily, and build as firmly upon it, as if it were certain knowledge'). But what we get from testimony is unlike the discovery that red is not black, according to Locke, because it is open to being undermined later by further experiences. (You thought you had no reason to distrust her, but the woman who told you she saw a man crossing the lake was making it up, and later you hear from ten other people that the ice was too thin to cross today.) Locke argues that this vulnerability to future contrary reports means that what we get from testimony doesn't literally count as knowledge.

We'll look at a challenge to Locke's reasoning shortly. But first it's worth taking a moment to appreciate how radical his position really is. If Locke is right, then the proper answer to the question 'Do you know where you were born?' is 'no' (assuming that your beliefs on this matter are, like most people's, determined by what your family has told you, or what is written on your birth certificate). You could say that it is very probable you were born in a certain place, but not having retained first-hand experience of the location, you won't have knowledge of this fact.

You also don't know that George Washington was once the president of the United States, or that Antarctica exists (assuming you haven't been there yourself). Lockean readers of this book can't even describe themselves as knowing that John Locke ever lived: they should at most consider it highly probable that he did.

It's clear that Locke is going against the way we ordinarily speak: we very freely describe people as gaining knowledge through testimony. ('Does Jones know that Smith got the job?' — 'Yes, he does—I just told him.')

However, we can recognize that in many situations the way we ordinarily speak is not strictly accurate (for example, when we talk of the sun rising or setting when really it is the earth that is rotating). Does Locke have good reasons of principle for saying that strictly speaking we do not gain knowledge from testimony? His argument about vulnerability to later doubts is questionable, in part because it seems to apply equally well to judgements grounded in perception and memory, which he does want to classify as knowledge. Locke thinks that perception enables you to know, for example, that you are now reading a book, and also to remember later that you were reading a book, retaining your perceptual knowledge through memory. However, it's possible here too that later on you will come to doubt yourself. Even if you are really perceiving something now, you might have doubts later on, perhaps wondering whether you were only dreaming. Locke doesn't seem to think that the possibility of stirring up later doubts should undermine your claim to know right now, as long as you are actually now perceiving and not dreaming. But a parallel argument could be applied to testimony: if someone knowledgeable tells you that Smith got the job, and you don't actually have any doubts about what they are saying right now, then you should now have the certainty needed for knowledge. If you start doubting later on, for example because of contrary testimony, you could lose that knowledge, but this is not proof that you never had it. If your informant was knowledgeable, then your later doubts couldn't show that what you originally judged was untrue: if your informant knew that Smith got the job, then it must be true that Smith got the job. Any report to the contrary is misleading. Of course, there could be situations in which you fail to have doubts, and take the word of a liar as if she were telling the truth, but these situations are parallel to situations in which you are taken in by a perceptual illusion. If there is a big difference between the knowledge-providing powers of perception and testimony, Locke hasn't shown us what it is.

The middle ground: reductionism

In claiming that testimony never supplies knowledge, Locke occupies a minority position. Most philosophers are more positive about it. The main moderately positive position is reductionism: we do gain knowledge through testimony, but the knowledge-providing power of testimony is nothing special. Whether we are reading, listening, or watching someone's gestures or sign language, we receive testimony through ordinary sense perception. Assuming everything goes well, sensory perception lets us know that a speaker has said a sentence. In order to gain knowledge of what the sentence itself says, and not just the fact that the speaker said it, we rely on our ordinary powers of inference and perception. There is still something Lockean about this: you look at the sort of factors on Locke's checklist for the likelihood of testimony being true (how well it fits with past experience, evidence about the integrity of the speaker, and so on), but when what you are hearing scores high enough, you come to know the communicated proposition. This way of thinking about testimony is known as 'reductionism', because the knowledge-providing power of testimony can be reduced to the knowledge-providing power of other sources, notably perception, memory, and inference.

Reductionism comes in two flavours: global and local. According to global reductionism, your own experience of the world gradually teaches you that testimony, in general, is a fine source of knowledge. As a youth, you ask for directions to the train station, someone tells you, and then even if you don't yet know the truth of what they said, you can follow the directions and confirm their truth for yourself. Because you can often double-check the truth of what people say, over time you gain knowledge of the track record of past testimony, which then works as a positive experience-based reason to accept present testimony. An ordinary adult can know that the train station is down the road and to the right as soon as he is told, not because testimony has any distinctive knowledge-generating power, but because his own past perceptions, memories, and inferences support accepting what he now hears. The global reductionist doesn't have to say that you should believe absolutely everything you hear: if you are in a situation where there are special undermining factors—for example, if you know that the person you are talking to has a strong incentive to lie—then you can take that into account. But when

there are no special warning signs, the global reductionist says you have a standing positive reason to believe what you are told.

The local reductionist tries something more finely tailored: rather than seeking a blanket positive reason to trust all testimony, the local reductionist suggests that you look for specific positive reasons, in any given situation, to accept the word of the person you are hearing on the topic she is speaking about. Is this person an expert? Has she told you the truth in the past? How plausible is her story now? Again the specific reasons we rely on ultimately come from perception, inference, and memory, rather than on testimony itself. If these ordinary reasons are strong enough in a given situation, you can know the truth of what you are being told.

Both forms of reductionism allow us to say that most adults know where they were born and know that Antarctica exists. When testimony comes from close and trusted witnesses (like your parents, telling you about your birthplace) or from appropriate experts (numerous map-makers and travel writers, telling you about a distant continent), then it can supply knowledge, according to both forms of reductionism. In cases where you don't have any special reasons to trust your informant—lost in a strange city, you ask a total stranger for directions—the global reductionist can say that you gain knowledge, but the typical local reductionist cannot. (If you are lucky, the local reductionist says, you gain a true belief.)

Local reductionism can sound very calculating: in practice, we don't often weigh the reasons to trust someone before accepting their word. However, local reductionism about testimonial knowledge is not a descriptive theory about how we actually form our beliefs in daily practice: it's a theory about the conditions under which those beliefs deserve to count as knowledge. Even if we tend to trust strangers blindly when we ask them for directions, the local reductionist suggests that we shouldn't think of ourselves as gaining knowledge on this basis. Attaining knowledge requires greater caution, if local reductionism is right. Taking this line requires explaining just why we need greater caution for testimonially grounded knowledge than for knowledge grounded in perception and reason. It's true that testimony can let us down (sometimes our informants are dishonest or confused) but perception can also let us down (sometimes our eyes play tricks on us). One possible reason why testimony could be special is that it involves free agents who have purposes of their own. Human communication differs from the communication among bees, for example, who reliably

signal to each other the location of nectar-bearing plants. A bee who learns the location of nectar from another bee is able to fly there as well as if it had witnessed that place: bee signalling gives bees the benefit of each other's experience, in what is sometimes called 'cognition by proxy'. Bee signals can be defective (if the bee is sick, or if the nectar-bearing plants are moved by an interfering researcher after the first bee's contact), but these defects are like the defects in our perceptual organs (when we are sick, or when things are moved behind our backs). One reason that bees can gain the benefit of other bees' experiences directly is that they cannot deliberately deceive. A local reductionist could stress that caution becomes important for communications between members of a sneakier species like ours.

Alternatively, local reductionists could dispute the suggestion that we often rely blindly on the advice of strangers. Perhaps we do ordinarily exercise caution, but in ways that are subtler than explicit weighing of the reasons to trust someone. Recent empirical work on 'epistemic vigilance' has advanced our understanding of how and when we actually accept the word of others. Even if we aren't explicitly thinking to ourselves about the reliability of the stranger we've asked for directions, we could be monitoring his facial expressions and speech patterns to assess how trustworthy he is. Better insight into our actual practices can help us see whether the local reductionist is in fact proposing an account that fits those practices quite closely, or instead proposing that our actual practices are sloppy and we don't have knowledge as often as we think.

Testimony as a distinctive source of knowledge

A still more generous approach to testimony is possible. Instead of seeing testimony as dependent on other ways of knowing, such as past experience and reasoning, you might think of it as a basic source of knowledge in its own right. According to the direct view of testimony (sometimes also called the 'default' view), when your knowledgeable co-worker tells you that Smith got the job, you know that Smith got the job, and your knowledge isn't dependent on your reasoning about the track record of testimony or the reliability of that particular co-worker. You just have to understand what a knowledgeable informant is saying in order to gain knowledge. It's true that the testimonial channel takes input from sensory perception (you have to be able to hear what someone is saying, or read what they have written); reasoning also takes input from

sensory perception (for example, you look at a half-solved sudoku puzzle and start to work out the answer). But testimony remains a distinct way of knowing something, just as reasoning is distinct from pure sensory perception. The way you think when you understand what someone says is different from the way you think when you see something with your own eyes, and different again from the way you think when you are engaged in reasoning or puzzle solving.

The direct view of testimony dates back a long way: it is defended by the Indian philosopher Akṣapāda Gautama in the 2nd century C.E. Gautama maintains that testimony is a special channel through which we gain knowledge, and emphasizes that testimony is not a form of inference. We do not think to ourselves: ‘Lee has said that Smith got the job, and Lee is a reliable person, therefore Smith got the job.’ We know, as soon as we understand what Lee has said, that Smith got the job (we can also focus on the fact that Lee was the person who told us about this, but that isn’t the main thing we pick up). Unlike local reductionism, the direct view has no problem at all with gaining knowledge from strangers: the classical Indian line is that knowledge can be gained directly not only from sages and ‘noble compatriots’ but also from ‘barbarians’, as long as they have knowledge and intend to share it.

Contemporary advocates of the direct view emphasize that trust in testimony plays a large role in the acquisition of language and in our everyday practices of communication. Where reductionists and Lockean think it is right to maintain a neutral stance towards public testimony until we can verify it with our private resources (our own perceptions and inferences), advocates of the direct view suggest that we do not have sufficient private resources available to manage that kind of verification. Vital knowledge of what words mean, for example, is made possible only if we can gain knowledge directly by being told something by another. We wouldn’t be able to understand each other in the first place if we didn’t start by trusting others to tell the truth and accepting what they say at face value. On this view, we drink in what others say, in something like the way bees do.

Even for the maximally generous direct view of testimony, there are still certain conditions that must be met: in order to gain knowledge, what your informant says must actually be true, and (at least according to most non-reductionists) your informant must also know that it is true, as opposed to just managing a lucky guess. Plato gives an example to illustrate this last point: he describes a slick lawyer who has to

defend a client against an assault charge. This client is in fact innocent, but has no witnesses to back him up. Although the lawyer himself has no idea whether what he is saying is true, he does a great job of telling the jury that his client is innocent, using his charisma to make the jury believe him. Plato then raises the question of whether the members of the jury actually *know* that the defendant is innocent, on the strength of what they have been told by the charming lawyer. The answer to this question seems to be ‘no’. If you are going to gain knowledge of a fact from an informant, the informant had better know that fact himself.

One of the leading contemporary theorists of testimony, Jennifer Lackey, uses the image of a ‘bucket brigade’ to illustrate this ‘take it from someone who knows’ condition on testimonial knowledge: In order to give you a full bucket of water, I must have a full bucket of water to pass to you. Moreover, if I give you a full bucket of water, then—spills aside—the bucket of water you now possess as a result of our exchange will also be full. The ‘spills’ here could include cases in which you don’t quite hear what I say, or cases in which someone has maliciously told you that I am a pathological liar, and you doubt me for that reason. Both in the classical and in the contemporary versions of the direct view, gaining knowledge by testimony can be blocked if you have doubts about the truth of what is being said, even if the speaker does have knowledge, and even if your own doubts are unreasonable.

Lackey herself raises doubts about the idea that testimonial knowledge arises only when we take it from someone who knows. Imagine a schoolteacher who has personal doubts about a true scientific theory that she is required to teach to her class (say, a ‘young earth’ creationist who is teaching the theory of natural selection, according to which humans have an ancient evolutionary heritage, sharing common ancestors with other primates). This teacher doesn’t know that the theory of natural selection is true—she doesn’t even believe it—but nevertheless she diligently teaches it to her class because it is part of the state curriculum. If the teacher’s trusting students come to believe the theory on the strength of their teacher’s testimony, couldn’t we say that these students now know the theory? If Lackey is right, this is a case in which someone with less than a full bucket manages to pass on more knowledge than she herself possesses.

There are other ways in which less-than-perfect sources might be able to pass on more than they individually know. One way is by working with others. Right now, I’m inclined to describe myself as knowing that the

Willamette River flows northward between the Oregon Coast Range and the Cascade Range in the American northwest. I know this because I just looked it up in Wikipedia (go ahead, double-check). It's possible that someone who actually knows this fact was the one who first updated the page about this river to report it. Perhaps someone with a full bucket has poured their knowledge through the Internet to me. It's also possible that the person who first reported the watershed boundaries had something less than a full bucket, and was, say, slightly unsure of the name of the mountain range on one side. However, over time, the entry as a whole has been vetted by so many people that the line about those mountain ranges is by now well-secured by the whole community of editors. This group may have succeeded in filling the bucket together, jointly generating an entry that is now able to provide the reader with knowledge. If the reliability of an informant is what counts, groups working together under the right conditions can outperform single authors.

This example would be handled differently by advocates of the different theories we have covered so far. A Lockean would say that what we really take away from reading the Wikipedia entry on the Willamette River is just highly probable opinion, rather than knowledge, even if there are many, many informants who are telling us something entirely plausible. A reductionist would say that any knowledge we gain here is really inferential in character: on this view, the article could transmit knowledge to those who are already aware that Wikipedia entries are generally reliable. For example, the entry could succeed in transmitting knowledge about the river's course to those who have read the 2005 *Nature* article reporting that entries in Wikipedia are comparable in accuracy to entries in the *Encyclopaedia Britannica*, or to people who have significant personal experience of double-checking the accuracy of Wikipedia entries. Meanwhile, a direct theorist of testimony could say that Wikipedia provides knowledge of the facts it reports (when its internal systems of quality control are working well), even to those who are unaware of the reliability of those systems of control. On this view, even a naive 12-year-old preparing a school report could come to know the names of the ranges flanking the river, just by reading the Willamette River entry. Traditional versions of the direct theory would require whoever wrote or edited the key sentence in the entry to know the facts it reports; however, the direct theory could also be stretched or modified to allow cases in which we gain knowledge not from a single knower but from a largely anonymous community including individuals with partial

confidence rather than full knowledge. As debate about the epistemic structure of testimony continues, new channels of information afford fresh opportunities for rival theories to offer competing explanations of the social transmission of knowledge.

Theories of knowledge generally turn to testimony only after they have examined perception and reason, but there are some philosophers who place it at the very heart of their approach to knowledge. Most notably, British philosopher Edward Craig argues that humanity came up with the concept of knowledge for the express purpose of managing the problem of testimony: we use this concept to mark people as good sources of information. Craig starts with the idea that all creatures struggling to survive need true beliefs about their environment. It helps us greatly if we are not restricted to what we have experienced personally but can also learn from others. It's imperative that we have a way of sorting out good informants, who can serve as our eyes and ears, from bad informants, who are likely to lead us astray. Good informants are identified as knowers.



Fig 7. The suffering of past generations

Craig reverses the usual direction of explanation: most epistemologists think that being a knower is something that makes you (potentially) a good informant. Craig, however, considers the notion of ‘good informant’ to be more fundamental, and vital to explaining both the value and the evolutionary origin of the concept of knowledge. Critics of Craig emphasize that knowers can sometimes be bad informants—knowers can be secretive or deceptive. In experimental settings, chimpanzees can’t distinguish between knowledgeable and ignorant informants who give them clues about where food is hidden. Chimpanzees can, however, keep track of who knows what when they are competing for resources: for example, subordinate chimpanzees are good at remembering whether a dominant animal knows where food is hidden. The connection between knowing and acting seems to be easier to spot than the connection between knowing and being a good informant. While testimony is an important topic in epistemology, it’s doubtful that it will work as our starting point.