forthcoming in *Synthese*

Testimony and Interpretation

Matthew A. Benton & Billy Dunaway

Abstract: Testimony can be a source of knowledge. This paper examines how misinterpretation, or the risk of it, can prevent a hearer from acquiring testimonial knowledge. Because unreliability in interpretation can arise in many ways, section 2 considers a variety of such cases. Section 3 sketches some desiderata for a successful account of the role of interpretation in testimony, by analogous consideration of inference. On our account, interpretation needn't proceed inferentially through knowledgeable belief about what is said. Finally, section 4 offers a safety-theoretic account of reliable interpretation which explains how and when misinterpretation prevents knowledge.

1 Preliminaries

Testimony can be a source of knowledge. The literature on the epistemology of testimony contains competing answers to the question of what, exactly, are the conditions under which testimony produces knowledge. Typically, it focuses on questions such as whether the receiver of testimony needs to know that the testifier is reliable, or whether the deliverer of the testimony needs to know or even believe the proposition testified. Yet some note that testimony involves an act of *communication*, and, in some circumstances, testimony can fail to produce knowledge when the communicative exchange goes awry.[1]

[1]Fricker 2003, Goldberg 2001, 2004, and 2015, Ch. 4, Longworth 2008 and 2018, and Peet 2016, 2018, and 2021 all raise this issue in some form. We will draw on these discussions as we sketch our own view below. We note in passing the related literature on referential communication (e.g. Loar 1976, Heck 1995, Buchanan 2013, Peet 2017, among many others); however, much of that literature is less focused squarely on the epistemological details, as we are here. Hence the literature on referential communication, or "linguistic understanding" from the philosophy of language, will be less relevant to what we have to say below.

Most of the time (though there are notable exceptions, discussed below) when communication in the context of testimony is raised, its relationship to whether the hearer can acquire knowledge is not discussed in detail. For example, Burge argues for an a priori, prima facie entitlement "to accept something that is prima facie intelligible and presented [in conversation] as true," acknowledging the novelty of his "claim that we are a priori entitled to rely on our *understanding* and acceptance of something that is prima facie intelligible."[2] Fricker (2006, 229) notes that "mutual understanding" where "a message must be got across and accepted," is a "precondition" for the spread of knowledge by testimony. And Lackey 2008 defends a view on which one can gain knowledge from someone's testimonial statement partly "on the basis of understanding and accepting the statement" (2008, 72).[3] In each case, the ways in which understanding can fail in testimonial contexts does not take center stage.

In this paper we examine the importance of interpretation to testimonial knowledge. Hearers can be unreliable interpreters for a number of reasons. The testifier and hearer might not be speaking the same language. Or even when sharing a language, a testifier might not use that shared language in order to produce a piece of testimony on a particular occasion. Finally, even when speakers are using their shared language on a particular occasion where testimony is produced, there can be multiple interpretations of what the testimony means in their shared language—testifiers can speak metaphorically, non-literally, or use expressions that have incomplete contents. In Section 2 we argue that there are commonplace cases where it is clear that a hearer fails to know due to unreliability in interpretation for these commonplace reasons.[4]

In Section 3 we sketch desiderata for a successful account of the role of interpretation in testimony by highlighting analogous features of *inference*. A

[2]Burge 1993, 472, incl. fn. 12; cf. Burge 2013, Chap. 11 for postscripts. Burge's focus is on the entitlement to rely on such understanding, though he later distinguishes comprehension from interpretation (cf. 2013, Chap. 14). See Bezuidenhout 1998, Malmgren 2006, and Longworth 2008 for important replies to Burge.

[3]Cf. also Graham 2000b, Audi 2006 42–43, and Sosa 2006/2011, Chap. 7, for similar claims.

[4]Other possible disruptions may arise due to difficulties over interpreting insinuations or conversational implicatures: see especially Fricker 2012 and Hawthorne 2012. We leave such cases aside.

good theory should not only explain why hearers who are unreliable interpreters typically cannot, for that reason, gain knowledge through testimony.[5] It should also explain this without requiring that speakers gain knowledge by first arriving at knowledge of what a testifier has said, and then *inferring* to a knowledgeable belief with the relevant content on the basis of that testimony. Yet we think that accommodating these desiderata only requires appeal to resources to which epistemologists are already committed.

Our account of this phenomenon highlights several respects in which testimony is analogous to other knowledge-generating and knowledge-transmitting processes. In particular, resources needed for a full account of knowledge by both perception and inference are also useful to provide a full account of testimony that accounts for the risk of misinterpretation. In Section 4, we illustrate one way the phenomenon can be accounted for within a simple, broadly externalist approach to knowledge-transmission through testimony, again by highlighting inference. Although we will not try to expand the account to be fully ecumenical, the resources we employ here can be appropriated by other approaches to the epistemology of testimony.

Central to this discussion is the connection between two components of forming beliefs on the basis of testimony. The first is that reliable interpretation of a testifier requires, at the very least, being in a position to know what it is to which the speaker has testified. Hearers who are not in a position to have this knowledge are typically thereby unable to have knowledge on the basis of the speaker's testimony. The second idea is that in successful cases of testimony, hearers need not *use*, inferentially, their knowledge of what the speaker has said to come to know. Rather, reliable interpretation operates in the "background" of the belief-forming processes deployed by a knowledgeable hearer.[6] Again, we aim to articulate one theory

[5] We allow that some cases of unreliable interpretation do not prevent knowledge of what was testified. But our focus is on cases where unreliable interpretation does prevent knowledge.

[6] One might fill out the details by appealing to quasi-perceptual understanding experiences as the basis for gaining knowledge of what a testifier says (Fricker 2003); or to a kind of "interpretative knowledge" (Sosa 2011, 135–136); or by appealing to the joint agency involved in interpretation presupposed by knowledge transmission (see Greco 2020, esp. Chap. 3); or to dispositions to correctly interpret another formed through a rich background of interpersonal relationship (Benton 2025). Apart from modeling the safety of such processes in Section 4, we shall remain neutral on these details.

that implements these ideas, but make no claims to exclusivity.

2 Examples

When does unreliable interpretation prevent knowledge acquisition by testimony? Peet 2019 adapts a case from Loar 1976 where the hearer is at risk of misinterpreting a speaker's use of a directly referential term:

> Suppose that Smith and Jones are unaware that the man being interviewed on the television is someone they see on the train every morning, and about whom in that latter role, they have just been talking. Smith utters 'He is a stockbroker' intending to refer to the man on the television; Jones takes Smith to be referring to the man on the train. Now Jones, as it happens, has correctly identified Smith's referent, since the man on the television is the man on the train; but he has failed to understand Smith's utterance. (Loar 1976, 357)

As Peet notes, it is plausible that in Loar's case that "Jones does not gain testimonial knowledge" (2019, 3309), even though the case was not introduced with testimonial knowledge in mind. Cases of this kind can easily be multiplied; any time context-sensitive expressions are used, there is always a possibility that a hearer misinterprets by fixing referents of singular terms in unreliable ways.

Such misinterpretation will usually rest on an avoidable mistake by the hearer. This is because, on plausible views of the reference-fixing mechanisms for singular terms, the hearer always has available a method guaranteed to fix the referent of the pronoun to the individual intended by the speaker by using a rigidified definite description referring to the individual the speaker intends to refer to. Thus Jones could have intended 'he' to refer to *dthat*: the person Smith is talking about.[7] Yet even if that sort of misinterpretation can be easily avoided, a speaker and hearer might entertain slightly different propositions owing to differences in their conceptual resources. The extent of

[7] Cf. Kaplan 1989 for the *dthat* operator and Manley and Hawthorne 2012 for a view of singular referring expressions along these lines.

these communicative mismatches depends in large part on what determines the content of a belief.[8] Nevertheless, if such mismatches occur, they threaten the ability of the hearer to gain testimonial knowledge, although arguably such mismatches do not guarantee that the hearer cannot acquire testimonial knowledge.

The cases in which misinterpretation interferes with testimonial knowledge to are not limited to those where context partly determines what the speaker said, or where speaker and hearer entertain slightly different contents. In this section we aim to expand the range of cases where interpretive problems prevent knowledge by testimony for ordinary and familiar sources of misinterpretation. A crucial feature of them is that interpretive problems should be the only barrier to the hearer gaining testimonial knowledge: so if the interpretive problems were not present, the hearer would have succeeded in acquiring knowledge by testimony. At a minimum, what the testifier intends to convey must be true. And ordinarily,[9] we don't gain knowledge from speakers whose true testimony is one that they fall short of knowing. So we will focus on cases where, even though there can still arise a mismatch in interpretation, the speaker also *knows* the proposition testified.[10]

[8] Goldberg 2007 favors externalism on the grounds that it provides for a wide range of shared content; internalists about content will require more nuance. See Pollock 2021 and 2024 for discussion. Peet 2019 raises cases where mismatch in content arises for other reasons. See also interpretation matters raised by Deigan 2022 on *stupefying*, and D'Ambrosio 2025 on *pied piping*; we do not consider such examples because they typically involves conversational manipulation, and at any rate such cases are less relevant to acquiring knowledge on the basis of testimony.

[9] Lackey (1999; 2008, esp. Chap. 2) argues that in many cases a hearer can acquire knowledge even when the speaker does not know, and thus that testimony can be a generative source of knowledge (see also Graham 2000a; cf. Graham and Bachman 2020 for an overview). But for our purposes it suffices to focus on cases where the speaker does know.

[10] In our examples we assume the speaker knows the truth of their testimony and knows what it is that they say when testifying. (It might be possible for a speaker accidentally to convey a different content than what they know, resulting in content mismatch between their knowledge and their testimony. But we do not consider such cases here, leaving those to future exploration.) Thanks to a referee here.

2.1 False friends

Begin with a simple case where a hearer fails to interpret a testifier correctly because they are not speaking the same language, but the hearer does not know this. The word 'angst' has slightly different meanings in English and German. In German, 'Angst' expresses a particular worry or fear: one can have *Angst* that is occasioned by a friend's poor health. In English, 'angst' has more psychological overtones, conveying a general feeling of anxiety or apprehension. An angst-y response to a specific thing—a friend's health, or heights—doesn't make sense in English. 'Angst' in German and English are "false friends."[11]

Imagine, then, a German speaker is in an international city at a bus station and has a friend with them, who is visibly unhappy. Having seen a concerned stranger observing, they point to the friend and say 'Angst', which is understood by the stranger to communicate the cause of the friend's unhappiness. Let us suppose the speaker knows that the friend is afraid of some specific thing—snakes, say—and knows that this fear is the source of the friend's unhappiness. (The speaker does not use any additional words, or speak with an accent, that would let the hearer know whether the speaker is using German or English.) The friend is *also* experiencing psychological feelings of anxiety, but these are unrelated to the fear of snakes, and to the speaker's attempt to explain the friend's unhappiness. As a matter of fact the testifier is speaking German, and uses the German 'Angst', intending to communicate (knowledgeably) that the source of the friend's unhappiness is fear of snakes.

A hearer who interprets the speaker as using the false friend of 'Angst', to communicate that the friend has a general sort of psychological anxiety, can believe something true on the basis of the testimony. By misinterpreting the testifier, even if the hearer has a true belief, intuitively he will not have knowledge. If the friend had not had general feelings of anxiety, the testimony and process of communication would have been exactly the same, and the hearer would have ended up with a false belief.

[11] This sort of case feels similar to cases from Pettit 2002, though his examples do not involve misinterpretation owing to 'false friend' terms.

2.2 Semantic underdetermination

Even when speaking the same language, failures of interpretation are possible. Suppose a testifier says 'there was a robbery on the corner of 8th and Main St.'. This testimony is semantically underdetermined in the following way: the definite article 'the', on orthodox views, semantically requires (or presupposes) that there is a *unique* thing which satisfies the description following the article. It is common knowledge that there is no unique corner of 8th and Main St., since in normal cases the intersection of two streets has four corners.[12]

Suppose that the claim which the speaker intends to communicate is that there was a robbery on the northeast corner of 8th and Main St. The hearer, however, does not grasp that intention, and interprets the testifier to have communicated that there was a robbery on the southwest corner of 8th and Main St. The cause of this error does not need to be a wild guess on the hearer's part; it could be the result of wrongly believing that 'the corner' in this sentence anaphorically refers to an earlier mention of the southwest corner of 8th and Main St. in the conversation. In this case, the hearer's belief that there was a robbery on the southwest corner of 8th and Main St. is not knowledge, even if there was in fact a robbery on the southwest corner in addition to the northeast corner, and so the belief is true.[13]

2.3 Non-literal speech

A final set of cases involves a testifier who testifies using speech that is to be interpreted non-literally, including uses of idiomatic meaning, and of metaphor. The sentence 'Sally has a bee in her bonnet' on its literal meaning entails that Sally is wearing something on her head, which has an insect in it. The expression 'has a bee in her bonnet' also has an idiomatic meaning,

[12] There are a number of competing views of the semantics of the definite article, and how (and whether) the missing content to achieve uniqueness gets filled in. See Graff Fara 2001; Soames 2005; and Hawthorne and Manley 2012, esp. Chap. 5.

[13] See Recanati 1993, 48–59, for interpretation of other underdetermination cases owing to *de re* communication using demonstratives or indexicals.

which implies that Sally is upset.[14]

Other non-literal meanings are possible, even for expressions that don't have idiomatic meanings in English. In a context where it is known that John's children are wild and unruly, but unknown to the hearer how many children John has, the testifier might say 'John has three Tasmanian devils'. One way to characterize what is communicated is to say that the literal content of the testifier's assertion is that John has three large carnivorous marsupials; yet the non-literal content of the assertion, which is what the testifier knows and intends to communicate, is that John has three rowdy children.

Misinterpretations of such non-literal speech are possible. If the hearer believes the testifier has intended to pick up on a previous conversational topic of John's family life, the hearer will interpret the testimony non-literally and believe that John has three rowdy children. The literal interpretation could in fact be true, and known by the testifier. But the hearer still wouldn't thereby know that John owns carnivorous marsupials, if they have misinterpreted the speaker. That the misinterpretation, which results from taking the testifier literally, is a proposition that is true and known by the testifier is a mere accident.

[14]Cases of *hermeneutical injustice* (Fricker 2007) can create situations of misinterpretation perhaps akin to idiomatic cases. In such cases a subject is prevented from generating meanings pertaining to some of their social experiences, owing to structural identity prejudice in the collective understanding, but they are nevertheless trying to communicate the contours of such experiences (as when women did not yet have the conceptual resources to name cases of sexual harassment: Fricker 2007: 149ff.). Where such hermeneutical resources are developing and still disseminating, it will be possible for some speakers to conceptualize such experiences and form the relevant beliefs, while many others as yet cannot; and when the latter are hearers, and the former testifiers, it will be plausible that misinterpretation is part of what blocks a hearer from believing (and thus knowing) the relevant content. (As Fricker notes, often hermeneutical injustice can be compounded by testimonial injustice (2007, 159–160), such as when women, who speak under conditions of identity-prejudicial credibility deficit, try to communicate their experiences of such harassment to men.) On the other hand, reversed cases where the speaker but not the hearer currently lacks the hermeneutical resources might allow for a speaker to intend a content which she herself cannot yet fully grasp, but which that hearer is able to understand. Such cases where a hearer's hermeneutical resources outpace those of the speaker may count as cases which generate knowledge in the hearer without knowledge in the speaker. Our safety-theoretic account in Section 4.3 substantiates one way this could be possible, although our focus there is on cases where the hearer misinterprets rather than correctly interprets the testifier.

2.4 A second route to the absence of knowledge through unreliable interpretation

The above examples show that knowledge can be absent purely because of an unreliable interpretation by the hearer. But they do not represent the only way in which misinterpretation can prevent knowledge from testimony. In fact, they might not be the best illustration of it, for the following reason: the hearer does not believe the proposition the speaker intends to communicate, even if the speaker does also believe the proposition the hearer comes to believe. On some definitions of testimony, such instances of communication do not count as genuine cases of testimony. On these definitions, a hearer only believes a proposition *by testimony* when the hearer believes the *same* proposition that the speaker asserts.[15]

A hearer might receive testimony from a knowledgeable testifier, believe the content of the intended testimony, and so come to have a true belief that is identical to what the speaker intended to testify. But even in this case the hearer does not necessarily have knowledge, if the hearer has a true belief about the content of what the speaker testified, but is an unreliable interpreter of the testimony.

False friends, semantically underdetermined expressions, and non-literal speech can all illustrate this point. For example, suppose in the example from Section 2.3, the testifier uses the sentence 'John has three Tasmanian devils' to communicate the literal meaning of the sentence. The hearer, moreover, receives the testimony by believing the literal meaning. Assuming the testifier has knowledge, the hearer has a true belief in a proposition identical to the one asserted. But this is not enough for knowledge if the hearer arrives at the interpretation in an unreliable way. As Goldberg (2007, 44) notes, an unreliable recipient of testimony might fail to have knowledge because the hearer would believe the same proposition even if the speaker had asserted something different. Suppose, for instance, that the hearer has a policy of always believing the literal interpretation of the sentence used to

[15]Others (cf. Peet 2019) claim that hearers sometimes arrive at knowledge even when believing a proposition that is distinct from the proposition asserted by the speaker. We do not take a stand on this issue here. In this subsection we present cases where unreliable interpretation prevents testimonial knowledge even though the hearer has believed the same proposition that the speaker has asserted.

communicate any piece of testimony, no matter what contextual clues are in place to suggest that a non-literal interpretation is intended. In this case, it is intuitive that even if the testifier did intend the literal interpretation of 'John has three Tasmanian devils', but could easily have intended the non-literal interpretation, then the hearer does not know the literal content of the testimony.[16]

Here the barrier to knowledge is not an actual false belief about what the speaker said, but rather the possibility of a false belief based on testimony. Similar possibilities exist when testifiers could have been speaking a different language, or when semantically incomplete testimony could have been interpreted incorrectly.[17] It is clear that the mere possibility of interpreting a testifier's audience incorrectly should not prevent knowledge; testimonial skepticism would result from such a principle. However, when the possible sources of misinterpretation are realistic sources of error, they can prevent knowledge.[18]

Characterizing this phenomenon in more detail is one task for the epistemology of testimony, and one which we will engage at length in the next two Sections. An important observation that is central to our discussion below is this: while the above cases highlight various ways that misinterpretation can prevent testimonial knowledge, it does not follow that successful acquisition of testimonial knowledge requires that the recipient of the testimony knows the facts about the correct interpretation of the testifier before acquiring

[16] As subsequent sections make clear, this is not merely a feature of the hearer's policy: perhaps in a community that never uses non-literal speech, the flat-footed policy of believing the literal content of a testifier's assertion can produce knowledge.

[17] The cases presented here differ from those in Goldberg 2004. Goldberg is interested in whether hearers could come to know what a testifier has said through a process of 'radical interpretation', which requires scrutinizing the meaning of a string of words on the basis of the ground-level non-semantic facts. It is clear, as Goldberg points out, that hearers are not in general in a position to do this. The argument here is different, since we are granting that speakers are in a position to know what sentences in their language mean, even if they are not in a position to work out what a sentence means from the ultimate grounds of the meaning-facts alone. Even with these assumptions in place, problems of interpretation can arise and prevent knowledge.

[18] On the view we will develop in Section 4, "realistic" possible errors will get cashed out as errors that make the belief unsafe, namely, where there is a close-enough world in which the believer, using the same (interpretive) method, goes wrong.

knowledge. That is, in some cases, a recipient of testimony can acquire knowledge from a testifier without having higher-order knowledge of what the testifier said. This is simply an intuitive observation about many successful cases of testimony. For example: if John delivers a piece of testimony with the speech "I have a car" and Suzy simply forms the belief that John owns a car on the basis of John's testimony without forming any higher-order belief about the content of John's speech, Suzy can have testimonial knowledge that John owns a car. That is, it appears that this knowledge can be had without Suzy forming any higher-order beliefs to the effect that John's utterance of "I have a car" on this occasion means that John owns a car. Likewise Suzy needn't form higher-order beliefs that rule out other interpretations, such as that John's speech means he is in possession of a car he can use (but does not necessarily own, e.g. because he stole the car). This is just an intuitive observation about some cases of testimonial knowledge-transmission. We explore this idea more systematically below, starting in Section 3 with analogies with other means of knowledge-production, and then in Section 4 by drawing out some implications from an account of knowledge in terms of a safety condition.

3 Testimony, memory, and inference

Testimony is often described as a means of knowledge transmission, to be contrasted with a source of knowledge generation.[19] In this respect, testimony can be understood along the lines of memory, which transmits knowledge gained by a person at one time to that same person at a later time. Testimony, the thought goes, serves a similar function when it transmits knowledge from one person to another (possibly, but not necessarily, at the same time). When transmissive, testimony passes on knowledge already had by another, and thus functions unlike perception, which typically generates new knowledge without being dependent on another knower.

[19] Cf. Lackey 2008 for arguments that testimony can be generative of knowledge rather than merely transmissive, since she argues knowledge can be acquired from testifiers who do not even believe what they testify. See Greco 2020 for testimony as transmission of knowledge (though his view of transmission is somewhat different, he nevertheless focuses on "Transmission… [which] concerns coming to know 'from someone else,' as when one is told from someone else who knows" (p. 1), and argues that knowledge transmission is irreducible to knowledge generation: cf. Greco 2020, 2–13, and 34–36).

3.1 Analogies with inference: basing

The role of successful interpretation in producing testimonial knowledge suggests another analogy, this time with a knowledge-generating process, namely *deductive inference.*[20] Inference is a knowledge-producing process, roughly, when one knows the premises of an argument and validly deduces a conclusion from the premises. There are complications, however:[21] for example, one must retain knowledge of the premises throughout the deduction. If one loses knowledge (or, for paradoxical reasons, cannot retain knowledge of the premises and simultaneously know the conclusion), then the derivation from known premises does not produce new knowledge of the conclusion.[22]

More importantly for present purposes is that new knowledge from inference has to be *based* on knowledge of the premises in the deduction. It is not enough simply to know the premises, and to believe the conclusion that is entailed by those premises. One's belief in the conclusion must be based on those premises in an appropriate way, though it may be difficult to adequately define this basing relation amounts to.[23] The intuitive idea is, however, clear: one knows that Socrates is mortal when one first knows that all men are mortal, that Socrates is a man, and believes, through syllogistic reasoning involving these premises, that Socrates is mortal.

Knowledge through testimony requires something similar. Failure to knowledgeably interpret the testifier is similar to a failure to base inference on known premises in the right way. When a testifier knows p and a hearer believes p in response to their testimony, it is natural to conclude that the hearer knows p. But this is not strictly entailed by the circumstances of the hearer's belief-formation, because it is possible for the testimony not to have p as its content. If the hearer mistakenly interprets the testifier as having

[20] Cf. Bezuidenhout (1998, esp. 269ff.), who argues against Burge that the interpretive processes underlying verbal communication are generative not simply preservative.

[21] Hawthorne 2004, esp. 31–50.

[22] Is there an analogous case for transmissive testimony? Perhaps: suppose I know I have never at any time made a knowledge-producing speech. If I say to you "I have never at any time made a knowledge-producing speech" then either you do not come to know what I have said, or I lose the knowledge I had before the speech. (Thanks to a referee for raising this question.)

[23] See Carter and Bondy 2019 for discussion.

asserted p, even if the testifier does know p, the hearer's acquired belief in p isn't based in the testifier's knowledge in the right way.

All of the instances of misinterpretation from Section 2 fit this pattern. In each case, the speaker knows some content: that the angsty friend is generally unhappy, that there was a robbery on the southwest corner of 8th and Main St., or that John owns three Tasmanian devils (meaning *marsupials*). What the speaker actually testifies, in each case, is something else: that the friend has a specific fear, that there was a robbery on the northeast corner of 8th and Main St., or that John has three rowdy children. The hearer's true beliefs are not based on what the testifier actually testifies; rather, these beliefs were not the contents conveyed in testimony. The source of this basing failure is misinterpretation, through either wrongly interpreting the testifier as speaking a different language, or in how to fill in the underdetermined content of the speaker's testimony, or in understanding whether the testifier is speaking literally.

3.2 Analogies with inference: processes vs. knowledge

The analogy with inference is helpful in a further respect. Inferential knowledge does not depend on knowing that the premises entail the conclusion, or on knowing that the inference is valid. If knowledge of the validity of an inference were required, then a regress threatens.[24] For if, in order to validly deduce q from p, it is not sufficient to know p, but also required that one know that p entails q, then one must know that p, and that p entails q, and know that: [(p & p *entails* q) entails q], and so on. Instead, one's belief-forming process can involve *inferring* the conclusion from the premises, which needn't involve a belief about an entailment. Thus one who knows *John is in the room* and *if John is in the room, then Bob is in the room* comes to know *Bob*

[24] See Carroll 1895, and Harman 1986.

is in the room by competently[25] deducing it via modus ponens. One who, through inattention or indifference, fails to form a higher-order belief about the validity of modus ponens does not thereby fail to acquire new knowledge. What matters is that the conclusion is believed via a *process* that involves deductively inferring it from the premises.

Similar points apply to testimony. In a good case, a hearer of testimony properly interprets it from a knowledgeable speaker who testifies that p; the hearer correctly interprets them as having said p, and comes to know p. What is the role of the correct interpretation in producing the knowledge? Knowledge that the speaker has said that p needn't be a premise in the hearer's reasoning producing an inference to p. This would be analogous to a requirement that someone who learns a conclusion through a valid modus ponens inference must use the knowledge that this instance of modus ponens is valid as an additional premise. In each case, it is *possible* that one arrives at new knowledge in this way, but we maintain that this is not *necessary*.[26]

One complication here is that competing views of successful testimony divide into what are commonly called *reductionist* and *non-reductionist*.[27] In this section we remain neutral between them (we'll drop the neutrality in the next section), but some comments will substantiate this neutrality. While proponents of either view can endorse the claim that knowledge of the correct interpretation needn't feature as a premise in an inference to gain

[25] If pushed on what this involves, we may appeal to Sosa's notion: a competence is a disposition to succeed in a given domain, with a performance undertaken in a particular situation, aimed at success, which manifests a skill (Sosa 2015, 95ff.; and 2021, 44–45, 145–147, 199–201). In the case of reasoning, it is a disposition to form a new belief on the basis of accepting some premises, with the aim of securing the truth. (Sosa's own discussion of competences often invokes safety-theoretic notions, e.g. 2015, 97 n. 5.) (Thanks to a referee here.)

[26] We note the (partial) analogy with inference, though we don't intend it carry to all of the weight. There is also the intuitive force of our earlier examples where it is stipulated that the hearer does not form a belief about what the testifier's speech means. If everything else goes well with the testimony I can, in some cases, still acquire testimonial knowledge on the basis of your speech. In addition, it falls out of the safety account (invoked later in Section 4) that with the right dispositions one can acquire testimonial knowledge of p without antecedently believing that the testifier's speech meant p. Insofar as safety is a reasonable account of knowledge, this is a further argument for our claim that antecedent knowledge of what is said isn't necessary. (Thanks to two referees here.)

[27] See Coady 1992; Fricker 1995; and Lackey 2006, among others.

knowledge of the testified proposition, the role of the interpretation will look very different, depending on each view.

According to reductionist views, to know from testimony one must possess non-testimonially based positive reasons for believing it, reasons grounded in basic sources such as perception, memory, and inductive inference. On the basis of these, including facts about the nature of the testimony and its circumstances, one must be in a position to support inferentially the testified claim. By contrast a non-reductionist, who thinks of testimony as a basic source of justification, denies this: one can know from testimony without being able to infer it from what one knows about it by non-testimonial sources. For a non-reductionist, it will be very natural to accept our thesis: if knowledge from testimony doesn't require inference from any non-testimonial premises, then it shouldn't require inference from a set of premises that include the (known) claim that the content of the testimony is such-and-such. As with other components of the non-reductionist account, it will need to provide details on the conditions under which testimony produces knowledge.

For a reductionist, the conditions under which testimony produces knowledge include the hearer having knowledge of some facts from non-testimonial sources. But even a reductionist will not require that one must antecedently know *all* of the non-testimonial facts about the occasion of testimony. So the reductionist is not committed to holding that the non-testimonial facts from which a hearer infers the testified proposition must include the fact that the speaker testified that such-and-such. Suppose one is a reductionist who holds that, in order to learn from testimony, one must know that the testifier is sufficiently reliable, has no reason to lie, and has expertise in the area of testimony. There is no obvious reason for a reductionist of this kind to hold that in addition to knowing these facts, one must also know that the speaker

testified that such-and-such.[28]

This points to the following parallel between inference and interpreting testimony. Knowledge gained through these sources must be based on competent inference from premises, or on reliable interpretation, respectively. But in each case, one needn't know the proposition that corresponds to the reliable inference or interpretation, and deploy *it* as a premise in an inference to the proposition one comes to know. In sum: there are two points of analogy between interpretation and inference in producing knowledge. (i) Being an unreliable interpreter or inferrer can prevent one from having knowledge on the basis of testimony or inference. (ii) In order to acquire knowledge on the basis of testimony or inference, one does not need to know and reason from facts about what the speaker meant, or facts about what the premises imply. In the next section we sketch one model of the role of interpretation in arriving at knowledge by testimony, which preserves both of these points.

4 A model: interpretation and safety

Here we drop the earlier neutrality and present one *externalist* and *non-reductionist* model of reliable interpretation. The model illustrates how the the role of reliable interpretation of testimonial knowledge can be accounted for with no additional resources. Some competing theories will be able to accommodate observations that are similar to those we make here.[29] But there is no guarantee that all will, and so the role of interpretation in testimony may be a deciding issue between competing theories.

[28] Note that a reductionist who required testimonial knowledge of p to involve a deduction of p from a set of known premises which includes "the testifier's speech meant that p" would seem committed to the following: One can fail to know that the testifier's speech meant that p either because (i) one fails to form any belief at all about whether the testifier's speech meant p, or (ii) one is an unreliable interpreter resulting in a false belief about what their speech meant, or (iii) one unreliably forms a true belief about what their speech meant. All three possibilities are not threats to having knowledge of p based on testimony on the safety-based non-reductionist view we will embrace in Section 4. So if the reductionist does require knowledge of what the testifier's speech means, the reductionist introduces a number of additional obstacles to acquiring testimonial knowledge. Maybe this reductionist view isn't all that unnatural. But it is a view that has undesirable consequences. (Thanks to a referee here.)

[29] Perhaps Longworth's (2018) 'entertaining view' can do so, discussed in Section 4.4 below.

4.1 Safety sketched

Knowledge requires a true belief that is believed in an appropriately reliable way;[30] we think of the appropriate notion of reliability in terms of *safety*.[31] A belief b held by a subject S is safe just in case in all "nearby" worlds w, or worlds that could easily have obtained, if S has the belief b in w, then S's belief b is true in w. This needs three refinements: first, in order for the belief b to be safe, it must also be that beliefs that are *similar* to b are true, if held. Even one who correctly guesses that $49 + 118 = 167$ does not know the relevant mathematical fact, since they easily could have believed a similar but false mathematical claim. Second, only those false beliefs in nearby worlds formed by *similar processes* to those that formed b (by S in w) can prevent b from being known. One actually believes that Caesar crossed the Rubicon on the basis of what one has read in history books. Suppose however that one could easily have met a convincing conspiracy theorist, who would have made a compelling case that the accounts of the late Roman Republic in history books have been fabricated. The nearby possibility of false beliefs in this situation is compatible with knowing that Caesar crossed the Rubicon. Believing on the basis of what one reads in history books is a very different process of belief-formation than believing conspiracy theorists.[32]

Finally, what counts as a "nearby" world? At a first pass, these are worlds that could easily have obtained: if one resolves to believe p if a fair coin lands heads, and to believe $\neg p$ if the coin lands tails, then if the coin lands heads,

[30] See Goldman 1986, Goldberg 2010, among others.

[31] See Sosa 1999, Williamson 2000, Pritchard 2005, among others. See Rabinowitz 2011 for an overview. Peet 2019 is one of the few to consider such modal conditions on apt interpretation and knowledge from testimony, and his proposal is somewhat similar to our own. But his account ("Hybrid Luck," 2019, 3322) centers on anti-luck conditions which he does not flesh out in detail (as he implicitly acknowledges: p. 3325).

[32] Reliabilists such as Goldman will hold that the method of reading history books is a distinct method from that of relying on testimony from a speaker, and so one believes that Caesar crossed the Rubicon by using a method that is reliable. Another way of implementing the idea is to treat different processes that produce beliefs as more-or-less similar. Two belief-formation episodes might both rely on testimony, which is prima facie, but not conclusive, reason to hold that they are the products of similar processes. What matters is not whether the processes fall under a general description (such as 'testimony' or 'perception') but whether the individual instances that produce the beliefs are sufficiently similar.

there is still a nearby world where one believes $\neg p$.[33] But there need not be a plausible analysis, in non-epistemic terms, of what constitutes a nearby world. For we may need to rely on intuitions about whether one knows in a particular case, in order to settle whether certain possibilities where one believes falsely count as nearby worlds. Thus a safety condition on knowledge cannot serve as a proper analysis of knowledge in the traditional sense; but this does not prevent the condition from illuminating central features of knowledge.[34]

This framework distinguishes between what has been called[35] *locally reliable* and *globally reliable* belief-forming processes, or methods. Roughly, a process is locally reliable if it will in fact, in the actual environment, produce a safe belief, if it is given safe inputs. For example, an abductive inference that the future will resemble the past is, in environments like ours, true in all nearby worlds if it is based on beliefs about past regularities that are true in all nearby worlds. The process of believing by an inference like this is locally reliable. Such an inference is not guaranteed to produce knowledge: there are distant possible-but-not-actual environments where similar inference would yield a false belief because it is possible that the laws of physics are radically different. A process is globally reliable only if, in every possible context in which it is deployed, it yields a safe belief from safe inputs. Deductive inference is the paradigm case of a globally reliable belief-forming process. In what follows we assume that very little knowledge rests on globally reliable processes, and in general we acquire knowledge by utilizing locally reliable processes.

[33] The belief in this nearby world will be very similar—the belief in p is similar to the belief in $\neg p$—and will be formed by a similar method—relying on the coin landing heads is similar to relying on the coin landing tails.

[34] Williamson 2000, Chap. 5. Analogous points apply to settling whether a belief in a nearby world, and the process by which it is formed, count as sufficiently similar. This conception of safety is circular, in the sense that it uses the concept of knowledge in the "analysis" of knowledge as safe belief (since nearby worlds are, in effect, worlds one cannot know one is not in). But the analysis is not circular in the sense that it is *trivial.* The safety account makes plausible, and non-trivial, claims about the connection between knowledge and possible false beliefs, as well as belief-forming processes. In section 4.2 below we shall consider different ways of explicating what counts as a "similar process" in the safety condition.

[35] Goldman 1986, 44–46, and especially Goldberg 2010, 12, and 50–54. Cf. also Plantinga 1993 and Bergmann 2006 on proper functionalist views.

4.2 Testimonial safety

Return to knowledge from testimony, setting aside for the moment the role of interpretation in gaining such knowledge. In a simple case where a testifier knows p, their hearer can know p on the basis of it for the following reason. Since the testifier knows p, the testifier's belief in p is safe—roughly the testifier could not easily have falsely believed p. When the hearer believes p on the basis of such testimony, their belief will, in a normal case, be free from the risk of error in the same way. On this picture, a hearer needn't know that the speaker knows p, or base her belief in the testimony on her beliefs about any specific facts involving the testifier and testimonial environment, in order to acquire knowledge of p by testimony. Often the hearer's belief, based on the testimony of a knowledgeable testifier, can be safe even if the hearer is not aware of the conditions that make it safe.

However, this (externalist, non-reductionist) outline of how knowledge can be acquired by testimony leaves several details to be filled in, corresponding to the refinements of safety that we sketched above. The most obvious needed detail concerns the process by which one might come to acquire similar beliefs in nearby worlds. Suppose one believes a knowledgeable testifier, who testifies that p. But one is not in a position to distinguish the knowledgeable testifier from those who don't know and, moreover, it could easily have been that one heard, and so believed, $\neg p$ on the basis of the testimony of an ignorant testifier. In that case, one would have a false belief in a claim that is very similar to p. Is it formed by a process that is similar to the process by which one actually believes p?

There are multiple ways to answer this question within the simple framework sketched here. One answer is that believing the ignorant testifier is a very similar process of belief-formation to the process of believing a knowledgeable testifier.[36] In this case, whether one can know by testimony will depend on certain local features of one's external environment, including whether any ignorant would-be testifiers are nearby. Call this a ***knowledge-***

[36]It is important in giving this answer not to rely on the intuition that, since one cannot tell what distinguishes the knowledgeable from the ignorant testifier, the processes must be similar. The externalist view sketched here will in general reject a test for process-similarity that requires accessibility of the distinctness of processes.

ignoring view of the similarity-relation on processes one might use to form a belief by testimony.

The alternative to a knowledge-ignoring view is a *knowledge-sensitive* view of the relevant processes. On this view, even if there are ignorant testifiers one could easily have encountered, one can still know that p if one relies on the testimony of one who knows p. According to the knowledge-sensitive view, believing someone who knows that p is a very different belief-forming process from the process of believing someone who does not know whether p. Thus on the knowledge-sensitive view, successfully coming to knowledge by testimony depends on some features of the external environment, including whether one is believing a knowledgeable testifier. But it does not depend on whether there are other, non-knowledgeable testifiers around: even if one could easily have formed a false belief by believing one of them, that belief would have been formed by a substantially different process. Thus on a knowledge-sensitive approach, the nearby presence of an ignorant testifier does not threaten the status of one's actual true belief as knowledge.

We will not try to settle the question of whether a knowledge-ignoring or a knowledge-sensitive view of the belief-formation process for testimony is correct. Instead we simply note it as one issue that is undecided by the general framework. Similar issues arise when one focuses on the role of interpretation in producing (or failing to produce) testimonial knowledge.

4.3 Interpretive safety

Suppose one forms a true belief by testimony, but has in fact misinterpreted the testifier. This case is fairly straightforward given our framework. Since one has misinterpreted the speaker, one typically could have easily have formed the exact same belief, in the exact same way, but believed falsely, and thus the interpretation is unsafe. And so our model predicts that this will not result in testimonial knowledge. Analogously, when one believes something true by an invalid inference, one does not know for similar reasons. One could easily employ the same method (invalid inference) with the same premises, and come to believe something false.

More interesting cases involve one who believes what a knowledgeable testifier has said, but is still an unreliable interpreter in the sense that one

could easily have believed something the speaker did not say. Suppose that in some of the nearby worlds where one misinterprets the testifier, what one believes is also false. Does this show that the risk of misinterpretation prevents knowledge? That is, does it show that, even if one actually believes p on the basis of a knowledgeable testifier having testified p, being an unreliable interpreter would prevent one from knowing? A simple application of a safety-based model suggests that the answer is 'yes'. If there is a nearby world where one misinterprets the knowledgeable speaker, then there is a nearby world where one believes falsely. Hence one fails to know.[37]

Yet this simple application ignores several refinements over what it takes for a belief to be at risk of being false, in the relevant sense. Take someone who interprets an utterance of 'there was a robbery on the corner of 8th and Main St.' correctly, believing it to mean that there was a robbery on the *northeast* corner of 8th and Main St. However, let us suppose that the hearer could easily have misinterpreted the utterance—perhaps by failing to pick up on conversational cues that the testifier intends to speak about the northeast corner—and would then have falsely believed that there was a robbery on its *southwest* corner.

This does not necessarily force us to conclude that the hearer does not know in the actual case, but the details will mirror the choice-point we face in how to characterize belief-producing processes in cases of testimony more generally. One might hold that a process leading to a correct interpretation is sufficiently dissimilar from a process resulting in a misinterpretation. After all, the first process involves things like the hearer knowing that certain earlier conversational cues determine the correct interpretation; the second process does not. This way of characterizing the similarity-relations between the processes allows us to hold that the hearer knows, despite the risk of

[37] Compare Sosa (2011, 135–136): "We must interpret our interlocutors, so as to discern the thoughts or statements behind their linguistic displays. ... Interpretative knowledge of what a speaker thinks (says) is thus instrumental knowledge that uses the instrument of language. Language is a double-sided instrument serving both speaker and audience. Hearers rely on the systematic safety of the relevant deliverances. Not easily would the speaker's utterance deliver that the speaker thinks (says) that such and such without the speaker's indeed thinking (saying) that such and such. ... If any of this is put in serious enough doubt, the supposed instance of testimony will be disqualified as a source of knowledge about its direct content, for that audience at that time."

a misinterpretation. It resembles in this respect the knowledge-sensitive approach to the relation between processes in the ordinary testimony case, since it takes features of the testimonial environment which are directly tied to the epistemic status of participants to generate substantial differences between processes that might have been used. Those that have no problem with a knowledge-sensitive individuation of processes might be happy with this result.

However, others inclined to a knowledge-ignoring characterization of the (relevant) differences between processes might be tempted to treat the risk of misinterpretation more seriously. The resulting view holds that the hearer's reliance on her knowledge of what occurred earlier in the conversation does not necessarily mean that in scenarios where she fails to rely on such knowledge, she is using a substantially different process. On this view, beliefs formed by misinterpretation of this kind are not ipso facto irrelevant to whether the hearer in the actual case has knowledge by testimony.

Regardless of how we come down on such details, they matter only for *how much* the risk of misinterpretation prevents testimonial knowledge, and not *whether* misinterpretation is relevant to testimonial knowledge. Some hearers might be in a position to know the salient facts about their conversational context, making a particular interpretation the right one; however they might fail to rely on this knowledge, arriving at a particular belief as a result. In many cases where the communicative environment is friendly, this is not a barrier to knowledge. (English-speakers who hear 'Timbuktu is in Africa' and come to believe the truth that Timbuktu is in Africa wouldn't fail to know simply because they have not formed additional beliefs about their conversational environment before forming this belief.) But when misinterpretation is a real possibility, such quick belief-forming processes introduce risk that may be incompatible with knowledge. In these cases, even when a hearer correctly interprets a knowledgeable speaker, the risk of misinterpretation will not involve a process that can be distinguished by the absence of a knowledge-state somewhere in the process. Thus a simple safety-based model of knowledge by testimony will deliver the result that, at least in cases involving quick, unreflective interpretation, unreliable interpretation can be a barrier to knowledge. For those not sympathetic to knowledge-sensitive

approaches, it will be natural to hold that the epistemic consequences of unreliable interpretation are even more widespread.

4.4 The role of inference

Reliable interpretation is typically needed to secure knowledge by testimony, since unreliable interpretation typically introduces the risk of false belief. But this doesn't mean that one needs prior knowledge of the correct interpretation of some testimony, and must use it to infer the content of the testimony in order to eliminate the risks arising from misinterpretation. As our discussion of the safety condition indicates, one can have a safe belief on the basis of testimony simply by having a sufficiently robust disposition to form a belief that matches, or is sufficiently close to, the content of the testifier's speech. Nonetheless, it does seem that the hearer who forms the correct belief, even if it is the result of an unreflective linguistic disposition, is in an epistemic position that is different from someone who misinterprets (or is at risk of misinterpreting) the testifier. Here we aim to say more about this state of the reliable interpreter who is thereby in a position to acquire testimonial knowledge.

Longworth 2018 distinguishes between a concept of linguistic understanding on which understanding requires a higher-order belief *about* what the speaker has said and a more minimal alternative, the "entertaining account," on which understanding simply consists in entertaining the proposition the speaker asserted, in the right way. While it is implausible that acquiring testimonial knowledge requires having a higher-order propositional attitude about what the testifier said, the hearer needs to have some kind of cognitive access to the speaker's assertion, so that it can be a source of distinctively *testimonial* knowledge.[38]

The model developed above enables an account of the epistemic position of hearers who learn from testimony. It will share with Longworth's account the concession that hearers need only be in a very minimal epistemic relation to what is said by the testifier: they need not form explicit, knowledgeable beliefs about what the testifier has said. However, instead of giving an account

[38] See also Fricker's (2003, 341–342) non-inferentialism, though she posits quasi-perceptual understanding experiences as the basis for gaining knowledge of what a testifier says.

of the more minimal state of "understanding," we propose to characterize speakers who have not misinterpreted a testifier in terms of what, in the right conditions, they are in a position to know.[39]

If a testifier utters

(U) There was a robbery on the corner of 8th and Main St.

and thereby says (owing to intentions and features of the conversational context) that there was a robbery on its *northeast* corner, a hearer can, in some circumstances, know the following claims:

(1) The testifier said that there was a robbery on the northeast corner of 8th and Main St.;

(2) What the testifier said (in this circumstance) is true.

From (1) and (2) they can deduce that there was a robbery on the northeast corner, and so if these are known, the conclusion is known as well.

On our model, this kind of inference relying on (2), a higher-order premise about what the testifier said, is not *necessary* in every case to form a true belief on the basis of testimony. The epistemic relationship between the hearer and what the testifier said is not a component of what makes the hearer capable of gaining knowledge from testimony. Instead it is downstream from the acquisition of testimonial knowledge. There is no need for an antecedent

[39]Longworth 2018 provides an extended argument that knowledge-transmission of p through testimony does not require higher-order knowledge of the fact that the testifier said that p. Thus his view has the same upshot as ours insofar as it only requires hearers who acquire testimonial knowledge to be in a position, in some sense, to know what was said (2018, 817–818). However Longworth also includes in his account some details we don't commit to here: hearers who acquire testimonial knowledge are in a state of "understanding" what was said, and the state of understanding plays a role in testimony that is analogous to the role seeing an object plays in acquiring perceptual knowledge of that object. We don't commit to the necessity of some prior state, like understanding, that one must be in (though we agree with him that a state of higher-order knowledge of what is said is not required) for acquiring testimonial knowledge. Moreover we do not wish to rely on seeing/perception as an analogy since our aim here is to press analogies with other knowledge-transmitting processes. However nothing we say here conflicts with Longworth's account, and he is explicit (p. 822) that his tools enable hearers to satisfy a safety requirement when arriving at beliefs based on testimony, without having higher-order knowledge of what is said.

"understanding of what was said" to secure this knowledge. Access to the content of the testifier's speech takes the form of an available[40] piece of knowledge.[41] A hearer who has learned through testimony that there was a robbery on the northeast corner of 8th and Main St. will have this proposition available as a premise in knowledge-generating reasoning. If such a hearer is asked what the testifier meant by uttering (U), in some cases the reasoning to the conclusion that they meant there was a robbery on the *northeast* corner of 8th and Main St. will be locally reliable.

But this account of what the testifier meant can be based partly on the hearer's knowledge which she acquired by testimony. The known premises in a non-deductive inference to what the speaker meant are:

(3) There was a robbery on the northeast corner of 8th and Main St.;

(4) The testifier uttered (U) ('There was a robbery on the corner of 8th and Main St.')

[40] Can we say more about the availability in play here? Suppose (a) I can see my cufflink but haven't yet recognized it and so don't yet know that it's there; all I need do is turn my attention in the right direction and my power of recognition will do the rest. Compare that with: (b) my cufflink is in front of me but my eyes are closed; I can easily open them, and if I do, I'll see the cufflink, and so be in the same position as in (a). In (a) my knowledge of the cufflink's whereabouts are available given my capacity to recognize the item I visually perceive; whereas in (b), my visual perceptual capacity is not even engaged yet, but it is needed to direct my recognitional capacity. We suggest that the interpreter who has come to know from testimony has a kind of access to what was said that is more akin to that in (a) than that in (b). A helpful analogy may be a Closure principle for knowledge (roughly, that by knowing p, competently deducing q and coming to believe q while retaining knowledge of p, one thereby knows q): when p entails q, simply knowing p isn't enough to know q, because one might not consider whether q, or one might not competently deduce it from p. But given Closure, one is in a position to know q, if one knows p. There are details to work out but perhaps the analogy is helpful for testimony in the following sense: one has to be in a position to know that the testifier said that p in order to be able to know p on the basis of the testifier's speech but need not actually know that the testifier said that p. Instead, as with Closure, one is in a position to know what one's knowledge entails only if one attends to the matter and uses the right process to reach the conclusion. Cf. Williamson 2000, 117; Hawthorne 2004, 34ff. and 2014; and Kvanvig 2006 for discussion of closure principles. See also Yli-Vakkuri and Hawthorne 2022 for a discussion of the flexibility of 'in a position to know.' (Thanks to a referee here.)

[41] Being in this position feels comparable to what some prefer to call "propositional justification" (or *ex ante* justification, as in Goldman 1979); we shall leave it to those with such preferences to decide whether that terminology is apt in this case.

In reasoning to the conclusion that the testifier's (U) meant (3), the hearer can know this if the reasoning process is locally reliable. The inference is not a deductive inference, and so is not guaranteed in all circumstances to generate knowledge of what the testifier has said.[42]

In order to have a deduction, one would need to know some stronger premises. But if one is in a position to know these stronger premises, one would also be able to use testimonial knowledge to arrive at what the speaker said via the globally reliable process of deduction. That is, if one knows

(3*) I know that there was a robbery on the northeast corner of 8th and Main St.;

(4*) The knowledge in (3*) is based on the testifier's utterance (U),

then one can deduce that (U) meant that there was a robbery on the northeast corner of 8th and Main St. This requires higher-order knowledge of what one knows on the basis of testimony, and in some cases this higher-order knowledge may be hard to come by. It may be better to rely simply on a merely locally reliable inference involving first-order facts, rather than to ascend to the higher-order premises (3*) and (4*) in order to use a deductively valid (and hence globally reliable) inference. But in either case, knowledge of what the speaker said relies on the hearer's testimonial knowledge.

Thus the explicit knowledge of what a speaker meant can, in a natural sense, be posterior to the knowledge arrived at by testimony. That one has this knowledge, or can easily acquire it, is a necessary condition for arriving at knowledge by testimony; hearers who are not even in a position to know what a testifier has said cannot learn from them. What is not necessary is that one *rely* on this knowledge of what is said, in order to arrive at the testimonial knowledge. This is both an intuitive claim about specific cases, and a consequence of a use of locally reliable belief-forming processes in the

[42] We shall assume that it is abductive in this case.

safety model sketched here.[43]

Even minimal accounts of a hearer's knowledge of what a testifier said can differ on this point. Plausibly, understanding the speaker's testimony requires having a set of linguistic abilities and concepts that are capable of being exercised in the circumstance of the testimony. And one sort of theorist might characterize such understanding as a prerequisite to acquiring testimonial knowledge of what the testifier said, even if they do not require a hearer to proceed through such understanding to knowledge of what they said, to come to know the content testified to through what they said. By contrast, while we agree that many such linguistic capacities will typically be engaged in interpretation, on our model the hearer can use such capacities by immediately forming the belief that was expressed by the testifier, rather than by relying on an intermediary representation of what was said in order to come to know it. The hearer can exhibit linguistic understanding by exercising a reliable disposition to (directly) form true beliefs in a content suitably related to the content of the testifier's speech. And as we've argued, our model allows that a hearer's being in position to know what is said can be accounted for partly in virtue of the fact that they have already acquired testimonial knowledge.[44]

5 Conclusion

Misinterpretation, or even the risk of misinterpretation, can prevent the acquisition of knowledge by testimony. Our discussion of how it does

[43] Our view may be open to challenge from those who depart from our externalist, non-reductionist approach, by insisting on further requirements (beyond satisfying certain modal conditions) for gaining knowledge by testimony, because (e.g.) the most reliable way of doing so is by relying on knowledge of what was said. While we do not want to focus on arguing against additional requirements on knowledge beyond safety, the intuitive possibility (and frequency) of cases where hearers do not form explicit meta-beliefs about what the testifier said but nonetheless acquire knowledge by testimony, suggest that there are no additional requirements on testimonial knowledge that require such meta-beliefs.

[44] Likewise Peet (2019, 3305ff.) considers the question of what the "purely communicative preconditions for the acquisition of testimonial knowledge" are. His answer, sketched in Peet 2019, §5, relies on a modal notion of luck in coordinating truth-values of what the speaker asserts and the hearer believes that is similar to the safety-theoretic resources we use here (though he does not consider the analogy with inference).

highlights some further features of the role of interpretation in producing testimonial knowledge: it can be the result of a process that is locally reliable, and does not need rest on prior knowledge of what the testifier has said. It exhibits an analogous basing requirement for inference, just as, with inference, one need not form an explicit belief about the validity of inference. Our safety-based model illustrates some central features of interpretation; other models may be able to accommodate these, or may contain new insights of their own. But to the extent that non-externalist views depart from (some analogue of) locally reliable processes as being epistemologically relevant, those views will need to cut their own path when explaining the role of interpretation in testimonial knowledge.[45]

University of Notre Dame
University of Missouri–St. Louis

References

Audi, Robert. 2006. "Testimony, Credulity, and Veracity." In Jennifer Lackey and Ernest Sosa (eds.), *The Epistemology of Testimony*, 25–49. Oxford: Oxford University Press.

Benton, Matthew A. 2025. "The Epistemology of Interpersonal Relations." *Noûs* 59: 92–111.

Bergmann, Michael. 2006. *Justification without Awareness: A Defense of Epistemic Externalism.* Oxford: Oxford University Press.

Bezuidenhout, Anne. 1998. "Is Verbal Communication a Purely Preservative Process?" *The Philosophical Review* 107: 261–288.

Buchanan, Ray. 2013. "Reference, Understanding, and Communication." *Australasian Journal of Philosophy* 92: 55–70.

[45]We are grateful to an audience at St. Louis University, and to reading groups at the University of Washington and Seattle Pacific University; and to Brittany Gentry, John Hawthorne, Conor Mayo-Wilson, Patrick McDonald, Leland Saunders, Ian Schnee, Rebekah Rice, and two anonymous reviewers for helpful comments.

Burge, Tyler. 1993. "Content Preservation." *The Philosophical Review* 102: 457–488.

——. 2013. *Cognition Through Understanding.* Oxford: Oxford University Press.

Carroll, Lewis. 1895. "What the Tortoise Said to Achilles." *Mind* 4: 278–280.

Carter, J. Adam and Bondy, Patrick (eds.). 2019. *Well-Founded Belief.* New York: Routledge.

Coady, C. A. J. 1992. *Testimony: A Philosophical Study.* Oxford: Clarendon Press.

D'Ambrosio, Justin. 2025. "Manipulative Underspecification." *The Philosophical Review* 134: 241–284.

Deigan, Michael. 2022. "Stupefying." *Philosophers' Imprint* 22: 1–27. doi:10.3998/phimp.2117.

Fara, Delia Graff. 2001. "Descriptions as Predicates." *Philosophical Studies* 102: 1–42.

Fricker, Elizabeth. 1995. "Critical Notice: Telling and Trusting: Reductionism and Anti-Reductionism in the Epistemology of Testimony." *Mind* 104: 393–411.

——. 2003. "Understanding and Knowledge of What is Said." In Alex Barber (ed.), *Epistemology of Language*, 325–366. Oxford: Oxford University Press.

——. 2006. "Testimony and Epistemic Autonomy." In Jennifer Lackey and Ernest Sosa (eds.), *The Epistemology of Testimony*, 225–250. Oxford: Oxford University Press.

——. 2012. "Stating and Insinuating." *Proceedings of the Aristotelian Society, Supplementary Volumes* 86: 61–94.

Fricker, Miranda. 2007. *Epistemic Injustice: Power and the Ethics of Knowing.* Oxford: Oxford University Press.

Goldberg, Sanford C. 2001. "Testimonially Based Knowledge from False Testimony." *The Philosophical Quarterly* 512–526.

——. 2004. "Radical Interpretation, Understanding, and the Testimonial Transmission of Knowledge." *Synthese* 138: 387–416.

——. 2007. *Anti-Individualism: Mind and Language, Knowledge and Justification.* Cambridge: Cambridge University Press.

——. 2010. *Relying on Others: An Essay in Epistemology.* Oxford: Oxford University Press.

——. 2015. *Assertion: On the Philosophical Significance of Assertoric Speech.* Oxford: Oxford University Press.

Goldman, Alvin I. 1979. "What is Justified Belief?" In George Pappas (ed.), *Justification and Knowledge.* Dordrecht: D. Reidel.

——. 1986. *Epistemology and Cognition.* Cambridge: Harvard University Press.

Graham, Peter J. 2000a. "Conveying Information." *Synthese* 123: 365–392.

——. 2000b. "Transferring Knowledge." *Noûs* 34: 131–152.

Graham, Peter J. and Bachman, Zachary. 2020. "Counterexamples to Testimonial Transmission." In Miranda Fricker and et al. (eds.), *The Routledge Handbook of Social Epistemology*, 61–77. New York: Routledge.

Greco, John. 2020. *The Transmission of Knowledge.* Cambridge: Cambridge University Press.

Harman, Gilbert. 1986. *Change in View: Principles of Reasoning.* Cambridge: MIT Press.

Hawthorne, John. 2004. *Knowledge and Lotteries.* Oxford: Clarendon Press.

——. 2012. "Some Comments on Fricker's 'Stating and Insinuating'." *Proceedings of the Aristotelian Society, Supplementary Volumes* 86: 95–108.

——. 2014. "The Case for Closure." In Matthias Steup, John Turri, and Ernest Sosa (eds.), *Contemporary Debates in Epistemology*, 40–55. Malden: Wiley Blackwell, 2nd edition.

Hawthorne, John and Manley, David. 2012. *The Reference Book.* Oxford: Oxford University Press.

Heck, Richard G. 1995. "The Sense of Communication." *Mind* 104: 79–106.

Kaplan, David. 1989. "Demonstratives." In Joseph Almog, John Perry, and Howard Wettstein (eds.), *Themes from Kaplan*, 481–563. New York: Oxford University Press.

Kvanvig, Jonathan L. 2006. "Closure Principles." *Philosophy Compass* 1: 256–267.

Lackey, Jennifer. 1999. "Testimonial Knowledge and Transmission." *The Philosophical Quarterly* 49: 471–490.

——. 2006. "It Takes Two to Tango: Beyond Reductionism and Non-Reductionism in the Epistemology of Testimony." In Jennifer Lackey and Ernest Sosa (eds.), *The Epistemology of Testimony*, 160–189. Oxford: Oxford University Press.

——. 2008. *Learning from Words: Testimony as a Source of Knowledge.* Oxford: Oxford University Press.

Loar, Brian. 1976. "The Semantics of Singular Terms." *Philosophical Studies* 30: 353–377.

Longworth, Guy. 2008. "Comprehending Speech." *Philosophical Perspectives* 22: 339–373.

——. 2018. "Understanding What was Said." *Synthese* 195: 815–834.

Malmgren, Anna-Sara. 2006. "Is There A Priori Knowledge by Testimony?" *The Philosophical Review* 115: 199–241.

Peet, Andrew. 2016. "Testimony and the Epistemic Uncertainty of Interpretation." *Philosophical Studies* 173: 395–416.

——. 2017. "Referential Intentions and Communicative Luck." *Australasian Journal of Philosophy* 95: 379–384.

——. 2018. "Testimonial Knowledge Without Knowledge of what is Said." *Pacific Philosophical Quarterly* 99: 65–81.

——. 2019. "Knowledge-yielding Communication." *Philosophical Studies* 176: 3303–3327.

——. 2021. "Testimonial Worth." *Synthese* 198: 2391–2411.

Pettit, Dean. 2002. "Why Knowledge is Unnecessary for Understanding Language." *Mind* 111: 519–550.

Plantinga, Alvin. 1993. *Warrant and Proper Function.* New York: Oxford University Press.

Pollock, Joey. 2021. "Linguistic Understanding and Testimonial Warrant." *Erkenntnis* 88: 457–477.

——. 2024. “Content Internalism and Testimonial Knowledge.” *Inquiry* 67: 1947–1968.

Pritchard, Duncan. 2005. *Epistemic Luck.* Oxford: Oxford University Press.

Rabinowitz, Dani. 2011. “The Safety Condition for Knowledge.” *Internet Encyclopedia of Philosophy,* ISSN 2161–0002, http://www.iep.utm.edu/safety–c/.

Recanati, François. 1993. *Direct Reference: From Language to Thought.* Oxford: Blackwell Publishing.

Soames, Scott. 2005. “Why Incomplete Definite Descriptions do not Defeat Russell’s Theory of Descriptions.” *Theorema* 24: 7–30.

Sosa, Ernest. 1999. “How to Defeat Opposition to Moore.” *Philosophical Perspectives* 13: 142–153.

——. 2006. “Knowledge: Instrumental and Testimonial.” In Jennifer Lackey and Ernest Sosa (eds.), *The Epistemology of Testimony*, 116–123. Oxford: Oxford University Press.

——. 2011. *Knowing Full Well.* Princeton: Princeton University Press.

——. 2015. *Judgment and Agency.* Oxford: Oxford University Press.

——. 2021. *Epistemic Explanations: A Theory of Telic Normativity, and What it Explains.* Oxford: Oxford University Press.

Williamson, Timothy. 2000. *Knowledge and its Limits.* Oxford: Oxford University Press.

Yli-Vakkuri, Juhani and Hawthorne, John. 2022. “Being in a Position to Know.” *Philosophical Studies* 179: 1323–1339.