

How lucky can you get?

Sanford Goldberg

Published online: 10 October 2006
© Springer Science+Business Media B.V. 2006

Abstract In this paper, I apply Duncan Pritchard's anti-luck epistemology to the case of knowledge through testimony. I claim (1) that Pritchard's distinction between veritic and reflective luck provides a nice taxonomy of testimony cases, (2) that the taxonomic categories that emerge can be used to suggest precisely what epistemic statuses are transmissible through testimony, and (3) that the resulting picture can make clear how testimony can actually be knowledge-generating.

Keywords Testimony · Transmission · Luck · Justification · Knowledge

1

Duncan Pritchard's recent book *Epistemic Luck* (Pritchard, 2005) is a most welcome addition to the contemporary scene in epistemology. In addition to being a systematic exploration of the doctrine that knowledge is incompatible with lucky belief (something often taken to be a truism), Pritchard exploits the anti-luck epistemology he develops to great benefit, examining at length such issues as the nature and challenge of skepticism, as well as internalism versus externalism in epistemology. It is a testament to the fruitfulness of the anti-luck epistemology he develops, that it has so much to say about these and other topics. In this paper, I propose to employ the anti-luck orientation to consider a domain of knowledge he does not systematically explore in the book: knowledge through testimony. I aim to show that the anti-luck epistemology complicates the picture of such knowledge, in ways that should elicit further discussion in the epistemology of testimony literature.

There are two notions of epistemic luck characterized in Pritchard (2005). Pritchard describes a subject *S*'s belief that *p* as *veritically lucky* (henceforth, *V*) just in case 'it

S. Goldberg (✉)
Department of Philosophy, University of Kentucky,
1427 Patterson Office Tower,
Lexington, KY 40506-0027, USA
e-mail: scgold@uky.edu

is a matter of luck that the belief is true' (p. 123). In a nutshell, veritically lucky belief is unsafe belief. The second notion is that of reflective luck: Pritchard describes S 's belief that p as reflectively lucky (henceforth, R) just in case 'given what the agent is able to know by reflection alone, it is a matter of luck that the belief is true' (p. 140). The rough idea here is that, whereas veritic luck is a matter of what could easily have happened, where this is determined by an ordering of possible worlds arranged according to how *objectively* similar each is to the actual world, reflective luck is a matter of what *by the subject's own best lights* could easily have happened, where this is determined by an ordering of possible worlds arranged according to how similar to the actual world *the subject herself* takes the various possible worlds to be.

So formulated, any belief can be assessed along the dimension of these two epistemic luck properties, V and R . But beliefs acquired through testimony would appear to offer the chance to see whether, and if so how, such properties can be *transmitted through speech*. In exploring this issue, we might begin with the question: how might it come to pass that one acquires a testimonial belief, which itself is (true but) reflectively or veritically lucky?¹ In general, there are two such ways: either S 's belief is lucky because she relied on testimony, which itself was lucky (where S 's belief inherits its luckiness from that of the testimony she consumed²), or else S 's belief is lucky because luck was generated somewhere in the process of communication. I will call the first sort of case a *transmission* case, and the second sort of case a *generation* case. I consider transmission cases in Sect. 2, and generation cases in Sect. 3. I make some programmatic conclusions in Sect. 4 about the utility of anti-luck epistemology for thinking about the nature of testimony.

2

To begin, let us consider transmission cases. Let us describe as a *V-transmission* case any case, in which a hearer's belief that p , acquired through a piece of testimony that is V , is itself V by inheritance. To illustrate, suppose S 's belief that p —a belief she expresses in her testimony that p —is true but lucky: it could easily have come to pass that S acquired the belief that p using the very method M she used in the actual world, yet $\sim p$. Now suppose further that in the actual world hearer A sees nothing remiss in S 's testimony that p , and so trustingly accepts that testimony and as a result acquires the testimonial belief that p . In that case A 's actual belief that p is veritically lucky as well: it could easily have come to pass that A acquired the belief that p using the very (testimonial) method M she used in the actual world, yet $\sim p$. This can happen as follows. In most of the nearby worlds in which $\sim p$, where S acquires the belief that p using the very method M she used in the actual world, S goes on to testify that p . In at least some of these worlds, the testimony S offers will strike A exactly as it struck A in the actual world, which is to say that in such worlds A trustingly accepts S 's testimony. But since p is false in such worlds, the result is that in each of these worlds A 's testimonial belief that p is false. Since those worlds are nearby, the result is

¹ Henceforth I will use 'lucky' to mean 'reflectively or veritically lucky'.

² I will assume that S 's testimony that p is V (veritically lucky) or R (reflectively lucky) if (a) S 's testimony that p reflects S 's belief that p , and (b) S 's belief that p is V or R . This is a sufficient condition, not a necessary one: there are other ways in which it can come to pass that S 's testimony that p is V or R . But as the sufficient condition presented is the paradigmatic case, this will be the only case I will consider.

that in the actual world A 's testimonial belief that p is veritically lucky—as we might put it, that belief has simply inherited the V -ness of S 's testimony.

The foregoing is perhaps the paradigmatic case of V -transmissive testimony. In this respect, reflective luck contrasts quite starkly with veritic luck: the former is not the sort of thing that can be transmitted by testimony (as opposed to *generated* by the unreflective acceptance of testimony; see Sect. 3 below). That reflective luck is non-transmissible can be established by showing that whether a hearer's testimonial belief is R does not turn on any of the features of the testimony she consumed, and so does not turn in particular on whether the testimony she consumed was R . This can be brought out as follows. A hearer A 's testimonial belief is R if, for all that A herself can tell by reflection alone, it could have easily happened that the proposition A believes is false yet A still forms the belief in that proposition through trusting acceptance of a piece of testimony. Now this condition holds just in case, for all A knows by reflection, there are nearby worlds in which A encounters testimony to the effect that p , A finds this attestation credible and so trustingly accepts what was attested, and yet it is false that p . But whether there are such worlds is independent of anything *the testifier* might know by reflection, and so is independent of whether *the testimony* is R or not. In this way we see that R -ness is not the sort of feature that is transmissible in testimony. Rather, R -ness is generated by, e.g., the unreflective consumption of testimony. To wit: if, for all that a hearer knows through reflection, it could easily be the case that the speaker got things wrong, or is lying or otherwise insincere, then (no matter what *the speaker* can and cannot rule out by reflection) the hearer cannot eliminate these possibilities by reflection, and so her belief is reflectively lucky if true at all.

Indeed, we can make the same point—that R -ness is not the sort of property that is transmissible through testimony—by contrasting what is at issue in R -ness and V -ness. V -ness is transmissible since what makes a belief veritically lucky (V) has to do with the modal profile of a belief, and in particular with its profile across a range of nearby possible worlds. That is, it has to do with whether it could easily have come to pass that the subject acquires the same belief, using the same method, in a world in which that belief is false. Now, the modal profile of a belief acquired through testimony is typically *inherited* from that of the testimony itself: if I accept your telling, and so come to believe what you tell me (where what you tell me is something you believe), and if I am a reliable discriminator of reliable from unreliable testimony, I acquire a belief that is approximately as safe as your testimony.³ This is precisely why V -ness can transmit in the first place: were I to accept your testimony, any lack of safety in your testimony would (typically; see below) infect my testimonial belief with a corresponding lack of safety. Of course, the same does not hold with respect to R -ness: were I to accept your testimony, any lack in the testimony along the R -dimension will not by itself infect my belief with a corresponding lack along the R -dimension.

This contrast is sufficiently important as to be worth putting the point a bit more formally. I begin with, what I will call the *infectious nature of epistemic perilousness*, where 'perilous' is used to mean 'not safe'. Let us designate the speaker S 's grounds for her belief that p as $G(S)$. These grounds will enable her to rule out certain error-possibilities (that is, certain possible scenarios, in which $\sim p$). Let $RA(S)$ designate the set of alternatives relevant to S 's belief that p . If $G(S)$ eliminates all of the alternatives in $RA(S)$, then S 's belief amounts to knowledge; if not, not. Now A 's

³ 'Approximately' as safe: the fact that I acquired my belief through a communication channel introduces a source of potential error that your belief (expressed in your testimony) did not face. I develop this point below.

belief formed through S 's testimony will have as its grounds $G(A)$. Letting $RA(A)$ designate the set of alternatives relevant to A 's belief that p , it is natural to suppose that $RA(A)$ has as a subset the members of $RA(S)$. That is, it is natural to suppose that the nearby error-possibilities in which S falsely testifies that p are nearby error-possibilities for A 's testimony-based belief (formed through accepting S 's testimony that p). I submit that this assumption is acceptable so long as we can vary the truth-value of S 's testimony without varying the features, on which A 's acceptance of S 's testimony supervenes. Now normally A 's testimonial grounds $G(A)$ cannot eliminate any error-possibilities in $RA(S)$ not already ruled out by S 's (non-testimonial) grounds $G(S)$. So if (as I have suggested) $RA(S)$ is a subset of $RA(A)$ any relevant error-possibility not ruled out by the testifier's grounds will ipso facto be a relevant error-possibility not ruled out by the hearer's grounds either. (At least this will hold in the normal cases; see Sect. 3 below for possible non-standard cases.) It is this, I submit, that underlies the infectious nature of epistemic perilousness.

If the foregoing provides the basis for V -transmissibility in testimony cases, consider that R -ness has to do not with the actual modal profile of a belief, but rather by the modal profile of a belief in a context, in which we are ordering possible worlds according to what the subject knows by reflection. Since a speaker S cannot control what a hearer A knows by reflection, S cannot control the possible-world-ordering determined by what A knows by reflection. It follows that this possible-world-ordering is not affected by whether S 's testimony is, or is not, itself R . Since it is this possible-world-ordering that determines whether A 's testimonial belief is R , we see that S is not in a position to affect this by, e.g., offering testimony that is non- R . Again, we see that reflective luckiness is not the sort of thing that can be testimonially transmitted.

So far I have given reason to think that V -ness but not R -ness can be testimonially transmitted. What then can we say about cases in which the testimony is clean, that is, neither V nor R ? Can cleanliness transmit? The foregoing results would lead us to speculate that there are cases in which non- V -ness (like its complement, V -ness) transmits, but no cases in which non- R -ness transmits (since, there are no cases in which R -ness transmits). And indeed this is exactly what we find.

Consider the following case of the transmission of non- V -ness (= safety). S 's testimony that p is clean. A knows S well, and so knows in particular that S 's testimony is invariably to be relied upon. A is also a very scrupulous consumer of testimony, remaining alive to signs of unreliability and insincerity in any testimony she observes. Indeed, A 's own dispositions here are based on her powers for reliably distinguishing reliable from unreliable (or otherwise incredible) testimony. Given, what she knows about S , and observing nothing in S 's present testimony that p , which leads her, A , to be suspicious of that testimony, A comes to accept S 's testimony that p . A natural description of this case is one in which the cleanliness of S 's testimony along the V -dimension transmits to the belief A formed through her acceptance of that testimony. Now, since A 's acceptance was itself informed by what she knows of S , and by what she discerned in S 's testifying on this occasion, we can say that A 's testimonial belief is not only non- V (= safe), it is also non- R (and so, by Pritchard's lights, is internalistically justified). At the same time, whereas the non- V -ness of A 's testimonial belief can be traced to the non- V -ness of the testimony she consumed (together with her own reliable capacity to discriminate safe from unsafe testimony), the non- R -ness of A 's testimonial belief should be traced to features of A 's reception of S 's testimony, rather than to features of S 's testimony itself. That is, A 's testimonial belief is non- R , not because of anything having to do with the non- R -ness of the testimony she consumed,

but rather because of *A*'s own careful scrutiny of incoming testimony. As this result does not depend on any features peculiar to this example, we can generalize from this case: we can say that non-*V*-ness is transmissible whereas non-*R*-ness is not.

(As an aside, it may be worth noting that this last point about the non-transmissibility of non-*R*-ness, together with two auxiliary doctrines Pritchard himself endorses in his (2005), have an important implication for his conception of testimonial justification. The two auxiliary doctrines Pritchard endorses in his (2005) are, first, that the elimination of reflective luck is tantamount to the possession of an internalist justification, and second, that internalist justification is the only notion of epistemic justification there is.⁴ Given these two auxiliary doctrines, the fact that non-*R*-ness is non-transmissible in testimony cases would suffice to show that *epistemic justification* is non-transmissible in testimony cases. Rather, what transmits is what others call 'externalist' justification or warrant—something that, so far as Pritchard (2005) is concerned, is not a type of epistemic justification at all. In any case, Pritchard is committed to the non-transmissibility of epistemic justification in testimony cases. Of course, if pressed regarding what *does* get transmitted in successful testimony cases, Pritchard could always say that *knowledge* does.⁵)

So far, the emerging picture of clean testimony seems to be this. Along the *V*-dimension, testimonial belief can be as safe (hence as free from *V*-ness) as the testimony on which it is based; but along the *R*-dimension, there need not be any interesting connection between the *R*-ness (or not) of the testimony, and the *R*-ness (or not) of the belief based on that testimony. This leads us to ask: does a hearer have any epistemic motive for preferring testimony that is both non-*R* and non-*V*, to testimony that is merely non-*V*?

One possible motive for preferring testimony that is both non-*R* and non-*V*, to testimony that is merely non-*V*, is as follows. Since the avoidance of reflective luck is a matter of being able to rule out, on the basis of what one can tell by reflection alone, all of what one takes to be nearby alternatives, and since what one can tell by reflection alone consists of a set of truths ('tell', like 'know', is factive), the avoidance of reflective luck will be strongly correlated with the avoidance of veritic luck. (The correlation will not be perfect, however: this is ensured in Gettier cases.) But given such a strong correlation, it would seem that hearers have a generic motive for wanting testimony that is non-*R*, insofar as they want testimony that is non-*V*. Now if a hearer somehow had easier access to the non-*R*-ness of clean testimony, than she had to the non-*V*-ness of clean testimony, then she would have an epistemic motive for preferring testimony that is both non-*R* and non-*V*, to testimony that is merely non-*V*. Her ultimate desire would be for testimony that is non-*V*, but her access to the non-*V*-ness of a piece of testimony would sometimes (typically? always?) be through her access to the non-*R*-ness of testimony. Now, I am not certain whether hearers do have easier access to the non-*R*-ness of clean testimony, than they do to the non-*V*-ness of clean testimony. Perhaps we can imagine some situations in which this is the case: a hearer might have little background information regarding a piece of testimony, beyond the

⁴ This latter doctrine is never explicitly stated in Pritchard (2005). My basis for ascribing it to the position developed in Pritchard (2005) is that (1) Pritchard's usage of 'justification' throughout the book is invariably qualified as 'internalist justification' (or sometimes this qualification is understood), and (2) that (when I asked him about this) Pritchard himself told me that he was assuming that justification *just is* internalist justification.

⁵ Does this then reflect a knowledge-first orientation (*à la* Williamson) in Pritchard's anti-luck epistemology?

information that the testifier is appropriately reflective in her acquisition of belief.⁶ In that case, the hearer's acceptance of the speaker's testimony might be grounded in the first instance on the (hearer's representation of the) speaker's reflectiveness, taken to be a likely indication of the safety of the testimony. In such cases, the hearer has a motive for wanting testimony that is both non-*R* and non-*V*, to testimony which is simply non-*V*. On the other hand, if there are no cases in which the hearer has easier access to the non-*R*-ness of a piece of testimony as opposed to its non-*V*-ness, then I would speculate that there is no epistemic motive for preferring testimony that is clean along both dimensions (i.e., non-*R* and non-*V*) to testimony that is clean merely along the *V*-dimension.

So far my claims have been about the modal properties of *V* and *R*: both *V*-ness and non-*V*-ness are transmissible (= capable of being transmitted) while neither *R*-ness nor non-*R*-ness is transmissible. But what, precisely, are the conditions under which *V*-ness or non-*V*-ness is in fact transmitted? I begin with non-*V*-ness, or safety. It is easy to see that not all cases in which a hearer consumes testimony that is safe (= non-*V*) will result in a testimonial belief that is safe. To take just one example,⁷ take the case of Polyanna. Polyanna, who is credulous to the point of gullibility, is in a room rife with expert liars who live for gulling the credulous. Fortunately, she happens upon Reliable Ronnie, the one person in the room whose testimony is reliable. Credulous as ever, Polyanna accepts Ronnie's testimony that *p*. And in fact Ronnie's testimony was indeed reliable (safe; non-*V*). Even so, Polyanna's belief that *p* is not safe (i.e., is not non-*V*): it could easily have come about that she acquired the testimonial belief that *p*, even though it was the case that $\sim p$. She would have acquired this belief had anyone attested that *p*, and one of the many liars in the room would have so testified even if it were not the case that *p*. If *p* is a contingent proposition that could easily have been false, Polyanna's testimonial belief that *p* is unsafe.

Indeed, it would seem that safety (non-*V*-ness) transmits only in those cases in which one's consumption of safe testimony is *itself* safe. Here an intriguing thought arises: we might model this phenomenon by regarding the acquisition of knowledge through accepting another's knowledgeable testimony as an instance of knowledge *iteration*—a case of *interpersonally* iterated knowledge.⁸ If this is the proper way to think about knowledge transmission in testimony, then we can cite something like Williamson's margin-for-error principle to explain why the acquisition of knowledge through testimony requires that the hearer be a *safe* consumer of (safe) testimony. In its 'crude' version, this principle states that "iterating knowledge is hard, and each iteration adds a layer of difficulty" (Williamson, 2000, p. 122). This can be applied to the case of testimony, conceived as (interpersonally) iterated knowledge, as follows. There are various possible scenarios in which the belief you report in your testimony would be false. As we saw above, assuming I would have accepted your testimony in those scenarios each of the scenarios are also possible scenarios in which

⁶ Of course, the hearer would also have background information that would let her form an all-things-considered plausibility judgement regarding the proposition attested to; but, we can imagine a case in which the hearer knows very little about the domain in question, and so has no real sense of the all-things-considered-plausibility of the attested proposition.

⁷ This example is taken from Goldberg and Henderson (*forthcoming*). But examples of this sort are familiar; see also Burge (1993, pp. 485–486); Plantinga (1993, p. 86); Faulkner (2000, pp. 591–592); the 'newspaper case' in Graham (2000a, p. 134); and Lackey 2003 and (*forthcoming*).

⁸ The *intrapersonal* case is perhaps the more familiar one: it is the case in which a single subject acquires the knowledge that she has some lower-order piece of knowledge.

my belief, formed through accepting your testimony, would be false. But in addition to these possible scenarios in which my belief based on your testimony would be false, we must add others: perhaps I misinterpreted your remark; perhaps I would have accepted testimony to this effect (whether proffered by you or by someone else) even if it had been false; and so forth. These introduce further error-possibilities—that is, error-possibilities that are relevant to my belief given that it is testimonially-based, but which are not relevant to your belief given that it was not testimonially-based. It is with these additional error-possibilities in mind that we speak of the need for the *safe consumption* of safe testimony. Any lack of safety in the consumption of safe testimony will itself *generate* veritic luck in a testimonially-based belief. (I will return to this in Sect. 3, when I pick up the topic of luck generation in testimony cases.)

What, then, can be said about the transmission of *V*-ness? We cannot think of this directly on the model of the transmission of safety, or non-*V*-ness. For whereas safety transmits only under conditions involving the safe consumption of safe testimony, the perilousness (or *V*-ness) of an unsafe piece of testimony transmits even if the hearer is a scrupulous consumer of testimony who reliably (but not invariably!) distinguishes safe from unsafe testimony, and even if the dispositions that underlie her scrupulous consumption of testimony were operative on the present occasion. Rather, *V*-ness transmits so long as the hearer's testimonial grounds, $G(A)$, fail to eliminate a relevant alternative in $RA(A)$. But since (as, we saw above) $RA(S)$ (the set of alternatives relevant to S 's testimony that p) is a subset of $RA(A)$ (the set of alternatives relevant to A 's testimony-based belief that p), and since in typical testimony cases $G(A)$ cannot eliminate any alternatives in $RA(S)$, which are not already eliminated by $G(S)$ (the speaker's grounds), the result is that $G(A)$ will fail to eliminate a relevant alternative in $RA(S)$ in most every case in which $G(S)$ fails to eliminate a relevant alternative in $RA(S)$. The point, once again, is that *V*-ness (unlike safety, i.e., non-*V*-ness) is not only transmissible, it is *positively contagious*.

3

I now want to consider generation cases, that is, cases, in which there is a change along the *V*- or *R*-dimension from the testimony to the testimony-based belief. We can easily see how it can come to pass that there can be a change for the better, or for the worse, along the *R*-dimension: these are cases, respectively, in which a reflective consumer of testimony acquires a testimonial belief, e.g., about a chick's sex from an unreflective chicken-sexer, or in which an unreflective consumer of testimony acquires a testimonial belief, e.g., about a chick's sex from a reflective chicken-sexer.⁹ In addition, we have already seen how it might come to pass that there can be a change for the worse along the *V*-dimension: these are cases, where, e.g., safe testimony gives rise to an unsafe testimonial belief (as illustrated by the case of Polyanna and the liar-filled room). In these cases *V*-ness itself is *generated* in the process whereby the

⁹ In the latter case, variants are imaginable, in which the testimonial belief is nevertheless safe, hence non-*V*, even as it is *R*. Just imagine that everyone in the hearer's community is endowed with hyper-reliable chicken-sexing abilities, that they never joke around about the sex of chicks, and that it would be near-impossible for the hearer to interact with anyone from outside her community. In that case, it could not easily have come to pass that the hearer acquired a false testimonial belief in a chick's sex; her belief is safe, even if reflectively lucky. See Goldberg (*forthcoming*) for a suggestion about how this point might help us think about the consumption of testimony by very young children.

hearer consumes testimony that was non- V . The hard case to theorize about is the possibility in which there is a change for the better along the V -dimension: a case in which a safe testimonial belief is acquired through testimony that is itself veritically lucky.

At first blush, the account developed so far would appear to tell against such a possibility. For consider the following natural generalization of the account presented so far: whereas a hearer's testimonial belief can be non- R even in cases in which the testimony on which it is based is R , and so in this sense the testimonial belief can *improve* upon the epistemic status of the testimony, even so the same does not hold along the V -dimension, where a testimonial belief cannot improve upon the safety (non- V -ness) of the testimony on which it is based. This natural generalization reflects an intuitive picture of what happens along the V -dimension in testimony cases. The intuitive picture is this. A hearer's fundamental epistemic aim when confronted with proffered testimony is to accept all and only those pieces of testimony that are safe (non- V). If she is perfectly reliable in discriminating testimony that is V from testimony that is non- V , she puts herself in a position to acquire beliefs whose non- V -ness is simply inherited from the non- V -ness of the testimony she is consuming: in that case, her beliefs will be exactly as safe as the testimony she has consumed. Of course, to the extent that she is not reliable in discriminating V from non- V testimony, there is a further source of potential V -ness in her own testimonial beliefs, and so her testimonial beliefs will not generally be as safe as the testimony she consumes. But in no case can the hearer's testimonial belief be *better off* along the V -dimension than the testimony she consumes.

But I want to argue that there are counterexamples to the thesis that:

- (H1) A testimonial belief cannot be better along the V -dimension than the testimony on which it is based.

Consider the following case, which arguably is a case in which a hearer acquires a testimonial belief that is safer than was the testimony she consumed. Ralph is a rabid fan of the New York Yankees. He religiously reads the Sports section of the various New York papers every day of the week, he follows internet trivia about the Yankees, and so forth. But Ralph is also given to wishful thinking, and though he acquires a good many of his beliefs about the Yankees from reliable sources (newspapers, official Yankees web sites, etc.), he also acquires a good many Yankee-beliefs through wishful thinking. Indeed, when the reliable sources give him information that is too difficult for him (i.e., reports of a Yankees' loss), he forgets what he read, and his wishful thinking takes over, providing him not only with the belief that they won, but also with the belief that he read so in this morning's paper! The result is that, taken as a class, Ralph's beliefs about the New York Yankees are an epistemically motley bunch: many are true and safe, but many are false (or sometimes true but V). This has an obvious affect on the testimony he offers: he offers about as much false testimony about the Yankees as true testimony about the Yankees. But there is something interesting about his testifying. Whenever he asserts beliefs originally formed through wishful thinking, his assertions are reliably associated with characteristics, which (in other speakers) are highly correlated with false (or at least unreliable) testimony: there is a certain lack of confidence exhibited in his speech on those occasions, he can't look you in the eyes, etc. But it is also true that his assertions expressing beliefs originally formed through reliable means are reliably associated with characteristics, which (in other speakers) are highly correlated with true (and reliable) testimony: on

these occasions he expresses himself with a good deal of confidence, he can look you in the eyes, and so forth. Ralph himself is oblivious to all of this. What is more, for any of his current beliefs regarding the Yankees, Ralph himself is unreliable in discriminating those formed in reliable ways (i.e., from his scrutiny of the Sports section from the New York papers) from those formed in unreliable ways (i.e., through his wishful thinking). From his perspective all of his beliefs about the Yankees are associated with the same characteristic phenomenology of memory. But those who observe his speech typically can make the discrimination: they are counterfactually sensitive to the tell-tale signs in his testimony.

Suppose that Mary can and does discriminate the relevant features of Ralph's Yankees-testimony, and that she correctly takes them for what they are, namely, indications of the (likely) truth or falsity of his testimony. Accordingly, Mary accepts Ralph's Yankees-testimony when it has the features indicating that it came from a reliable source (such as the morning's paper), but does not accept his testimony when it has the features indicating that it came from an unreliable source (wishful thinking). Since there is a strong correlation between the presence of these features and the truth-value of Ralph's testimony, the result is that most (if not all) of Mary's testimonial beliefs (formed through her acceptance of Ralph's Yankees-testimony) are safe. Yet on at least many (if not all) of these occasions Ralph's testimony that p is not safe: it could easily have come to pass that he attested that p , in a situation in which the attested proposition was false (although believed by him to be true via wishful thinking). Of course, in one of the situations in which Ralph's testimony was formed through wishful thinking (rather than having read the morning paper), Ralph's testimony would have had the tell-tale signs of unreliability, and hence would not have taken Mary in. I submit that this is a case in which *the hearer's testimony-based belief is safer, and so less susceptible to veritic luck, than is the testimony she herself consumed.*

The case for the foregoing depends on the distinction between the epistemic status of Ralph's *testimony* that the Yankees won last night, from the epistemic status of Ralph's *belief* that the Yankees won last night. My claim is that Ralph's *testimony* is not safe; I make no claims about whether *the belief* Ralph expresses in the testimony is safe. This is important, since the point I am making still stands even if it is granted that, in those cases in which Ralph's belief (e.g., that the Yankees won the previous evening) is formed through his reliance on the morning Sports section, he counts as safely believing—stronger, *knowing*—e.g., that the Yankees won the previous evening. Even so, the point would remain that his *testimony* was not safe: easily it could have come to pass that he testified that the Yankees won the previous evening, under conditions in which the Yankees did not win the previous evening.¹⁰ And it is the safety of the testimony, rather than that of the belief manifested by the testimony, that is relevant to the epistemic status of the hearer's belief formed through accepting the testimony. [For an extended argument to this effect, (see Lackey *forthcoming*).]

The case of Mary and Ralph, then, is a case in which Mary's testimonial belief improves along the V -dimension over that of the testimony she has consumed. Admittedly, such cases are peculiar, and (if actual) will be rare. What is more, it is somewhat strange to think that, as consumers of testimony, we can be like the naive (unreflective) chicken sexer, making reliable discriminations without having reflective access to the grounds on which we make the relevant discriminations. (It seems that, when it comes to testimonial belief, cognitively mature agents are not so blind: we have

¹⁰ Here I disregard the fact that the Yankees, a great ballclub, win much more often than they lose.

access to some of the grounds on which we make the discrimination.) But it is also strange to think that there are cases in which the safe consumption of testimony yields beliefs that can eliminate error-possibilities not eliminated by the testimony's grounds themselves. Strange or not, though, such cases, I submit, are possible. And if nothing else, their existence suggests the possibility (documented elsewhere) of knowledge that is *generated* (as opposed to transmitted) by testimony.¹¹

The existence of such cases is important in that it ought to force a modification in our application of the (interpersonal) knowledge-iteration model to the case of testimonial knowledge. I suggested above that the aim of the hearer is to safely discriminate safe from unsafe testimony, that is, that her task is the safe consumption of safe testimony. We now see that this is not quite right. Her task is actually the safe consumption of testimony, *period*—by which I mean that her aim is to consume testimony in such a way as to acquire a safe (non-*V*) testimonial belief. In most cases, one consumes testimony safely only if the testimony one consumed is itself safe. But, we have just seen that there can be cases in which one consumes testimony safely even when the testimony one consumes is not itself safe. So the interpersonal knowledge-iteration model is not quite right: one can acquire testimonial knowledge *other than* via iteration of knowledge had 'at the first level', so to speak.

This argument may not convince. Those who wish to preserve the interpersonal knowledge-iteration model of testimonial knowledge, in the face of the alleged example of testimonial knowledge through unsafe testimony, might do so by returning to a distinction introduced above, between the epistemic status of the speaker's testimony and that of her underlying belief. The claim might be made, for example, that Ralph's testimony is unsafe, but his belief was not (since, it was safely formed and sustained). If, we are entitled to relativize our notion of safety to methods of belief-formation (for which see Pritchard, 2005, p. 135), we might well make some headway with this proposal. As noted above, given a case in which Ralph acquires a belief through the testimony of the local papers, the fact that he might have acquired that belief in some other unsafe way—for example, through wishful thinking—is (given the relativization of safety to methods) simply irrelevant to the question whether his belief is safe. On such a view, Ralph does have the piece of knowledge he purports to express in his testimony, it's just that this knowledge is not safely expressed in the testimony (in that he would have attested to the truth of the proposition even in nearby worlds in which it is false).

Unfortunately, it is not altogether clear whether in such a case it is permissible to describe Ralph as knowing whereof he speaks. At the very least the relevant alternative account of knowledge would have us conclude otherwise. Since Ralph's wishful thinking is often elicited in connection with his forming beliefs regarding the doings of the Yankees, it would seem that on any particular occasion where he forms a belief regarding the Yankees' doings, he counts as knowing that, e.g., the Yankees won, only if he can rule out that this belief was formed by wishful thinking. The trouble is that Ralph can't rule it out in such cases. In cases of wishful thinking, his memorial sustenance of the belief is such that, on recalling the belief, he (wrongly) thinks that the belief in question was acquired through the testimony of the morning paper. Given a case in which his belief was actually formed via wishful thinking, he fails to know this. Worse, given a case in which his belief was actually formed via his reliance on a trustworthy source (the Sports section of this morning's paper), Ralph fails to be able

¹¹ See, e.g., Lackey (1999), Graham (2000), and Goldberg (2005).

to discriminate this case, from the case in which it was formed via wishful thinking: the characteristic phenomenology of memory is the same in both, as is his belief that he acquired his first-order belief about the Yankees' victory from the morning paper. The trouble is that, given the prevalence of wishful thinking in Ralph's psychological economy, the hypothesis that a particular Yankees-belief of his was formed by wishful thinking remains a relevant alternative. Since he can't rule this out even when the belief was formed through reliance on the morning papers, he fails to know that the Yankees won the previous evening.¹²

It would seem, then, that we have reason to think that, even in those cases Ralph acquires the belief (e.g., that the Yankees won last night) in a safe way, such as via the testimony of the paper, still, when he recalls this belief, he does not count as knowing that the Yankees won last night, for failing to be able to rule out a relevant alternative (e.g., that the belief was formed through wishful thinking). Interestingly, Mary can rule out this alternative, as she is sensitive to features in Ralph's testimony, which, she knows, correlate strongly with those cases where Ralph has *not* exhibited wishful thinking. In this way, we can see how it can come to pass that, even though Ralph fails to know that the Yankees won last night (as he testifies to this effect), even so *Mary* comes to know that the Yankees won last night through her consumption of that testimony. The account of how this can come to pass suggests that, while the interpersonal knowledge-iteration model of testimonial knowledge may be unacceptable, still there is an important parallel between knowledge iteration on the one hand, and the acquisition of testimonial knowledge on the other: both involve an increase in the margin-for-error relative to the first-order belief.

The point can be developed as follows. Knowledge requires the elimination of a certain margin for error. By present hypothesis, Ralph's belief about the Yankees fails to have the margin-for-error required for it to count as knowledge: Ralph does not know. Now when Mary accepts Ralph's testimony, the result is that she acquires a belief whose margin-for-error (reflecting the conditions on her counting as knowing) must be *wider* than the margin-for-error pertaining to Ralph's belief (reflecting the conditions on his counting as knowing). In other words, if she is to count as knowing, Mary must be such that the grounds of her testimonial belief eliminate more error-possibilities than those that Ralph's grounds would have had to eliminate (in order for him to have counted as knowing). And what we find is that, indeed, her grounds do cover the required margin-for-error: she does know. In particular, she *can* rule out what Ralph *can't* rule out, namely, the hypothesis that Ralph's testimony is based on wishful thinking. She can rule this hypothesis out, since she has observed those features of Ralph's testimony—the lack of confidence he has, the nervousness

¹² It is interesting to ask how this result, obtained via appeal to the relevant alternative account of knowledge, squares with Pritchard's own anti-luck epistemology. If safety is method-relative, and if it is granted that the methods employed in the two cases (relying on the morning newspapers; wishful thinking) differs, then (assuming that the morning papers can give knowledge) we reach the conclusion that Ralph does know in the case where he relies on the morning newspaper—and this, despite the fact that his memory systematically deceives him in the substantial number of cases in which he forms his belief through wishful thinking. This strikes me as an unhappy result. Given (1) the prevalence of wishful thinking in Ralph's psychological economy and (2) the hyper-unreliability of memory in those cases, in which wishful thinking was operative in belief-formation, I want to say that, insofar as Ralph relies on his memory, he does not know that the Yankees won the previous evening even when he formed this belief via trusting the morning papers. I hope to return to this topic at some future time.

he exhibits, the characteristic failure to look one in the eyes—that obtain *when and only when* he is expressed a belief formed through wishful thinking.

In sum, Testimonial knowledge should not be thought of on the model of knowledge-iteration. Although testimonial links generate ever-increasing margins-for-error (ever more *V*-generating possibilities), a hearer can acquire testimonial knowledge through testimony which itself is not knowledgeable. This marks an important disanalogy between testimonial knowledge and knowledge-iteration: the latter but not the former requires knowledge at the ground level (so to speak). The latter requires it since knowledge-iteration involves the acquisition of knowledge that one knows that *p*, and if one does not know that *p*, then the iterated knowledge cannot be attained. Acquiring testimonial knowledge that *p* is not the interpersonal version of this, however, since testimonial knowledge can be acquired by a hearer even when the speaker does not have the knowledge she seeks to communicate—with the trivial result that the hearer need not know that the speaker knows, in order to acquire testimonial knowledge through her speech. At the same time, though, testimonial knowledge does exhibit one of the features typical of knowledge iteration, namely, an increased margin for error as one ascends from the lower- to the next-higher-order. This is the core truth in the interpersonal knowledge-iteration model of testimonial knowledge.

4

I submit that there are three main lessons to be learned by our application of Pritchard's anti-luck epistemology to the case of testimony. I conclude, first, that, whereas *V*-ness is transmissible, *R*-ness is not; second, that the transmissibility of *V*-ness depends on the safe consumption of testimony, something which typically *though not invariably* occurs in cases involving testimony which itself is safe; and third (and relatedly), that the interpersonal knowledge-iteration model of testimonial knowledge should be rejected, although we should preserve the core idea of an increase in the margin-for-error of testimonial belief over that associated with the belief expressed in the testimony. It is a mark of the fruitfulness of Pritchard's anti-luck epistemology, that it can be used to further our efforts to attain an adequate understanding of the epistemology of testimony.¹³

References

- Burge, T. (1993). Content preservation. *Philosophical Review*, 102, 457–488.
- Faulkner, P. (2000). The social character of testimonial knowledge. *Journal of Philosophy*, 97, 581–601.
- Fricker, E. (Forthcoming). Testimony and epistemic autonomy. In J. Lackey, & E. Sosa (Eds.), *The epistemology of testimony*. Oxford: Oxford University Press
- Goldberg, S. (2005). Testimonial knowledge from unsafe testimony. *Analysis*, 65, 302–311.
- Goldberg, S. (Forthcoming). Testimonial knowledge in early childhood, revisited. *Philosophy and Phenomenological Research*.
- Goldberg, S., & Henderson, D. (Forthcoming). Monitoring and anti-reductionism in the epistemology of testimony. *Philosophy and Phenomenological Research*.
- Graham, P. (2000). Transferring knowledge. *Noûs*, 34, 131–152.
- Lackey, J. (1999). Testimonial knowledge and transmission. *The Philosophical Quarterly*, 49, 471–490.

¹³ Thanks to all the participants in the Author-Meets-Critic session for Pritchard's book *Epistemic Luck* at the Society for Skeptical Studies meeting of the 2005 Pacific APA, and to Duncan Pritchard and Tim Black for very helpful comments on an earlier draft of this paper.

- Lackey, J. (2003). A minimal expression of non-reductionism in the epistemology of testimony. *Noûs*, 37, 706–723.
- Lackey, J. (Forthcoming). Knowing from words. *Philosophy and Phenomenological Research*.
- Plantinga, A. (1993). *Warrant and proper function*. Oxford: Oxford University Press.
- Pritchard, D. (2005). *Epistemic luck*. Oxford: Oxford University Press.
- Sutton, J. (2005). Earn their trust, *manuscript*, <http://faculty.smu.edu/jsutton/earnit.pdf>
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.