

## TESTIMONY, TRANSMISSION, AND SAFETY

Joachim Horvath

### Abstract:

Most philosophers believe that testimony is not a fundamental source of knowledge, but merely a way to transmit already existing knowledge. However, Jennifer Lackey has presented some counterexamples which show that one can actually come to know something through testimony that no one ever knew before. Yet, the intuitive idea can be preserved by the weaker claim that someone in a knowledge-constituting testimonial chain has to have *access to some non-testimonial source of knowledge with regard to what is testified*. But even this weaker claim has a counterexample which I develop in close connection with a safety-account of knowledge. Thus, testimonial statements can sometimes enable us to know something for which none of our informants has any source of knowledge available. I conclude that my counterexample nevertheless does not affect the core of our intuitions about testimony, although it establishes that testimony can indeed be a fundamental source of knowledge.

It is widely believed that testimony is not a fundamental source of knowledge, but only a way to transmit knowledge from an already knowing speaker to a yet ignorant hearer who thereby acquires the knowledge that is properly testified to him. For, testimonial chains that merely go in circles or terminate in informants who are themselves ignorant about what they testify just do not seem to be plausible sources of knowledge, let alone fundamental sources (cf. e.g. Audi 2002, Fricker 2006). Also, in most paradigmatic, everyday cases of testimonial knowledge there is in fact a non-testimonial source available somewhere down the chain. Just consider my knowledge that there are kangaroos in Australia, or that Fermat's last theorem is true – presumably, I know about the first fact because someone has had non-testimonial perceptual knowledge of kangaroos living in Australia at some point, and I know about the second fact because someone, namely Andrew Wiles, has actually proven it. Nevertheless, Jennifer Lackey (1999) presents some tricky counterexamples to the seemingly modest claim that at least *the first link* in every testimonial chain must be an instance of *knowledge*. There is, however, a better way to formulate the intuitive idea that testimony cannot be a fundamental epistemic source, namely that at least *the first link* in every testimonial chain has to have *access to some non-testimonial source of knowledge with regard to what she testifies*. I will elaborate on this idea below and explain why it nicely captures our intuitions about testimony and transmission and, most importantly, why it is not

susceptible to Lackey-style counterexamples. However, I think that there is another counterexample even to this weaker claim which I will develop and defend in close connection with a safety-account of knowledge. According to it, testimonial statements can sometimes put the hearer in a position to know something for which the speaker herself has no source of knowledge available. Finally, I draw the comforting conclusion that this further counterexample neither challenges the core of our intuitions about testimony and epistemic transmission, nor affects our usual epistemic dealings with testimonial sources that much. Nevertheless, if it goes through, then testimony can actually be a fundamental source of knowledge, contrary even to a fairly weak formulation of our respective intuitions.

### **1. Why is it intuitive that testimony is not a fundamental source?**

Visual perception is a paradigmatic fundamental source of knowledge. It enables us to learn ever more new and surprising facts about our environment and in many cases there is no alternative way available to learn all these things. Why is testimony intuitively so different from visual perception? For, we learn lots of new and surprising facts through testimony, too, and in many cases there is no other viable way to learn all these things, either. For example, there is no way to learn, for me, facts about every far away country on this planet or about each of the highly specialized branches of contemporary science – no way, except by the combined testimony of many first-hand informants and the chains of information that they instigate. So, what actually is it that makes testimony so different from perception? Let me try to illustrate this with a little parable. Imagine a far away country where only blind people live.<sup>1</sup> Not only are they blind, but they also love to talk to each other and they tend to tell their friends and neighbors practically everything they know almost all of the time, so that more or less everyone knows what everyone else in the community knows, too. In other words, we are imagining a kind of “Testimonial Paradise” here. But, as a matter of principle, even these highly communicative citizens of Testimonial Paradise cannot ever come to know that blood is red or that ripe bananas are yellow<sup>2</sup> because they all suffer from the same defect: they are as blind as a bat. To put it in more abstract terms, no visual information ever enters into

---

<sup>1</sup> Readers of H.G. Wells’ “The Country of the Blind” (1938) will be familiar with such a place.

<sup>2</sup> I ignore the complication here that they may not even have the relevant phenomenal concepts.

their chains of communication and so no one ever comes to learn anything about the colors, and other visual properties, of things. Here is an apt way to conceptualize the intuitive point: Testimony is a *purely reproductive* source of knowledge. This is supposed to mean that it does not generate any *new, original knowledge* – knowledge that no other interlocutor ever did possess before. Visual perception, on the other hand, is a *productive* source of knowledge because it enables the sighted to learn things which no one ever knew before.

Thus, I take our key intuitive idea about testimony and the transmission of knowledge to be the following:

- (IT) One cannot come to know through testimony what no one in the testimonial chain in question ever knew before in some other, non-testimonial way.

From this idea it follows that testimony may well be practically indispensable for finite creatures like us, but that it is not a fundamental epistemic source in the sense that a being with unlimited resources, cognitive capacities and time would be in need of any testimonial input. For, if such a being wants to know if there are any kangaroos in Australia, then it just goes there and finds out for itself, or if such a being wants to know if Fermat's last theorem is true, then it just spends a few years and proves it by itself. Thus, while such a being would still be in need of, for example, perception and logical reasoning it could conceivably do without any testimonial knowledge at all – or so the intuition goes.

## 2. A Lackey-style counterexample

However, here is a counterexample to (IT) in the spirit, even if not quite in the letter, of Jennifer Lackey (1999):

René is a tall and handsome French philosopher who cannot help but believe that an evil demon creates, in his mind, the perceptual illusion of an external world. Besides that, he is a very charming, kind-hearted and helpful guy. Now John, who is a rather short English philosopher, meets René in front of a high wall behind which a soccer game is going on. John desperately wants to know what happens in the game, since his favorite team really needs a victory today. Therefore, he asks René, who is barely able to look beyond the wall, to tell him exactly what happens in the game. Although René actually believes that the whole game, including John, is nothing but a clever perceptual illusion, he nevertheless reports to John every move with supreme accuracy. As a consequence,

John comes to know each and every minute detail of yet another Waterloo of his unfortunate team.

Two things are intuitively suggested by this case: First, that John really *comes to know* all the details about the game because René's reports are not only true, but also based on his impeccable faculty of vision; second, that the acquisition of John's knowledge is not impaired by the fact that René himself does not believe, and therefore *fails to know*, any of the things he kindly tells John (although he does believe that his reports are fairly accurate descriptions of what would be the case, if there really were an external world). To make the case even more compelling, it can be further stipulated that John actually knows that René, despite his somewhat strange philosophical views, is a highly trustworthy informant with regard to mundane matters, like soccer games, and that his faculty of vision is in excellent shape. Then, it seems, epistemological internalists and externalists alike, as well as reductionists and anti-reductionists in the epistemology of testimony, should find it equally intuitive that John comes to know many facts about the soccer game based on René's testimony. In the following, I will simply assume the validity of this intuition. Thus, the case of René and John is a counterexample to both of the following two claims: (1) That every informant in a knowledge-transmitting chain of testimonies has to have the respective knowledge herself, and (2) that at least the first link in such a chain has to have the knowledge in question. The falsity of these two claims, in turn, entails the falsity of (IT).

### 3. A weaker claim about testimony and transmission

My Lackey-style counterexample to (IT) trades on the fact, if it is indeed a fact, that knowledge implies belief. Our skeptical French philosopher René, as well as Lackey's creationist teacher who selflessly teaches evolutionary theory<sup>3</sup>, are both reliable informants even in the absence of belief in what they successfully testify. Thus, they themselves do not know – because they do not believe – what they testify to their more credulous hearers. Lackey has a second type of counterexample where the speaker does not know what she testifies because she possesses a defeater for her respective belief which, however, her hearer does not have<sup>4</sup>. The basic idea here is that only the belief

---

<sup>3</sup> Lackey (1999: 477).

<sup>4</sup> Lackey (1999: 484-485).

itself and its positive epistemic status is transmitted in the act of testimony, yet not the defeater for that belief because the speaker never mentions it. For the sake of the argument, I will simply accept that second type of counterexample as valid, too. I think, however, that the idea that testimony is a purely reproductive epistemic source can be formulated in such a way that it still does justice to our intuitions about Testimonial Paradise, yet escapes both types of Lackey-style counterexamples. It is based on the observation that the first informant in Lackey-style cases, although she does not *believe* or *know* what she testifies to her hearer, nevertheless has *access to some highly reliable source of knowledge with regard to what she testifies*. This is quite obvious in Lackey's own case of the creationist teacher who prepares for her selfless lectures on evolution "from reliable sources, and on this basis develops a set of reliable lecture notes" (Lackey 1999, 477). In fact, it is this very feature of the case that makes it palatable to us even in the face of the intuitive appeal of (IT). Equally, what makes it acceptable that John comes to know many things about the soccer game through René's testimony is precisely that I have described René's faculty of vision as "impeccable" and "in excellent shape". If we leave out this crucial emphasis on *reliable sources* from Lackey-style cases, then their power to convince us even in the face of contrary intuitions crumbles. Therefore, the following reformulation of (IT) stands to reason:

- (ITR) One cannot come to know through testimony that  $p$  if no one in the testimonial chain in question ever had access to some reliable non-testimonial source of knowledge concerning  $p$  before.

This reformulation essentially preserves the intuitive idea that testimony is not a reproductive epistemic source while it avoids, at the same time, any counterexamples of the kind that Lackey has introduced into the literature.

#### 4. Safety and a further counterexample

Are there any conceivable counterexamples to (ITR)? Asked slightly differently, what exactly would it take for a hypothetical (or actual) case to count as a counterexample to (ITR)?

It would have to be a case where some hearer  $H$  of a testimony acquires knowledge that  $p$  based on that testimony alone, while the speaker  $S$  of that testimony not

only lacks knowledge that  $p$  but also does not have any source of knowledge with regard to  $p$  at his disposal. In addition, one would have to make the comparatively unproblematic assumption that this is the limit case of a testimonial chain with S as its only link, because (ITR) merely requires that at least one link in the testimonial chain has had access to some source of knowledge concerning  $p$ .

The immediate worry that arises from these constraints is that they specify a case where there is no sufficient epistemic connection between H's resulting belief that  $p$  and the fact, or state of reality, which makes it the case that  $p$  is true. For, if speaker S has no source of knowledge concerning  $p$  at his disposal (and, we may assume, no relevant evidence or source of justification either), then how can it be anything but an *accident*, or a matter of luck,<sup>5</sup> if hearer H nevertheless acquires a true belief with the content that  $p$  solely based on S's testimony? Indeed, this still appears somewhat mysterious even if we are told that S is, in general, a trustworthy and reliable informant. It looks, at any rate, like a straightforward case of a merely *accidentally true belief* and thus it would not qualify as *knowledge* – which, however, it must in order to falsify (ITR).

We can thus refine the challenge for finding a counterexample to (ITR) as follows: What, if anything, may conceivably prevent our hearer H's belief that  $p$  from being merely *accidentally true* when it is solely based on a testimony that is neither an expression of knowledge that  $p$  nor backed up by any source of knowledge concerning  $p$ ? It can, I submit, only be the fact that at least S's true testimonial *statement* with content  $p$  is related to whatever makes it the case that  $p$  in some non-accidental way. This shift from mental states to utterances seems possible because Lackey-style counterexamples to (IT) already suggest that the epistemic status of a speaker's own beliefs and the epistemic force of her verbal statements may indeed come apart. Thus, Lackey herself concludes that the primary sources of testimonial knowledge are the *statements* of informants – and *not* their *beliefs*<sup>6</sup>.

But how can the epistemic resources of S concerning  $p$  and the epistemic properties of her respective testimonial statement with content  $p$  come apart in such a radical way? And, moreover, how can this happen without making the statement merely acci-

---

<sup>5</sup> I am only concerned here with the type of accidentality that Pritchard (2005) dubs “veritic luck”. The more subjective variety of accidentality which he calls “reflective luck” is less important for our discussion, because knowledge is not a (merely) subjective, or internal, epistemic property.

<sup>6</sup> Lackey (1999: 489; 2006: 96-97).

dentally true? To clarify this issue we need some kind of standard, or criterion, that helps us to distinguish accidental truth from non-accidental truth. For this purpose, we can draw on the resources of *safety-accounts* which lend themselves quite naturally to anti-luck conceptions of knowledge<sup>7</sup>. Simply put, the basic idea behind safety is that if someone *knows* that *p*, then she *could not easily have been wrong*. Thus, I propose the following two safety-principles as necessary conditions on knowledge in order to rule out merely accidentally true beliefs<sup>8</sup> and I also assume that these two principles are sufficient for knowledge if conjoined with the two standard conditions that the agent in question believes that *p* and that it is actually true that *p*<sup>9</sup>:

(SA1) If an agent knows that *p*, then, in all nearby possible worlds in which she forms her belief that *p* in the same way as she forms her actual belief that *p*, *p* is true.

(SA2) If an agent knows that *p*, then, in all nearby possible worlds in which she forms a different belief that *q* in the same way as she forms her actual belief, *q* is true.

Not only do these two principles exclude Gettier-scenarios of all kinds but they also handle a number of other epistemological problem cases very well, e.g. the lottery paradox<sup>10</sup>. Furthermore, simple safety-accounts have a notorious problem with beliefs in necessarily true propositions because these beliefs could *never* have been wrong, no matter how unreliably one may acquire them. The above conditions, however, avoid this problem first, by focusing on the safety of the actual *way* in which the belief in question was formed, and second, by also requiring the truth of all the *alternative beliefs* that the actual way of belief formation might have easily produced. Then, for example, a lucky guess of a mathematical truth is ruled out as a case of knowledge because it violates condition (SA2). For, by just making a lucky guess one would have believed some mathematical falsehood in many nearby possible worlds. Moreover, a simpler safety-account with (SA1) as its only condition cannot be simply cured by restricting it to our knowledge of *contingently true* propositions.<sup>11</sup> Just consider my belief that there is

<sup>7</sup> As elaborated by Pritchard (2005).

<sup>8</sup> Sainsbury (1997), Sosa (2002), Pritchard (2005: chapter 6).

<sup>9</sup> Any worries one might have that this assumption of a pure safety-account of knowledge makes the scope of my counterexample to (ITR) too limited will be addressed at the end of this section.

<sup>10</sup> Pritchard (2005: 162-173).

<sup>11</sup> I think that this is a serious shortcoming of Pritchard's ultimate formulation of the safety-principle (Pritchard 2005, 163), since it is not even sufficient (together with truth and belief) for knowledge of *contingent* propositions. Furthermore, it also does not suffice to rule out *accidentally true beliefs* in contingent propositions (i.e. "veritic luck"), as my carbon-example in the text illustrates. Thus, Pritchard's

carbon on earth. Surely, the proposition that there is carbon on earth is contingent. However, a lot of things would have to be very different for such a belief to be false – including the inexistence of human beings on earth who themselves consist of carbon to about 20 percent. Thus, there is no nearby possible world in which I could falsely believe that there is carbon on earth, no matter how *irrational* my belief may be formed. Since it is very implausible that each of those beliefs constitutes knowledge, condition (SA2) has to be added in order to rule out such cases.

So, given both (SA1) and (SA2) as our criteria for non-accidental truth and knowledge, can the epistemic resources of S with respect to *p* and the epistemic force of her corresponding testimonial statement come apart as radically as indicated above? Only if, it seems, the *safety* of S's statement is not affected by the lack of any (non-testimonial) source of knowledge concerning *p* on the part of S.<sup>12</sup> And this, in turn, can only be so if that *statement* (or, alternatively, a belief based on that statement) has some epistemically relevant feature which is lacking in S's own epistemic situation. The following seems to be a case which intuitively fills that 'job description':

Suppose that Peter is a perfect informant. He only tells other people that *p* if it is really the case that *p*, regardless of whether he himself actually believes that *p* or not. Now suddenly, a cosmic event C causes Peter to verbally utter the following statement in the living room of his friend Thomas: "Sound waves with frequencies between 20 Hz and 20 kHz occur." C is such a rare event that it happens only once in the history of the universe, that is, it only happens to Peter at the time when he is sitting comfortably in Thomas's living room. Since Thomas has no reason to distrust Peter, he comes to believe what Peter tells him, namely that (T) sound waves with frequencies between 20 Hz and 20 kHz occur. Now, it is in fact true that sound waves with frequencies between 20 Hz and 20 kHz occur because that is exactly what happens whenever a human being produces a verbal statement. Therefore, Thomas comes to know that (T) solely based on Peter's testimonial statement.

If this is supposed to be a real counterexample to (ITR), then basically two things have to be the case:

- (1) Thomas really comes to know that (T) based on Peter's testimony, that is, his resulting testimonial belief ought to be true and safe according to (SA1) and (SA2).

---

safety-principle fails with regard to its central aim: to give an account of *non-accidentally true belief* at least for beliefs in contingent truths.

<sup>12</sup> Goldberg (2007) argues for the related, yet importantly different claim that there can be safe "consumptions" of testimonial statements which are themselves not safe. Our case of Peter and Thomas below is different, though, for it involves a safe "consumption" of a *safe testimonial statement* that is not the expression of a safe belief.

- (2) Peter is the first link in the testimonial chain in question and he has not had access to any non-testimonial source of knowledge concerning (T).

Let me address these two points in turn. As far as (1) is concerned, Thomas comes to know that (T) based on Peter's testimony because he satisfies both (SA1) and (SA2). There is no nearby possible world where Thomas bases his belief that (T) on Peter's testimony and (T) is false. For, Peter is a perfect informant and his statement of (T) is self-verifying as a verbal utterance. So, Peter's statement, as well as Thomas's belief, could not easily have been false.<sup>13</sup> Consequently, Thomas satisfies (SA1). Yet, he also satisfies (SA2) because any other belief he might have acquired by listening to Peter would have been true as well. Furthermore, Thomas could not easily have acquired a different and false belief instead of (T) due to some other cosmic event like C that might have caused Peter to testify something *false*. For, we have conceived of C as an extremely rare kind of event which may even be nomologically impossible to repeat. So, C could not easily have happened again or in some other, similar way – and thus nothing could easily have caused Peter to make a false statement in Thomas's living room. Therefore, Thomas does not form a different, and false, belief in any nearby possible world and thus he satisfies (SA2) as well.

Concerning (2), we have simply set up the case in such a way that Peter is the only link in a (maximally short) testimonial chain that results in Thomas's belief that (T). Furthermore, Peter's statement with content (T) is not backed up by any source of knowledge whatsoever, testimonial or otherwise. But why can't we just say that cosmic event C *itself* counts as Peter's source of knowledge? Because, for something to count as a source of knowledge it ought to be capable of producing more than just a single output in some close possible world. Event C, however, only causes Peter to produce one and the same statement in all close possible worlds, including the actual world. Therefore, its epistemic output is simply too meager for it to qualify as a *source* of knowledge in any substantial sense of the concept of a source. Nevertheless, if event C were to cause Peter himself to believe that (T) – or, rather, some mental equivalent of (T) –, he might likewise satisfy (SA1) and (SA2). But such a case would imply that *source-*

---

<sup>13</sup> Please note an important feature of the case here: that Thomas bases his belief that (T) *solely* on Peter's testimonial statement. For, if Thomas had e.g. come to believe that (T) because of his realization that Peter's statement was self-verifying, then this would, of course, not be a case of testimonial knowledge *proper* anymore.

*less* knowledge is possible. However, if sourceless knowledge were actually possible, then (ITR) would be very easy to falsify. For then, Peter could in principle know something that is not based on any source. But since there are no special problems for Peter to testify something that he already knows, it would of course be very easy for him to *transmit* that knowledge. Yet, this transmitted piece of knowledge, let's call it (MS), would not be such that Peter has had access to some reliable source of knowledge concerning (MS), as (ITR) requires. So, if Peter were the only link in the testimonial chain in question, then every piece of *sourceless* knowledge that he testifies to someone would constitute a counterexample to (ITR). Note, however, that such cases would not also falsify (IT), the first intuitive formulation of the transmission principle. Thus, we would then have counterexamples to (ITR) that are not also counterexamples to (IT), and we would have Lackey-style counterexamples to (IT) that are not in turn counterexamples to (ITR). Both formulations of the transmission principle would still be wrong, then, but because of very different cases – whereas the original case of Peter and Thomas falsifies both (IT) and (ITR) in a single stroke. Moreover, there might be an independent argument against the possibility of sourceless knowledge. The original case would still be more effective, then, for it does not involve any sourceless knowledge because Peter himself only makes a *testimonial statement* and Thomas's belief has *Peter's testimony* as its source.

So, if we accept the two safety-principles (SA1) and (SA2) as necessary (and, together with the belief-condition and the truth-condition, as sufficient) conditions on knowledge, that is, on non-accidentally true belief, then we also have to accept the case of Peter and Thomas as a counterexample to (ITR). Testimonial knowledge is possible, it seems, even in the absence of a non-testimonial source as its ultimate ground. In the following, I will discuss some worries and objections that may be prompted by this counterintuitive result.

First, the somewhat underdetermined notion of 'ways of belief-formation' plays a crucial role in (SA1) and (SA2). But what actually is Thomas's way of forming his belief that (T)? We typically referred to his belief as a *testimonial* belief. So, is *being based on testimony* the way in which Thomas forms his belief? In that case, can there *ever* be a safe testimonial belief? For, if one simply bases a belief on *testimony*, without any further qualifications or restrictions, one could easily have formed a false belief

because of the many unreliable informants in one's social environment, and without some criterion of trustworthiness at hand one could not easily have avoided one of those highly fallible informants. However, there are good reasons why we have to individuate ways of belief-formation rather *fine-grained* anyway when it comes to knowledge. Because otherwise, not even perception would qualify as a source of knowledge. For example, if I individuate my way of coming to believe that there are two rabbits in front of me simply by *looking at the things before my eyes*, then even this Moorean belief will not be safe, and thus fail to count as knowledge, because there are, of course, many nearby possible worlds where merely *looking at things*, without any further qualifications, will eventually deliver a false belief because of unfavorable perceptual circumstances. So, in order to avoid external world skepticism, we have to individuate perceptual ways of belief-formation pretty fine-grained as well, e.g. by adding qualificatory clauses concerning lighting conditions, the distance of things, the state of mind of the observer, and so on. But on any plausible more *fine-grained* individuation of Thomas's way of coming to believe that (T) his belief will actually be safe, and therefore count as knowledge, too. If we individuate it, for example, as *trusting Peter's testimony*, then Thomas's belief will be safe because Peter is a perfect informant. Yet, even if we individuate it extremely fine-grained as *trusting Peter's verbal statement of (T)*, then Thomas's belief will still be safe because there simply is no nearby possible world where Peter's statement is false, for every verbal human statement essentially involves the production of sound waves with frequencies between 20 Hz and 20 kHz. Therefore, given human physiology, Peter's statement is *self-verifying* and thus guarantees its own truth. To sum up this point, any non-skeptical epistemology must individuate ways of belief-formation in a *fairly fine-grained* way when it comes to knowledge, and any plausible fine-grained way in which Thomas may come to believe that (T) is a safe way according to (SA1) and (SA2).<sup>14</sup>

Second, since the ultimate cause of Thomas's belief that (T) is some cosmic event C, why isn't his belief *merely accidentally true*? Although prima facie plausible,

---

<sup>14</sup> What the exact criteria of individuation of ways of belief-formation may be is, of course, a difficult question that I cannot deal with satisfactorily within the limited scope of this paper. One might try to individuate some of the more 'natural' ways of belief-formation, like e.g. perception or memory, by their biological function in the type of environment for which they have been selected. Or one might concede some kind of context-relativity for the individuation of ways of belief-formation. For the purposes of my argument it suffices, however, that Thomas comes to know that (T) on any prima facie plausible way of forming his belief that (T).

this worry turns out to be extremely difficult to pin down. It won't do to claim that Thomas's belief is merely accidentally true because it could easily have been false – this explanation has already been ruled out by the fact that it satisfies (SA1) and (SA2). It also won't do to point out that there is some *distant* possible world where event C occurs but nevertheless causes Peter to state something false instead of (T). For, there almost always is *some* distant world where the actual way in which one forms a true belief results in some false belief or other. This seems *metaphysically possible* in almost any case. All of my current perceptual beliefs, for example, are false in a distant possible world where my brain was just envatted by evil scientists five minutes ago. Thus, in some sense all perceptual beliefs are merely accidentally true because they all could have been false in some metaphysically possible world. However, if we insist on spelling out non-accidentality in such a rigorous way, anything less but certainty will not be good enough for knowledge. Yet, in that case human beings will not be good enough for knowledge either and total skepticism ensues.

In reaction, one may try to spell out non-accidentality in an even stronger way, e.g. by demanding a *robust metaphysical relation* between a true belief and the fact, or state of reality, which makes it true. The most obvious candidate here is causality. But, as epistemologists have learned from the failure of the causal theory of knowledge, a causal relation between a true belief and its corresponding fact is neither necessary nor sufficient for knowledge. What if we merely demand *some* robust metaphysical connection or other between belief and reality? On such a vague construal of non-accidentality, however, it is anything but clear that Thomas's belief that (T) does not satisfy it. For, the ground on which he bases his belief, namely Peter's statement with content (T), *is* indeed metaphysically tied to the fact which makes it true (i.e. that sound waves with frequencies between 20 Hz and 20 kHz occur) by the relation of *being self-verifying*. Peter's statement is related to its own occurrence in such a way that it must be true if it occurs. So, given that it actually does occur, it simply could not have been false. What more can we expect of an epistemic ground to make it the proper basis for a non-accidentally true belief?

Third, even if it may be hard to pin down what makes the case of Peter and Thomas bothersome, one may still have the strong *intuition* that Thomas just cannot come to know that (T) in such an awkward way. However, if one were to insist on that

intuition, then one would have to accept the case not as a counterexample to (ITR) but as a counterexample to the safety-account of knowledge itself, that is, as a counterexample to (SA1) and (SA2). What makes this move quite troublesome, however, is that safety is probably our best attempt to spell out the notion of non-accidental truth, and since many philosophers agree that knowledge is nothing but non-accidentally true belief, safety is one of our best theories of knowledge to date. Thus, we face the following dialectical situation here: The case of Peter and Thomas satisfies (SA1) and (SA2) and *if* these principles (together with the belief-condition and the truth-condition) make for a good theory of knowledge, then (ITR) must go. On the other hand, if we *insist* on our intuition that Thomas does not come to know that (T), then we can keep (ITR), yet our safety-account of knowledge is doomed.

Let me remind you at this point how theoretically powerful the above safety-account of knowledge with principles (SA1) and (SA2) actually is, for it deals so well with many central epistemological problem cases, e.g. Gettier scenarios or lottery paradoxes.

Still, one may worry that I have only shown (ITR) to be incompatible with a *pure* safety-account of knowledge, that is, an account that takes (SA1) and (SA2) – in conjunction with the belief-condition and the truth-condition – to be already *sufficient* for knowledge. Thus, one may worry that this limits the scope of my argument significantly, for many epistemologists might well agree that something like safety is *necessary* for knowledge (in order to rule out non-accidental truth), yet they might not agree that it is also *sufficient* for knowledge.<sup>15</sup> While I do think that a pure safety-account is indeed quite attractive, I want to consider some plausible additional conditions on knowledge in the following and see if they somehow jeopardize my counterexample to (ITR).

So, what if we impose *justification* as a further condition on knowledge, in addition to belief, truth and safety along the lines of (SA1) and (SA2)? That is, does Thomas' belief that (T) not only count as safe and true, but also as *justified*? This depends, to some degree, on the assumption of reductionism or non-reductionism about testimonial justification. A non-reductionist will probably accept the case as an instance of testimonial justification as it stands, for Thomas' informant Peter is in fact highly

---

<sup>15</sup> Special thanks to Aidan McGlynn for pressing this point, and for suggesting a feasible solution as well.

reliable and Thomas has no defeaters with regard to his testimony (he has “no reason to distrust Peter”, as I have put it above). A reductionist about testimonial justification, however, will presumably demand that Thomas possesses some non-testimonial reasons that make the truth of Peter’s testimony likely. Thus, it might suffice that we stipulate that Thomas has had enough positive experiences with Peter’s testimonies in the past in order to justify him in believing that he is indeed a credible and reliable informant. Therefore, it seems, even if we impose an additional justification-condition on knowledge we can easily construct a mild variant of the case that still constitutes a counterexample to (ITR).

Similar considerations apply if we demand, as a further condition on knowledge, that Thomas bases his belief that (T) on good *evidence* for the truth of (T). For, on most non-reductionist conceptions of testimony the sincere statement of an objectively reliable informant will already constitute good evidence to believe what the informant testifies, given the absence of any counterevidence. All these constraints seem to be satisfied by the case of Thomas and Peter as it stands. If one is a reductionist about testimony, however, then one can simply stipulate, as above, that Thomas actually has good evidence to believe that Peter is a reliable and trustworthy informant. So, even if we add some kind of evidentialism as a further constraint on knowledge – in addition to safe true belief – the case of Peter and Thomas still looks like a counterexample to (ITR), or can easily be amended in the required way.

I cannot, of course, consider all possible additional constraints on knowledge that any epistemologist has ever imposed (or might plausibly impose) within the limited scope of this paper. But I have argued that two of the theoretically most important and influential candidates, namely *justification* and being based on good *evidence*, merely present some minor complications for my intended counterexample to (ITR). Thus, I can see no reason why other plausible, non ad-hoc constraints on knowledge should lead to a significantly different result here.

Would it, then, not be quite reasonable to simply accept our case as a counterexample to (ITR)? At any rate, it seems that we face the following *dilemma* here:

- (H1) The case of Peter and Thomas is a counterexample to the highly intuitive principle (ITR).
- (H2) The case of Peter and Thomas is a counterexample to the theoretically powerful safety-account of knowledge with principles (SA1) and (SA2), and it is also a

counterexample (with slight modifications) to any plausible conjunction of such a safety-account with further non ad-hoc conditions on knowledge.

Whichever horn we choose, it seems, there will be a significant price to pay. However, in the concluding section I offer some reasons why the intuitive costs of accepting the first horn are much lower than one might expect. Given that the theoretical costs of an outright rejection of the safety-account of knowledge with (SA1) and (SA2), and of any conjunction of that account with plausible further conditions on knowledge, are indeed substantial, we should then simply give up (ITR).

### 5. Comforting conclusion

We should, I claim, simply accept (H1) and concede that even a fairly weak formulation of our intuitions about testimony and transmission, like (ITR), is not free of counterexamples and therefore not a necessary truth about testimonial knowledge. It follows that testimony can indeed be a fundamental source of knowledge although typically it is not. Our intuitions about Testimonial Paradise are therefore *not* intuitions of *necessity*, yet only intuitions of *typicality*. As the empirical research on the psychology of concepts has shown, such intuitions of typicality are actually fairly pervasive<sup>16</sup> and our tendency to mistake them for intuitions of necessity is thus unsurprising – especially if a certain feature is so *overwhelmingly typical* as the purely preservative use of testimony. For, the above counterexample to (ITR) notwithstanding, in the vast majority of knowledge-constituting testimonial chains there will actually be someone down the chain who has (or has had) access to some non-testimonial source of knowledge with regard to what is testified. Thus, in almost all actual instances of testimonial knowledge the intuitive principle (ITR) will in fact hold. That is, (ITR) has an objective probability of almost 1 – which is the grain of truth in our intuitive reactions to the case of Testimonial Paradise. Accordingly, all of the following inferences regarding testimonial knowledge are indeed highly probable (with increasing probability from I1 to I3), although none of them is actually necessary:

- (I1) If *S* knows that *p* based on testimony, then someone in the testimonial chain preceding *S*'s belief that *p* also *knows* that *p*.

---

<sup>16</sup> Murphy (2004).

- (I2) If  $S$  knows that  $p$  based on testimony, then someone in the testimonial chain preceding  $S$ 's belief that  $p$  also *believes* that  $p$ .
- (I3) If  $S$  knows that  $p$  based on testimony, then someone in the testimonial chain preceding  $S$ 's belief that  $p$  also has (or has had) *access to some non-testimonial source of knowledge*.

So, even if (ITR) is strictly speaking false because of counterexamples like the one discussed in section 4, it still remains *typically true*. Therefore, the intuitions that led to our acceptance of (ITR) as a necessary truth about testimony in the first place can largely be preserved under the slightly different guise of typicality intuitions, even if we must ultimately deny the necessity of (ITR). However, compared to an outright rejection of all plausible accounts of knowledge with safety along the lines of (SA1) and (SA2) as an ingredient, this seems to be a far more reasonable theoretical price to pay.<sup>17</sup>

**Joachim Horvath**

*University of Cologne*

joachim.horvath@uni-koeln.de

---

<sup>17</sup> This paper grew out of Thomas Grundmann's and David Schweikard's graduate seminar on social epistemology at the University of Cologne in the winter semester of 2007/2008. Special thanks to Thomas Grundmann for incisive criticism, helpful suggestions and permanent encouragement throughout the process of writing this paper. I would also like to thank Alva Noë and the other participants of the *1<sup>st</sup> Bochum Graduate Conference in Philosophy* on February 14, 2008 for stimulating discussions. I am especially grateful to Aidan McGlynn (and some anonymous referees for the *2008 University of Texas at Austin Graduate Philosophy Conference*) for helping me to eliminate a serious weakness in the presentation of my argument.

## References

- Audi, R. (2002) 'The Sources of Knowledge' IN P. K. Moser (ed.), *The Oxford Handbook of Epistemology*, Oxford: Oxford University Press, pp. 71-94.
- Fricker, E. (2006) 'Testimony and Epistemic Autonomy' IN J. Lackey and E. Sosa (eds.), *The Epistemology of Testimony*, Oxford: Oxford University Press, pp. 225-250.
- Goldberg, S. (2007) 'How lucky can you get?', *Synthese* 158, 315–327.
- Lackey, J. (1999) 'Testimonial Knowledge and Transmission', *The Philosophical Quarterly* 49, 471-490.
- Lackey, J. (2006) 'Learning from Words', *Philosophy and Phenomenological Research* 73, 77-101.
- Murphy, G. (2004) *The Big Book of Concepts*, Cambridge (Mass.): MIT Press.
- Pritchard, D. (2005) *Epistemic Luck*, Oxford: Oxford University Press.
- Sainsbury, R. M. (1997) 'Easy Possibilities', *Philosophy and Phenomenological Research* 57, 907-919.
- Sosa, E. (2002) 'Tracking, Competence, and Knowledge' IN P. K. Moser (ed.), *The Oxford Handbook of Epistemology*, Oxford: Oxford University Press, pp. 264-284.
- Wells, H.G. (1938) *The Country of the Blind and Other Stories*, New York: Garden City.