

Fabrication and Testimony

Andreas Stokke

Swedish Collegium for Advanced Studies
Uppsala University

Penultimate draft

Forthcoming in E. Michaelson & A. Stokke (eds.), *Lying: Language, Knowledge, Ethics, Politics*, Oxford University Press

1 Introduction

A large part of what we know, we know because someone told us so. And in many cases our informant was someone who sincerely related her own knowledge. It is not surprising, therefore, that the traditional view of testimony considers testimony to be a mechanism for preserving, or transmitting, knowledge from a source to a receiver and relatedly holds that one can only obtain testimonial knowledge from sincere testifiers.

This traditional view of testimony, and of testimonial knowledge, has recently been challenged by critics who argue that testimonial knowledge may be new knowledge, that is, that testimony does not always serve merely to preserve old knowledge. Hence, according to these arguments, testimony is a generative source of knowledge. Correspondingly, these writers have claimed that testimonial knowledge can be obtained from insincere testifiers.

I agree with the critics about this conclusion. Testimonial knowledge may be new knowledge, and insincere testimony may be a source of testimonial knowledge. Yet insincerity at least sometimes blocks testimonial knowledge. Part of the motivation for the traditional view was the observation that insincere testimony is typically incapable of grounding testimonial knowledge. So we should ask, what are the conditions under which insincerity blocks testimonial knowledge and what are the conditions under which testimonial knowledge may be acquired in the face of insincere testimony?

This paper argues that when insincerity blocks testimonial knowledge, the insincerity involved is a kind of unreliability. I argue that acquiring testimonial knowledge requires that the testimony be given on a reliable basis. That is, testimonial knowledge relies on the testifier passing on information from a reliable source. I show that, in cases where listeners acquire testimonial knowledge from insincere testifiers, the testimony is given on a reliable basis, whereas in cases where insincerity prevents acquisition of testimonial knowledge, the reason is that the testimony is given on an unreliable basis. Specifically, insincere testimony - in particular, lying - typically involves what I call *fabrication*, that is, making something up. Fabrication, in this sense, is not a reliable basis for testimony, and hence this explains why lying testimony typically does not yield testimonial knowledge.

An alternative is to think of insincerity as a potential *defeater* for testimonially based belief, that is, as a factor that can undermine or rebut the subject's warrant or justification for the relevant belief. One can distinguish views according to which insincerity on the part of the testifier is what is often called a "factual" defeater for a belief based on her testimony and views according to which insincerity is a potential "psychological" defeater for such a belief.¹ Roughly, according to the former notion, insincerity on the part of the testifier defeats justification regardless of whether the recipient is aware of the insincerity, while according to the latter, insincerity defeats justification only if the recipient believes or has reason to believe that the testifier was insincere. Depending on how one thinks of these issues, one can add a suitable requirement of no undefeated defeaters on testimonial knowledge. I will not discuss justification or warrant for testimonially based belief in this paper. The omission of this factor in testimonial knowledge, and others, is the reason I confine myself to discussing necessary, but not sufficient, conditions on testimonial knowledge in what follows.

Section 2 reviews the traditional view of testimony and the reasons for rejecting it. In Section 3 I argue that testimonial knowledge requires testimony given on a reliable basis, and I show that lying often involves fabrication, and hence testifying on an unreliable basis. Section 4 argues that the requirement of reliable bases is superior to an alternative view according to which testimonial knowledge merely requires reliable testimony.

2 The Preservative View and its Critics

¹ See Lackey (2008: 44-45)

2.1 Transmission and Sincerity

A common way of spelling out the traditional view is as endorsing a principle like the following:

Transmission Condition

A knows *p* testimonially on the basis of *B*'s testimony that *p* only if *B* knew *p* at the time of testifying.

This necessary condition on testimonial knowledge, or some version of it, has been endorsed by, among others, Keith Lehrer (1990), Robert Audi (1997), (2002), (2011), Tyler Burge (1993), Alvin Plantinga (1993), Micheal Dummett (1994), John McDowell (1994), David Welbourne (1994), and David Owens (2006). Even though all of these writers on testimony recognize that there are cases in which the condition is not met, they all agree that testimony is, at least paradigmatically, a preservative source of knowledge. Audi states the view succinctly as follows:

Testimonially based knowledge is received by transmission and so depends on the attester's knowing that *p*. (Audi 1997: 410)

In turn, the Transmission Condition is often taken to motivate a Sincerity Condition on testimonial knowledge of the following kind:

Sincerity Condition

A knows *p* testimonially on the basis of *B*'s testimony that *p* only if *B*'s testimony that *p* was sincere.

In describing the traditional view, which she rejects, Jennifer Lackey (2008) writes,

[N]early everyone takes sincerity on the part of the speaker to be a necessary condition for testimonial knowledge. In order to properly learn from a speaker's belief, there needs to *be* a belief present from which to learn. Thus, if a speaker is insincere and expresses what she herself does not believe, there is nothing for her to pass on to a hearer. (Lackey 2008: 38)

The strength of the connection between these two ideas depends on how one thinks of sincerity. If one thinks that sincerity is simply a matter of saying something one believes, then the Transmission Condition entails the Sincerity Condition, since knowledge entails belief. There are reasons to think that

sincerity is not always just a matter of saying what one believes.² Still, it is arguably true that, in many cases, saying something one believes is to speak sincerely. So we may take the Sincerity Condition to be at least strongly motivated by the Transmission Condition.

Both the Transmission Condition and the Sincerity Condition remain unspecific on what it is to testify that p to someone. Some philosophers, such as Robert Audi (2011), think that testifying that p must involve *asserting* that p , that is, roughly, saying that p in an affirmative manner. Others, such as Jennifer Lackey (2008), allow that one can testify that p in other ways, for example, by conversationally implicating that p . I do not take a stand on this debate here. In what follows, we will be concerned only with cases where someone testifies that p by asserting that p .

The primary kind of insincerity in speech is lying, even though insincerity in speech also encompasses other things, such as falsely implicating.³ To lie, according to the view I favor, is to assert something one believes to be false.⁴ Consequently, we will be concerned here only with cases in which someone testifies insincerely by lying, that is, by asserting that p while believing not- p .

2.1 Generating Knowledge

According to the Transmission Condition, to have testimonial knowledge requires having been told by someone who themselves knew the relevant proposition at the time. This view is motivated by everyday cases in which what one knows by testimony, one was told by someone who knew it themselves. One's source might in turn also have been told by someone else, who may themselves have been told by someone else in their turn, etc. In simple cases of such chains of testimony, each link in the chain testified while knowing the relevant proposition.

Yet reflection on one's own testimonial knowledge quickly suggests that not all of it has come from chains in which every link shared the knowledge in question.⁵ It seems wildly implausible to think that, of the vast number of things we know because we have been told, not one of them has reached us

² See Stokke (2014). Given that one may know p while believing that one does not believe p , one may arguably be insincere in asserting p , even though one knows p . If so, the Transmission Condition does not entail the Sincerity Condition.

³ Stokke (2014).

⁴ See Stokke (2013).

⁵ Dummett (1994: 265) makes this point.

through a chain in which some testifier, for whatever reason, did not know the proposition in question at the time of testifying.

A striking example is our knowledge about the distant past. For example, I take myself to know that Charlemagne was crowned Holy Roman Emperor in 800. Yet it seems incredible to assume that along the tortuous road by which this information has been passed on - presumably through manuscripts that have been copied and translated many times, through word of mouth, and a myriad of other connections - all the testifiers on the way have known it.⁶ It is surely too optimistic to think that not one of these hundreds of individuals testified while falling short of knowing. For example, they may have failed to believe the information in question, they may have been unaware of undefeated defeaters, they may have lacked justification, and so on.⁷

As an alternative, therefore, it might seem plausible to suggest that testimonial knowledge merely requires that the first link in the chain of testimony knew.⁸ For example, it might be thought that for me to know testimonially that Charlemagne was crowned Holy Roman Emperor in 800, the chain must have started from a first testifier, perhaps an eye-witness, who passed on her own knowledge about Charlemagne's coronation.

As this suggests, this view implies that the first link in the chain knew the relevant proposition non-testimoniaally. Suppose you know p on the basis of the testimony of an individual A_n who is the n th link in a chain of testifiers, in which A_1 is the first link, that is, the first person to testify that p . According to the view under consideration, since you know p , A_1 knew p . Could A_1 have known p non-testimoniaally? According to this view, she could not. For if so, this view implies that there must have been a prior testifier, say A_0 , who knew p and on whose testimony A_1 's knowledge that p is based. The same argument applies to A_1 and so on until the buck stops at a first testifier who knew p on some basis other than testimony.

Hence, according to the knowledge-preservative view of testimony, even if it is merely required that testimonial knowledge be based on a chain of testimony in which the first link knew, what is passed on by testimony can only be

⁶ Perhaps we can distinguish between links in the chain that are testifiers and links that merely pass on information. For example, translators might be of the latter sort, or more extremely, xerox machines, or the like. But even if one screens off such non-testifier links, the point still stands: it is implausible to think that not one of the genuinely testifying links in the chain testified while not knowing.

⁷ See Lackey (2008) for a range of detailed examples of this kind.

⁸ Dummett (1994: 265) suggests a view of this kind. See also Graham (2006: 108-109), Faulkner (2009: 481-482).

knowledge that originally started off as non-testimonial knowledge. For example, the first link in the chain may have come to know that p by perception, or by inference, or from some other non-testimonial source. Hence, on this view, testimony is a mechanism for transmitting, or preserving, knowledge that started off as non-testimonial knowledge.

It has been argued that there are instances of testimonial knowledge where there are no prior knowing links. Here is a version of an example given by Peter Graham (2006):⁹

Fossil

Jones is a schoolteacher with an extensive familiarity with evolutionary theory. Although privately a creationist, he is diligent and believes in the importance of teaching the curriculum. On a field trip he finds a fossil and correctly deduces that it is a fossil of a creature that lived in the area millions of years ago. So he tells the students that the fossil is the remains of a creature that lived in the area millions of years ago.

Jones does not believe what he is testifying to, that is, that the fossil is the remains of a creature that lived in the area millions of years ago. Hence, at least in this sense, Jones's testimony is insincere. Moreover, in testifying to the students, Jones is asserting something he believes to be false. So, at least in the generic sense, Jones is lying to the students.

In support of this judgment, suppose that, outside school, Jones is talking privately to one of his relatives who shares his religious beliefs. The relative might confront Jones with, "So you are up there telling these kids all these things you know to be false. You are lying to them!" Given that Jones truly does believe that the things he tells the students are false, he is obliged to agree with this description.¹⁰ To be sure, there may be ways in which Jones can argue that his lies (as he thinks of them) are permissible lies, since they are what the curriculum requires him to teach the students. But this does not mean that he is not lying. To assert something one is aware of disagreeing with, in the sense of

⁹ This example is a strengthening of Lackey's (2008, 48) much-discussed "Creationist Teacher" case. See also Plantinga (1993: 87) for a different kind of case of testimonial knowledge from a chain in which no one has testimonial knowledge of the relevant proposition.

¹⁰ Thomas Carson (2010) thinks that you lie only if what you assert is false. On this view, Jones is not lying, when what he tells the students is true. However, the view that lying requires asserting something false - as opposed to merely asserting something one believes to be false - is an outlier view, and is rejected by most other views of lying. See, for example, Fallis (2009), Stokke (2013).

it being something one would dissent from if asked, is to lie, even if one can be excused for doing so.¹¹

Since Jones does not believe what he testifies to, he does not know it. But nevertheless, we want to say that the students come to know what Jones tells them. Being told by Jones, in this way, is a perfectly good way of coming to know that the fossil is the remains of a creature that lived in the area millions of years ago. So examples like this one show that both the Sincerity Condition (understood as involving belief) and the Transmission Condition are false.

But moreover, examples like Fossil show that even the weak version of the knowledge-preservative view of testimony, on which testimonial knowledge merely requires a first knowing link in the chain, is incorrect. As Graham points out, on the intended reading of the example,

Since the children come to know something no-one has ever known before, they are the first to know. Testimony has *generated* knowledge. (Graham 2006: 113)

In other words, there are chains of testimony that sustain testimonial knowledge, but in which the first link does not know. And, as the Fossil example makes vivid, there are chains that yield knowledge but in which no prior links knew the relevant propositions. Hence, testimony is wrongly characterized as a mechanism for preserving old knowledge. Testimonial knowledge may be new knowledge.

3 Insincerity and Bases for Testimony

3.1 Preserving Information and Reliable Bases

One of Graham's central observations is that, although testimony is not a mechanism that functions to transmit knowledge, testimony is a mechanism for transmitting information. However, Graham's (2000), (2006) view is couched in terms of a specific theory of information, inherited from Fred Dretske (1981), according to which information carrying is understood as counterfactual dependence between a signal and an event. On this view, information *per se* constitutes grounds for knowledge, and hence testimony is seen as a mechanism for transmitting grounds for knowledge.

¹¹ Stokke (2014) argues that insincerity involves communicating that *p* while mentally assenting to not-*p*. Hence, on this view, if one is unconscious of believing not-*p* while one mentally assents to *p*, one is not insincere in asserting *p*.

Instead, I propose to think of information in a more intuitive sense. On this way of thinking about information, a piece of information is simply a proposition.¹² Doing so allows us to point out that the core of Graham's observation about the information-preservative nature of testimony is independent of any particular commitment concerning the nature of information-carrying and the relation between information and knowledge. Indeed, any theory of testimony should acknowledge the simple point that testimony is preservative with respect to information, regardless of whether it is sympathetic to a view of knowledge as, in Graham's technical sense, "belief based on the information that P." (Graham 2006: 118)

Even on a loose understanding of information, testimony works by passing on information from a testifier to a receiver. In this respect, traditional theorists are right to classify testimony as being like memory and unlike perception.¹³ Even if testimony is not a mechanism for transmitting knowledge from a source to a receiver, testimony is a mechanism for transmitting information from a source to a receiver. Graham writes,

Perception extracts information from the world. [...] Testimony, on the other hand, conveys or transmits information that a speaker has already extracted via perception (or by some first-person means). (Graham 2006: 121)

What I want to draw attention to is that, since testimony is a mechanism for transmitting information, testimony is always taken from a source, or as I will say, testimony is always *given on a basis*.

Consequently, there is a parallel with the traditional view according to which testimony is preservative with respect to knowledge. In particular, information passed along in a chain of testimony must originally have entered the chain from a non-testimonial source. This source may itself be an information-

¹² This view is compatible with different accounts of the nature of propositions. One might think of propositions as structured entities, as sets of possible worlds, or in some other way. All that is assumed here is that there is a natural sense in which the information that *p* is to be analyzed in terms of the proposition that *p*, rather than, for example, in terms of a relation between signals and events.

¹³ See Dummett (1994), Audi (2002) for views on which memory is knowledge-preserving. See the appendix to Lackey (2008) for a knowledge-generative view of memory. It seems plausible to think, as Lackey does, that one may know *p* on the basis of having remembered that *p*, even though no-one, including oneself, has previously known that *p*. Yet, as with testimony, even if memory is ultimately to be regarded as generative with respect to knowledge, this does not detract from the plausibility of taking memory to be preservative with respect to information. But see Dummett (1994: 268) for some discussion.

preserving source, such as memory. But, ultimately, the first source for the relevant information must be capable of generating - or extracting, to use Graham's phrase - new information. For any proposition p that is passed along through a chain of testimony, the buck must stop at a first testifier who transmitted the information that p from a non-testimonial source of information.

It is straightforward to illustrate this by ordinary cases. For example, you look out the window and see that it is snowing. You report this to me, and I thereby come to have that information, too. In ordinary circumstances, moreover, I will come to believe that it is snowing on the basis of your testimony. In other cases, I may not come to believe what you say, but I will nevertheless be in a position to entertain the proposition that it is snowing because you told me so.¹⁴ Your testimony was given on the basis of your perceptual experience. That experience was your source for the information that it is snowing. In turn, my source for the same information is your testimony. If I subsequently tell my friend that it is snowing, she can come to have the information in her turn. While you testified on the basis of perception, I testified on the basis of your testimony to me. Both our acts of testifying were acts of passing on information.

Given that testimony is always given on a basis, a natural thought is that the quality of the testifier's source of information plays a role in whether her testimony can support knowledge of what she testifies to. Indeed, testimony given on an unreliable basis, even if true, typically does not yield knowledge. Here is an example:

Mail

Mitch wants to know whether there's any mail for him at the office. He calls the front desk, and Marvin answers. When asked whether there's any mail for Mitch, lazy Marvin decides to flip a coin, letting heads be "yes" and tails be "no." He flips the coin, gets tails, and tells Mitch, "No, there's no mail for you." As it happens, what Marvin said was true; there is no mail for Mitch.

¹⁴ According to a long tradition, reaching back to Frege (1997 [1918]), Geach (1965), and others, propositions can be grasped, or entertained, without being judged, asserted, or believed. This observation was part of the motivation for the distinction between force and content in the Fregean tradition. An alternative view, recently defended by Hanks (2015), rejects the idea of a neutral attitude towards propositions, and correspondingly rejects the force-content distinction. I am sympathetic to the former view, but note that, even if one rejects the force-content distinction, one should acknowledge that propositions may occur in thought or speech without being believed or asserted, as for example when merely supposed or hypothesized. Hanks (2015: ch 4) argues that these are cases of force-cancellation. See Soames (2010) for a position that shares many of the features of Hanks's view, but allows for a notion of neutrally entertaining propositions.

Mitch does not come to know that there is no mail for him at the office in this case, even though he does come to have a true belief. A plausible explanation for this is that Marvin testified on an unreliable basis. The basis for Marvin's testimony, that is, the coin flip plus his decision on how to interpret heads and tails, is not a reliable guide to the facts.

Is Marvin lying? It is natural to think that Marvin does not believe that there is mail for Mitch, nor does he believe that there is no mail for Mitch. In other words, Marvin is most likely agnostic about what he testifies to. For this reason, he does not count as lying on the generic view according to which to lie is to assert something one believes to be false. Instead, it might be more plausible to think of Marvin's assertion as a case of *bullshitting*, in Harry Frankfurt's (2005 [1986]) sense of asserting something while being indifferent towards whether what one says is true or false. Marvin does not care about whether what he tells Mitch is true or false.¹⁵

Regardless of the issue of how to classify Marvin's assertion - as lying, bullshitting, or perhaps something else - it is clear that Marvin's testimony is given on an unreliable basis. This I take to be the reason his testimony does not support knowledge of what he testifies to, even though it is true. While Marvin's testimony is true, and Mitch believes him, Mitch's belief does not constitute knowledge because what Marvin tells him is taken from an unreliable source of information. By contrast, when one acquires testimonial knowledge, the testimony one receives is given on a reliable basis. When you tell me that it is snowing outside on the basis of having seen that it is snowing outside, your testimony is given on a reliable basis. That is a plausible explanation for why I can come to know that it is snowing outside on the basis of your testimony.

Motivated by these considerations, I propose the following condition on testimonial knowledge.

Reliable Basis Condition

A knows *p* testimonially on the basis of *B*'s testimony that *p* only if *B*'s testimony that *p* was given on a reliable basis.

One can impose refinements on this view in order to take into account chains of testimony. I will forego discussion of this here, though, and simply focus on simple cases of just one testifier.

¹⁵ See Stokke and Fallis (2016) for a view on which asserting *p* while being agnostic about *p* is typically a way of bullshitting, as opposed to lying.

The Reliable Basis Condition goes some way in explaining the role of bases, or sources of information, in supporting testimonial knowledge. We will see next that the Reliable Basis Condition also helps explain how insincerity sometimes prevents testimonial knowledge and sometimes does not.

3.2 Lies and Fabrication

Examples like Graham's Fossil case show that one can acquire testimonial knowledge from insincere testifiers. But we should nevertheless be interested in investigating the way insincerity may influence the acquisition of testimonial knowledge. In many cases lying involves asserting something false. In these cases it is trivial that one cannot acquire knowledge of what is said. If you ask me what time it is, and I tell you it is 1 pm even though I know it is 2 pm in order to make you late for a meeting, you cannot come to know what I told you, since what I told you was false.

Yet lying can block testimonial knowledge, even when what is testified to is true. This can be seen by considering cases like the following:

Slander

Sue hates her boss and wants to give Bob a bad impression of her. So, while Sue in fact believes that her boss would never do anything of this sort, she makes up the story that her boss moved funds illegally, and she tells Bob this. Bob has no reason to think that Sue is lying, and so as a result of her testimony comes to believe that her boss moved funds illegally. As it turns out, the story is true - Sue's boss did move funds illegally.

Bob does not end up knowing that Sue's boss moved funds illegally in this case, even if he does end up with a true belief that she did.

Even though what Sue says is true, Sue's utterance in Slander is a paradigm case of lying. By suggesting that insincerity itself prevents testimonial knowledge, in this way, cases like Slander arguably constitute a central part of the motivation for the Sincerity Condition. More particularly, it is often pointed out that, unlike sources of knowledge like perception, testimony involves the voluntary actions of another agent. In turn, this is sometimes seen as motivating the view that one cannot get testimonial knowledge from someone who lies. For example, Robert Audi writes,

There is a sense in which testimonially based belief passes through the will - or at least through agency: the attester must select what to attest to and in the process can also lie, in which case the belief does not constitute knowledge [...]. (Audi 2002: 79)

Recognition of this feature of testimony, along with observation of cases like Slander, motivates the Sincerity Condition.

As we have seen, however, the Sincerity Condition is false. Yet, even though testimonial knowledge can be acquired from insincere testifiers, the Reliable Basis Condition captures the motivation for the traditional Sincerity Condition. In particular, a plausible explanation for why Bob does not come to know in Slander is that Sue testified on an unreliable basis. What is Sue's basis for testimony in Slander? Put differently, what is her source for the story that her boss moved funds illegally? The intuitive answer is that Sue *made up* the story about her boss. As I will say, in such cases, the speaker's source, and hence her basis for testimony, is *fabrication*.

By "fabrication" I mean the everyday phenomenon of making something up. I will not offer a full account of the nature of fabrication here. All I am assuming is the obvious truth that fabrication - making things up - is among our capacities. Indeed, Slander exemplifies a common and well known phenomenon. Yet it is worth remarking briefly on one aspect of fabrication, in the sense of making something up. In particular, I want to note that fabrication, in this sense, is not straightforwardly identifiable with imagination, at least as the latter has been understood by some philosophers, even though imagination is arguably a central component of fabrication.

For example, Timothy Williamson (2007) argues that one can acquire knowledge of metaphysical modalities by using imagination to evaluate counterfactuals, an ability that he argues is "moderately reliable." (Williamson 2007: 155) However, Williamson also acknowledges that imagination can operate in an unconstrained way as when one imagines a rock rising vertically into the air. I am not suggesting that fabrication, of the kind I have in mind in this paper, can be equated with this kind of less constrained imagination, nor am I suggesting that fabrication is not constrained in many ways.

From a slightly different perspective, Elizabeth Camp (2009: 107) distinguishes between "make-believe" as the attitude of "supposing a set of propositions to be true," and "imaging" as an "experiential state of imagining a scenario as if it were before one [...]." On this terminology, the way I use "fabrication" bears similarities to the former category. However, fabrication - as exemplified by the

Slander case - is not always accompanied by a supposition that what one makes up is true. In other words, fabrication is not identifiable with make-believe, in Camp's sense, since, for example, Sue does not pretend or suppose that her story about her boss is true.

Fabrication - the capacity to make things up - is a generative source of information. One can make up things that no one has ever thought of before. But, unlike information-generative sources like perception, fabrication is not a reliable source of information. Since Bob's belief, in Slander, is based on Sue's testimony, the unreliability of fabrication as a basis for testimony provides a plausible explanation for why he does not come to know what she tells him.

For convenience, we can think of reliability, roughly, in terms of being right in nearby situations.¹⁶ Then here is a more precise way of explaining the case. There are nearby situations in which Sue makes up the story that her boss moved funds illegally but in which the story is false.¹⁷ The story could easily have been false. Sue's capacity for fabrication is an unreliable source of information. Since Sue's testimony about her boss is actually based on fabrication, we may assume that situations in which Sue testifies on the basis of fabrication of the story about her boss are not more far-fetched than situations in which the story is false. So there are nearby situations in which Sue testifies that her boss moved funds illegally but in which the story is false. Since Bob's belief is actually based on Sue's testimony, we may assume that situations in which his belief is based on her testimony are not more far-fetched than situations in which her testimony is based on her fabrication but false. So there are nearby situations in which Bob bases his belief on Sue's testimony but in which the belief is false.

More simply, since the story that Sue made up could easily have been wrong, her testimony could easily have been wrong, and hence so could Bob's belief. Hence, the Reliable Basis Condition explains why one typically cannot get knowledge from insincere testifiers. The reason is that insincerity - in particular, lying - typically involves fabrication.

¹⁶ This way of thinking about reliability was proposed by Sainsbury (1997). Subsequently it has been spelled out as what Sosa (1999) called *safety*. As a condition on knowledge, safety has been endorsed by, among others, Sosa (1999), Williamson (2000), Pritchard (2005) .

¹⁷ On a view like Graham's (2000), (2006), Sue's testimony does not carry the information that her boss moved funds illegally, since, on this view, a signal r carries the information that p if and only if in nearby situations, if r , then p . Hence, Graham will explain the Slander case by saying that Bob's belief is not based on information. By contrast, on my view, it is both more intuitive and less theoretically committing to acknowledge that Sue's testimony is a source of information, namely the information about her boss, and instead explain the failure of knowledge by observing that Sue's source for the information, that is, fabrication, is unreliable.

Moreover, the Reliable Basis Condition explains why one can sometimes get testimonial knowledge from insincere testimony. Consider the Fossil case. A plausible explanation for why Jones's testimony yields knowledge is that it is given on a reliable basis. What is Jones's basis for testimony? What is his source for the information that the fossil is the remains of a creature that lived in the area millions of years ago? Intuitively, Jones's source for this information is a valid and sound inference he makes from evolutionary theory plus his observation of the fossil evidence. This way of extracting information is a reliable guide to the facts. Hence, we explain why the students come to know by pointing to the fact that, despite its insincerity, Jones's testimony is given on a reliable basis.

The Reliable Basis Condition is motivated, on the one hand, by the information-preservative nature of testimony, and on the other hand, by facts about when insincerity does and does not block testimonial knowledge. In the next section I argue that it is superior to an alternative, weaker, way of construing the reliability condition on testimonial knowledge.

4 Reliability and Testimonial Knowledge

4.1 Reliable Testimony

Some critics of the traditional view of testimony have argued for an alternative to the Reliable Basis Condition. For example, Lackey argues that

What really matters for the epistemic status of testimony is whether the speaker is a *competent testifier*, where this is understood in terms of the reliability - or other form of truth-conduciveness - of *the statement* in question. (Lackey 2008: 73-74)

This position can be seen as endorsing the following condition on testimonial knowledge:

Reliable Testimony Condition

A knows *p* testimonially on the basis of *B*'s testimony that *p* only if *B*'s testimony that *p* was reliable.

The Slander case fails the Reliable Testimony Condition. As we said, there are nearby situations in which Sue testifies that her boss moved funds illegally, but in which the story is false. Hence, Sue's testimony is not reliable, and hence the Reliable Testimony Condition implies that Bob does not come to know. And

similarly, in the Mail case. Marvin's testimony is unreliable, and hence Mitch does not come to know. Moreover, this view also makes the right prediction in the Fossil case. Jones's testimony is reliable - it is true and it could not easily have been false.

So why should we opt for the Reliable Basis Condition over the Reliable Testimony Condition? In order to answer this question, we need to consider the difference between the two views more carefully.

We have seen that, since testimony is preservative with respect to information, testimony is always given on a basis, that is, testimony is always taken from a source. Correspondingly, we explained the unreliability of Sue's testimony in Slander by pointing to the fact that it is based on an unreliable source, that is, fabrication. According to the basis reliabilist, this is the explanation for why Bob does not come to know. By contrast, the testimonial reliabilist can be seen as claiming that we do not need this last step in order to explain why knowledge fails in Slander. Instead, the testimonial reliabilist wants to stop at the fact that Sue's testimony is unreliable.

Generally, the testimonial reliabilist can be seen as claiming that the basis for testimony does not matter for whether it can yield testimonial knowledge. What matters is whether the testimony itself is reliable. As Lackey puts it in the passage quoted above, what matters is whether "the statement" is reliable. One way of looking at this view, then, is to see it as maintaining that we do not need to be as fine-grained as to consider bases for testimony in order to predict when testimony sustains testimonial knowledge and when it does not.

However, as we will see next, this view has problems accounting for a particular kind of case.

4.2 Testifying Reliably on Unreliable Bases

One can testify reliably on unreliable bases. In such cases, testimonial knowledge is typically blocked. Intuitively, in these cases, the speaker is not a competent testifier. Here is an example, inspired by a similar case given by Steven Luper-Foy (1984):

Lights

Each day before Larry goes home from the office, his boss calls him to check whether the lights are on in the hallway. Sometimes, the lights are on; sometimes they're off. In the morning, the boss comes in so late that, if the lights are on when she shows up, she has no way of knowing

whether the lights have been on all night or whether they've been switched on earlier in the morning. After months of this, Larry gets too lazy to go over and check the lights when the boss calls. Instead, each time, Larry flips a coin, having decided once and for all that heads is going to mean the lights are on and tails that they're off. At the same time, unbeknownst to Larry, Melvin, who is keen to prevent Larry from getting into trouble, is watching his coin-flips and is secretly adjusting the lights, so that each day, what Larry tells his boss is true.

Many will think that Larry's boss does not get testimonial knowledge about the lights from Larry's testimony in this case. There is a challenge, then, for the testimonial reliabilist to explain why the boss does not get testimonial knowledge in Lights without appealing to the basis for Larry's testimony.

In what follows, I will consider what I take to be the two most prominent proposals for meeting this challenge on behalf of the proponent of the Reliable Testimony Condition.

4.3 Ruling Out Accidental Correctness

One proposal is to claim that Larry's testimony, although reliable on this particular occasion, is not reliable *in itself*, and therefore it does not yield knowledge. As motivation for this suggestion, one might appeal to the common idea that a basis for belief is reliable in itself only if it rules out accidental correctness *in normal circumstances*.¹⁸ For example, Luper-Foy writes,

If we come to a true belief via a completely unreliable method such as by tossing a coin, our belief is correct only coincidentally even if we happen to be in cooperative circumstances. The reason is that as far as the method is concerned, it is sheer coincidence that we are in cooperative conditions. Therefore, even when we are in (abnormally) cooperative ones, the processes through which we know must be at least as reliable as is required to eliminate the possibility of accidentally correct belief in normal circumstances. (Luper-Foy 1984: 32)

Analogously, one proposal is that Larry's testimony in Lights does not rule out accidental correctness in normal circumstances, that is, without Melvin's correction of the lights. As far as Larry's testimony is concerned, it is coincidental that he is in cooperative circumstances. Hence, it might be argued, the boss's belief is not held on a reliable basis.

¹⁸ A similar view of reliability is endorsed by Graham (2000).

However, to make this argument, the testimonial reliabilist needs to specify the range of circumstances in which the reliability of Larry's testimony is to be measured. In other words, the testimonial reliabilist faces the question, what are the normal circumstances in this case? That is, what are the circumstances in which Larry's testimony "in itself" does not rule out accidental correctness?

The answer must be that the normal circumstances are those in which Larry testifies about the lights *on the basis of his coin-flip* without Melvin's correction. If it is not specified that, in the relevant cases, Larry's testimony is based on the coin-flip, there is no explanation for why his testimony, in those cases, does not rule out accidental correctness.

Yet this is just to say that the reason Larry's testimony is not reliable "in itself" is because it is based on the coin-flip. In other words, this argument ultimately explains the lack of knowledge by appealing to the basis for Larry's testimony. Hence, taking this route is simply to sneak in the Reliable Basis Condition by the backdoor.

4.4 Knowledge and Testimonial Knowledge

Another response on behalf of the testimonial reliabilist to cases like Lights is to reject the challenge. One can claim that Larry's boss *does* get knowledge from Larry's testimony. As motivation, one can appeal to the fact that the combination of Larry's testimony and Melvin's correction of the lights establishes a reliable link between the facts and the boss's belief, and insist that this is sufficient for it to establish knowledge.

Some will disagree with the claim that Larry's boss comes to know. Others will think that there is intuitive support for the claim that the boss does get knowledge in Lights. But as I argue below, even if one is in the latter camp, one should not give up the Reliable Basis Condition on testimonial knowledge.

If one thinks that Larry's boss comes to know, one's motivation is likely to be that her belief is based on Larry's testimony about the lights, which, as corrected by Melvin, is a reliable indicator of the facts. The fact that Larry is testifying on the basis of the coin-flip and that Melvin is diligent enough in his corrections of the lights is what establishes the sense in which the boss's belief is true and could not easily have been false.

More particularly, we can imagine that situations in which Melvin corrects the lights are not more far-fetched than situations in which the boss bases her belief

on Larry's testimony. If required, we can flesh out the case so that this becomes more plausible, for example, by specifying more about Melvin's reasons for correcting the lights, etc. In other words, it is hard to deny that there are examples of this kind in which, in nearby situations where the boss bases her belief on Larry's testimony, the belief is true. In such cases, the boss's belief is true and could not easily have been false. This is at least one reason for thinking that the boss, in these cases, does come to know.

However, given that it is the combination of Larry's testimony and Melvin's correction of the lights that secures the reliability of the boss's belief, it can be argued that the boss's knowledge is not an instance of testimonial knowledge *per se*. The epistemology of testimony is interested in accounting for the phenomenon of knowledge based on the testimony of other people. A case in which knowledge is based on someone's correction of certain features of the environment to make them match what someone says is not an instance of this phenomenon.

More particularly, all theories of the epistemology of testimony must distinguish between testimonial knowledge and what is sometimes called knowledge "by way of" testimony.¹⁹ For example, Audi writes,

your attesting to p could cause a machine to produce the belief that p (perhaps even knowledge that p) directly in me; but this would at best be a case of knowledge *due to*, not *on the basis of*, testimony. A mere cause of my knowing something is not a source of knowledge. (Audi 2002: 79)

Standardly, theories of testimony implement this distinction by imposing the condition that a belief amounts to testimonial knowledge only if it is held solely on the basis of the content of the testimony in question.²⁰ Call this the Content Condition:

Content Condition

A knows p testimonially on the basis of B 's testimony that p only if A 's belief that p is held solely on the basis of the content of B 's testimony that p .

¹⁹ See, for example, Audi (1997), (2002), Lackey (2008: 55–56).

²⁰ See, for example, Audi (1997), Lackey (2008: 55–56). There are complications here that one will want to correct for in a full version of the view. For example the Content Condition rules out as testimonial knowledge for which the corresponding belief is held both on testimonial and non-testimonial grounds, such as Plantinga's (1993: 87) Australia case. I refrain from discussing this here, but a final version of this kind of principle should spell out more precisely the way in which testimonial knowledge need be based on the content of the relevant testimony. Thanks to an anonymous referee for Oxford University Press for pressing this point.

The Content Condition correctly predicts that the belief-inducing machine in Audi's example does not give rise to testimonial knowledge. The belief, in this case, is not based on the content of the testimony, but rather it is based on the machine's input.

Audi's example is intended to illustrate factors of testimonial knowledge having to do with the pedigree of the relevant belief.²¹ Correspondingly, the requirement specified by the Content Condition that a belief that is to count as testimonial knowledge be held solely on the basis of the content of the testimony is intended to make sure that the relevant beliefs are based on the content of the testimony, that is, roughly, on understanding the testimony in question.²²

An additional suggestion is that testimonial knowledge, as opposed to knowledge merely by way of testimony, is had only if the reliability of the relevant belief is due to the reliability of the testimony in itself, in the sense we articulated earlier. If one's belief is based on testimony, and yet the reliability of one's belief is not due solely to the testimony's being reliable in the sense of ruling out accidental correctness in normal cases, but instead is due to some factor external to the testimony itself, one's knowledge is not of the kind that is to be captured by a theory of testimonial knowledge.

We have already seen that, in order to articulate the notion of testimony being reliable in itself, one needs to appeal to the reliability of the basis for testimony. Hence, the suggestion from above amounts to the claim that testimonial knowledge requires that the reliability of the relevant belief be due to the reliability of the basis on which the testimony is given. This condition is not satisfied in Lights. Hence, even if one thinks that Larry's boss comes to know, one can maintain that her knowledge is not testimonial knowledge, since one reason why she knows - the reliability of her belief - is not explained solely by features of the testimony itself.

If the testimonial reliabilist wants to argue that Larry's boss does get knowledge in Lights, the basis reliabilist can agree. But, in addition, one can plausibly claim that the boss's knowledge is not testimonial knowledge, the reason being that it is not based on testimony that is reliable in itself. Put differently, testimonial knowledge, as opposed to knowledge merely by way of testimony, requires that the testimony was given on a reliable basis.

²¹ See Lackey (2008: 42) for another example.

²² This is what Graham (this volume) calls "acceptance," or "comprehension-based beliefs."

5 Conclusion

Testimony is not a mechanism for transmitting knowledge, but a mechanism for transmitting information. For this reason, testimony is always given on a basis. Testimonial knowledge, as a consequence, requires that the testimony be given on a reliable basis. Many cases of insincere testimony involves testifying on unreliable bases. Most prominently, lying often involves testifying on the basis of fabrication. This explains why one typically cannot get knowledge from lying testifiers. Moreover, in complex cases, the requirement of reliable basis for testimony may distinguish testimonial knowledge from knowledge by way of testimony.

Acknowledgements

I am grateful to Daniele Sgaravatti, Eliot Michaelson, Federico Luzzi, Jessica Brown, Kirstine La Cour, Mikkel Gerken, Torfinn Huvenes, and two anonymous referees for Oxford University Press for valuable suggestions and discussion on this and earlier versions.

References

- Audi, R. (1997), 'The place of testimony in the fabric of knowledge and justification', in *American Philosophical Quarterly* 34/4: 405–422.
- Audi, R. (2002), 'The sources of knowledge', in P. Moser (ed.), *The Oxford handbook of epistemology* (Oxford University Press), 71–94
- Audi, R', (2011), 'Testimony as a social foundation of knowledge', in *Philosophy and Phenomenological Research* 87/3: 507-531.
- Burge, T. (1993), 'Content preservation', in *Philosophical Review* 102: 457–488.
- Camp, E. (2009), 'Two varieties of literary imagination: Metaphor, fiction, and thought experiments', in *Midwest Studies in Philosophy* XXXIII: 107–130.
- Carson, T. (2010), *Lying and deception: Theory and practice* (Oxford University Press).

- Dretske, F. (1981), *Knowledge and the flow of information* (MIT Press).
- Dummett, M. (1994), 'Testimony and memory', in B. K. Matilal & A. Chakrabarti (eds.), *Knowing from words* (Kluwer), 251–272.
- Fallis, D. (2009), 'What is lying?', in *Journal of Philosophy* 106: 29–56.
- Faulkner, P. (2009), 'Review of: *Learning from Words: Testimony as a Source of Knowledge* by J. Lackey', in *Mind*, 118/470: 479–485.
- Frankfurt, H. (2005 [1986]), *On bullshit* (Princeton University Press).
- Frege, G. (1997 [1918]), 'Thought', in M. Beaney (ed.), *The Frege reader* (Blackwell), 325–345.
- Geach, P. (1965), 'Assertion', in *The Philosophical Review* 74: 449–465.
- Graham, P. (2000), 'Transferring knowledge', in *Noûs*, 34/1: 131–152.
- Graham, P. (2006), 'Can testimony generate knowledge?' in *Philosophica* 78: 105–127.
- Hanks, P. (2015), *Propositional content* (Oxford University Press).
- Lackey, J. (2008), *Learning from words* (Oxford University Press).
- Lehrer, K. (1990), *Theory of knowledge* (Westview Press).
- Luper-Foy, S. (1984), 'The epistemic predicament: Knowledge, Nozickian tracking, and scepticism', in *Australasian Journal of Philosophy* 62/1: 26–49.
- McDowell, J. (1994), 'Knowledge by hearsay', in B. Matilal & A. Chakrabarti (eds.), *Knowing from words* (Kluwer), 195–224.
- Owens, D. (2006), 'Testimony and assertion', in *Philosophical Studies* 130: 105–129.
- Plantinga, A. (1993), *Warrant and proper function* (Oxford University Press).
- Pritchard, D. (2005), *Epistemic luck* (Clarendon Press).
- Sainsbury, M. (1997), 'Easy possibilities', in *Philosophy and Phenomenological Research* 57/4: 907–919.

- Soames, S. (2010), *What is meaning?* (Princeton University Press).
- Sosa, E. (1999), 'How to defeat opposition to Moore', in *Philosophical Perspectives*, 13: 141–153.
- Stokke, A. (2013), 'Lying and asserting', in *Journal of Philosophy*, 110/1: 33–60.
- Stokke, A. (2014), 'Insincerity', in *Noûs* 48/3: 496–520.
- Stokke, A., & Fallis, D. (2016), 'Bullshitting, lying, and indifference toward truth', MS.
- Welbourne, D. (1994), 'Knowledge, testimony, and belief', in B. Matilal & A. Chakrabarti (eds.), *Knowing from words* (Kluwer), 297–313.
- Williamson, T. (2000), *Knowledge and its limits* (Oxford University Press).
- Williamson, T. (2007), *The philosophy of philosophy* (Blackwell).