

Telling as Inviting to Trust*

EDWARD S. HINCHMAN

University of Wisconsin-Milwaukee

How can I give you a reason to believe what I tell you? I can influence the evidence available to you. Or I can simply invite your trust. These two ways of giving reasons work very differently. When a speaker tells her hearer that *p*, I argue, she intends that he gain access to a *prima facie* reason to believe that *p* that derives not from evidence but from his mere understanding of her act. Unlike mere assertions, acts of telling give reasons directly. They give reasons by inviting the hearer's trust.

This yields a novel form of anti-reductionism in the epistemology of testimony. The status of testimony as a *sui generis* source of epistemic warrant is entailed by the nature of the act of telling. We can discover the nature of this illocution, and its epistemic role, by examining how it functions in the real world of human relations.

In "Other Minds" J. L. Austin compares testifying with promising:

If someone has promised me to do A, then I am entitled to rely on it, and can myself make promises on the strength of it: and so, where someone has said to me 'I know,' I am entitled to say *I know* too, at second hand. The right to say 'I know' is transmissible, in the sort of way that other authority is transmissible. Hence, if I say it lightly, I may be *responsible* for getting you into trouble.¹

It is a characteristically Austinian remark. As a description of what we ordinarily take for granted, it's beyond question. If I'm not careful in what I tell you I know, you *will* hold me responsible if taking my word brings trouble. Yet as a description of what we should say upon reflection, the remark couldn't seem more complacently wrongheaded.

* This paper had an extremely long gestation and benefited from commentary at several stages. In the mid-1990s—when the philosophical literature on testimony was much smaller than it is now!—conversations with Krista Lawlor and Ian Rumfit helped keep proto-drafts evolving. A few years later, when it was a chapter of my dissertation (University of Michigan, 2000), Steve Darwall, Steven Everson, David Hills, Jim Joyce, Peter Railton, Nishi Shah, David Velleman and Andrea Westlund gave helpful correctives. An audience at the Central APA in 2001 (with David Eng commenting), along with some painfully frank anonymous referees, helped me at last turn the chapter into a paper. Thanks to all. (A companion paper on practical reason, 'Advising as Inviting to Trust,' is forthcoming in *Canadian Journal of Philosophy*; and a companion paper on doxastic judgment, 'Judging as Inviting Self-Trust,' is in preparation.)

¹ J. L. Austin, 'Other Minds' (1946), in his *Philosophical Papers*, third edition (Oxford: Oxford University Press, 1979), 100.

How could my telling you that *I* know something *eo ipso* entitle you to say that *you* know it? How could your entitlement to believe something derive from my entitlement to believe it? You may have no idea what entitles me to believe what I tell you I know. Moreover, the basis of my belief may be utterly unlike your putative testimonial basis. So it's hard to see how my entitlement to believe could simply become your entitlement to believe. Perhaps you derive a distinctive entitlement of your own from my knowledge claim. Yet this entitlement must, it seems, also include such a thesis as that when I claim to know that *p*, or when someone claims to know that *p* in this context, it is reliably the case that *p*. Your entitlement to that thesis, should you have it, must not itself depend on my testimony, or the same question of entitlement will arise and take you in a circle.² But if you take my word on the basis of accepting a thesis for which you don't simply take my word, it would seem I'm off the hook:³ the sense in which I may get you into trouble is merely causal, not epistemic. Why should I take responsibility for your belief that I am reliable?

In this paper I will defend Austin against the force of these reflections. I'll focus not specifically on knowledge claims but more generally on acts of telling, whether or not the speaker claims to know what she tells her hearer.⁴ Austin is right that a speaker can be responsible, not merely causally but epistemically, for getting her hearer into trouble. My thesis is that a speaker takes on this burden by addressing her hearer in the way distinctive of telling. I argue by articulating a key distinction between the act of telling someone that *p* and the act of merely asserting that *p* in his presence.

How can I entitle you to believe what I tell you? One way is by influencing the evidence available to you, perhaps by making an assertion or otherwise manifesting a belief, which still makes you epistemically responsible for the belief I want you to form.⁵ Another is by inviting you to trust me,

² Set aside for now the likelihood that your entitlement to the thesis will rest on the testimony of others.

³ Unless, of course, the thesis in question is Austin's thesis. But this paragraph is presenting a series of objections to that thesis.

⁴ Some philosophers hold that any assertion is an implicit claim to know. If this is right, then my argument will entail the truth of that Austin actually said. (For this view, see Peter Unger, *Ignorance* (Oxford: Oxford University Press, 1975), 250-65; and Michael Slote, 'Assertion and Belief,' in Tom Baldwin and Jonathan Dancy (eds.), *Papers on Language and Logic* (Keele: Keele University Library, 1980). For an account of the relation between assertion and knowledge from which this thesis follows, see Timothy Williamson, 'Knowing and Asserting,' *Philosophical Review* 105 (1996), revised as Chapter 11 of his *Knowledge and its Limits* (Oxford: Oxford University Press, 2000).)

⁵ Philosophers have often argued that belief cannot be *forced* in the way that action can (see, for example, Locke's *Letter Concerning Toleration*). For a plausible suggestion why they have been wrong, see Gerald Dworkin, 'What Can We Be Forced to Do?,' *Journal of Social Philosophy* 22 (1991), esp. 45-6. But even if we cannot be forced to believe, we can surely be forced to confront evidence, which is all I have in mind here.

thereby taking part of that responsibility onto my own shoulders. These two ways of giving an epistemic entitlement work very differently. When a speaker tells her hearer that *p*, I'll argue, she acts on an intention to give him an entitlement to believe that *p* that derives not from evidence of the truth of '*p*' but from his mere understanding of the act she thereby performs. The objection two paragraphs back is wrong because, unlike acts of mere assertion, acts of telling give epistemic warrant directly. A hearer may come to be entitled by accepting a speaker's word without accepting any general thesis about when he is so entitled, since a speaker gives her word that *p* by *eo ipso* giving her hearer a *prima facie* entitlement to believe that *p*.⁶ Once we understand the nature of the acts, we cannot even notionally separate the speaker's act of giving her word by telling her hearer that *p* from the act of giving him an entitlement to believe that *p*.

I do not claim that Austin could have foreseen this defense of his position. But I'll argue for it from two starting points that are recognizably Austinian in spirit. I follow the followers of Austin who argue that a speaker performs an illocutionary act just if the appropriate hearer recognizes the intention on which she acts in performing it.⁷ In standard cases of sincere telling, I'll argue, the speaker acts on an intention to make an epistemic entitlement available to her interlocutor. She intends to make this entitlement available, moreover, simply through his recognition that that's what she intends. It is thus possible, I'll argue, for a hearer to come to have an entitlement simply by recognizing his interlocutor's intention to make an entitlement available to him.

I am also Austinian on a methodological point, that what reveals the nature of an illocutionary act is an inquiry into the phenomenon Austin termed 'infelicity.' If an illocutionary act amounts to an interpersonal transaction, we can understand the nature of a given kind of illocutionary act by inquiring into what's at stake for interlocutors as they attempt to bring the transaction off. By looking at the risks they run of infelicity, we come to understand the nature of the joint performance. We cannot understand the nature of telling without understanding how it functions in the real world of human relations.

⁶ I use the terms 'warrant' and 'entitlement' interchangeably. The formulations in this paragraph will be refined.

⁷ This was not quite Austin's own view. For explicit subscriptions to it, see John Searle, *Speech Acts* (Cambridge: Cambridge University Press, 1969), 47; and Jennifer Hornsby, 'Illocution and Its Significance,' in S. L. Tsohatzidis (ed.), *Foundations of Speech Act Theory* (London: Routledge, 1994), 193-4. Both argue that Austin ought to have adopted this view, given his aims.

What's at stake for S and A when S tells A that p? There is, it seems, something more and importantly different at stake for S and A than there would be had A merely heard S tell someone else that p. What's at stake for S when she tells A that p is A's recognition of her status as worthy of his trust.⁸

Contrast two scenarios. First, imagine that when S tells A that p, A does not regard himself as having thereby acquired any entitlement to believe that p. Imagine A manifestly looks as if he needs to learn the time, so S tells him it's noon, but A doesn't regard himself as having thereby acquired any entitlement to believe it's noon. Imagine not that A regards himself as having acquired an entitlement to believe it's noon that gets defeated by such background knowledge as that S's watch tends to run fast, but that A regards himself as not having acquired *any* entitlement, not even a (now defeated) prima facie entitlement, to believe it's noon.⁹ Second, imagine that A hears S tell someone else it's noon, but A does not regard what he hears as giving him any entitlement to believe it's noon. Finally, assume that in each scenario S comes to know that A does not regard S's act of telling as giving him any entitlement.

Now there is, I think, this intuitive difference between the scenarios: in the first A is failing to acknowledge S—he is, as we say, 'slighting' S¹⁰—whereas in the second A is not necessarily slighting S. In the second scenario, after all, S does not tell A it's noon; she doesn't address A at all. We can explain the difference if we view telling as an invitation to trust. In the first scenario S invites A to trust her that it's noon. In the second scenario A merely hears, or overhears, S invite someone else to trust her that it's noon. In the first scenario, if A doesn't regard himself as having an entitlement to believe that it's noon on the basis of S's invitation he quite explicitly (however silently) turns the invitation down: he refuses to trust her. And that is the explanation of S's sense of having been slighted: she has tendered an invitation to A to trust her and explicitly been rebuffed.¹¹

⁸ In my schematic examples, S (the speaker) will always be female and A (the addressee) always male. In the non-schematic example I'll use in section II, I'll redress the gender imbalance by making the speaker ('Stephen') male and the addressee ('Anna') female.

⁹ Of course, if the entitlement is defeated it isn't really an entitlement. I'm imagining that A does not believe that S's telling him the time *would* give him a reason *even if* it were not defeated by (say) his knowledge of the condition of her watch.

¹⁰ By 'slighting' here I do not mean *insulting*. Not all failures of recognition are insulting. Sometimes a failure of recognition is merely disorienting.

¹¹ The idea that telling is inviting to trust figures implicitly in David Simpson, 'Lying, Liars, and Language,' *Philosophy and Phenomenological Research* 52 (1992). Simpson argues that lying involves more than the intention to deceive; it also involves the intention to exploit a trust relation with one's interlocutor and therefore a species of betrayal. But on Simpson's account the reason a liar intends to convey is a reason concerning what the liar intends the interlocutor to believe, against a mutual expectation of the liar's truthfulness, not directly a reason to form that belief. So his account of the trust relation in ques-

We can, of course, imagine a case along these lines in which S would not be entitled to feel slighted: perhaps A doubts not S's word but merely the accuracy of S's watch—in a way that does not reveal any mistrust of S herself. (Perhaps A has seen S unwittingly bang her watch against a lamp-post, and A doubts whether it now runs accurately.) Or perhaps S's belief that it's noon itself derives from the say-so of someone else, S*, and A does not trust S* despite thinking that it was reasonable for S to trust her.¹² We can imagine numerous cases in which S would not be entitled to feel slighted by A's refusal to treat her telling as a source of epistemic warrant. But in each such case we are implicitly assuming that A's refusal to take S's word for it manifests mistrust in something or someone other than S herself. Without that assumption, S is entitled to feel slighted by A's refusal.

Without that assumption, in declining S's invitation A treats S's telling as a telling he merely hears, or overhears, and not as a telling addressed to him. In other words, he treats S's telling as a mere assertion.¹³ He recognizes S's intention to tell him that it's noon and in recognizing that intention satisfies it: S thereby succeeds in telling him that it's noon. But he then fails to live up to his end of the transaction. If he had failed to recognize the intention on which S acts when she tells him that it's noon, S's attempted telling would, in Austin's term, have 'misfired.' But while the act does come off, there is nonetheless an infelicity: Austin terms this sort of infelicity, which presupposes that the act has come off, an 'abuse.'¹⁴ It's an abuse of his status as addressee if he treats S's telling as one he has merely heard or overheard, that is, if he treats it as a mere assertion. Equivalently, it's an abuse of the relation that now exists between them if A fails to regard himself as having, in virtue of S's act, a prima facie entitlement to believe that it's noon.

When S tells an untrusting A that p, we naturally tend to hedge the claim that she told him. Yet when we say, 'She tried to tell him but he wouldn't

tion lacks the epistemic dimension that I emphasize. (Unfortunately I was not aware of Simpson's paper when I wrote the present paper.)

¹² A can think this only if he assumes an internalist view of the reasonability of trust, however. I discuss the difference between internalist and externalist approaches to this issue in sections VI through VIII.

¹³ By 'mere assertion' I don't mean an act that is *nothing but* an assertion; I mean an act that may be a telling but is not a telling addressed to the addressee in question. Here, S's act is a mere assertion from A's perspective.

Speakers seldom merely assert that p. A speaker who asserts that p typically means to tell *someone*—the person whose eye she has caught, the readers of a particular journal, perhaps merely herself. If *all* S is doing is asserting that p, if she does not mean thereby to tell anyone (even herself) that p, then her act does not, I believe, fall into the kind that Austin meant to pick out with the term 'illocutionary.' But to describe an act as an assertion is typically to underdescribe its illocutionary nature.

¹⁴ J. L. Austin, *How to Do Things with Words* (1962), second edition (Cambridge: Harvard University Press, 1975), 16. Austin continues: "do not stress the normal connotations of these names!"

listen,' we don't mean that she tried to tell him but he didn't hear. If A doesn't hear, or doesn't understand what she's trying to do, then S's attempt to tell him misfires: there is no telling. If A does understand what she's trying to do but simply doesn't trust her, S succeeds in telling him in one sense but not in fully implementing or consummating the act.¹⁵ If telling is inviting to trust, the act is performed when S extends her invitation but not consummated until A accepts it. It is a point of Austinian wisdom about speech acts that not all illocutionary unions are consummated.

A consummates S's act by accepting her invitation only when he regards himself as having an entitlement to believe what she told him based in trust. When S tells A that it's noon, she intends that A have an entitlement to believe that it's noon not by virtue of whatever evidence A may possess of her reliability but simply through his recognition of her intention that he thereby acquire an entitlement to believe that it's noon. Since S counts as telling when A recognizes the intention on which she acts in telling, this suggests the following quasi-Gricean analysis:

(T) S tells A that p (sincerely) iff A recognizes that S, in asserting that p, intends A to gain access to a prima facie entitlement to believe that p through this very recognition.¹⁶

¹⁵ These are Austin's terms. When the act does come off but there's an abuse, it is "not implemented or not consummated, rather than void or without effect" (*ibid.*).

¹⁶ More fully: S tells A that p iff A recognizes that S is making an utterance to the effect that p intending thereby to represent herself as intending A to gain access to a prima facie entitlement to believe that p through A's recognition of S's intention that A gain access to a prima facie entitlement to believe that p. That's a mouthful. For simplicity we can take it for granted that S is asserting that p. And we can assume that S is sincere. It also helps to formulate A's recognition as self-referential—or to formulate S's intention as self-referential ("... through recognition of this very intention") or as having a self-referential content ("... through recognition that that's what she intends"), as I do below. For appeals to self-reference in emendation of Grice, see Gilbert Harman, 'Review of Schiffer's *Meaning*,' *Journal of Philosophy*, 71 (1974); Simon Blackburn, *Spreading the Word* (Oxford: Oxford University Press, 1984), 114-118; and Francois Recanati, *Meaning and Force: The Pragmatics of Performative Utterances* (Cambridge: Cambridge University Press, 1987), Ch. 7. For a formulation à la Grice, see section IX below. (I use self-reference merely to ease formulation; I do not consider any philosophical rationale for it.) I call this a *quasi-Gricean* account because on Grice's analysis of speaker's meaning (see the canonical papers, especially 'Meaning' (1957), in Paul Grice, *Studies in the Way of Words* (Cambridge: Harvard University Press, 1989)) the speaker aims to produce a response in the speaker—to get him to believe something, or to do something—not merely to give him an entitlement to believe or do it. Moreover, I do not aim to analyze speaker's meaning but instead the nature of an illocutionary act.

Some Griceans (e.g. P. F. Strawson, 'Meaning and Truth' (1969) in his *Logico-Linguistic Papers* (London: Methuen, 1971)) would view S as intending to get A to believe that she, S, believes that p. (Grice himself accepts the modification of his 1957 view in 'Utterer's Meaning, Sentence-Meaning, and Word-Meaning' (1968), *Studies in the Way of Words*, 123.) But here I sympathize with McDowell's complaint ('Meaning, Communi-

We can now account for the nature of A's abuse when he fails to regard himself as coming to have an entitlement to believe that it's noon—when he fails, in effect, to take S's word for it. A recognizes that S is telling him that it's noon. By (T), A recognizes that S intends that he gain access to an entitlement to believe that it's noon through this very recognition. But then A refuses to acknowledge the entitlement. In telling A that p, S offers him something, an entitlement to believe that p, which she conceives as his for the taking. In recognizing her intention to tell him that p, A satisfies that intention, and S thereby counts as telling him that p. But in refusing to acknowledge the entitlement, he refuses the offer she makes in telling him that p. The telling comes off, but S's aim in telling is thwarted.

II

The quasi-Gricean analysis explains nicely part of what's constitutively at stake for S when S doesn't merely assert that p in A's presence—perhaps telling someone else that p—but tells A that p. If what's at stake for S when she tells A that p is his recognition of her as worthy of his trust, what's at stake for A? Quite simply: his status as reasonable in trusting her. While A can abuse the illocution by failing to regard S as worthy of his trust, S can abuse it by failing to live up to that trust, should A accept her invitation and give it.

If you've told someone that p, that person is now under certain conditions—for example, if he's challenged whether p—entitled to hold you accountable for producing a reason to believe that p. If you can't do so, you are subject to criticism not only as an *asserter*, that is, as one who has represented herself as believing that p yet when pressed cannot produce a reason to believe that p, but as a *teller*, that is, as one who has culpably misled your addressee.¹⁷ Inability or refusal to produce a reason would amount to an abuse of your status as teller. What shows that you do take on this additional burden by telling as opposed to asserting? One thing that shows it is that you can pre-

cation, and Knowledge,' in Z. van Straaten (ed.), *Philosophical Subjects* (Oxford: Oxford University Press, 1980), 127) that "the primary point of making assertions [I would say: in telling people things] is not to instil into others beliefs about one's own beliefs, but to inform others—to let them know—about the subject matter of one's assertions (which need not be, though of course it may be, the asserter's beliefs)." How else can she do this if not by giving him an entitlement to believe what she tells him?

¹⁷ Cf. Robert Brandom, *Making It Explicit* (Cambridge: Harvard University Press, 1994), 171: "In asserting a claim, one not only authorizes further assertions (for oneself and for others) but undertakes a responsibility, for one commits oneself to being able to vindicate the original claim by showing that one is entitled to make it." If 'telling someone that p' is substituted for 'asserting a claim' this is similar to the thesis I'm endorsing here. But the one to whom one is responsible is just the one told, that is, the one entitled to complain that one has misled him if one cannot vindicate the claim. (It is a separate question what counts as vindication sufficient to render complaint unreasonable. In everyday contexts, the standards are not very high. Cf. note 51 below.)

vent your addressee from holding you accountable by asserting what you want to tell him while preventing him from construing your assertion as an effort to tell him that. It is toward this end that you might avail yourself of a species of conversational indirection.

Consider two protagonists, Stephen and Anna, returning from a visit to Anna's parents. Stephen wants to tell Anna that she has been cruel to her mother, but he knows that if he does she'll simply accuse him of accusing her and they'll be off on one of their rows. Stephen is in a not uncommon bind. He wants to convince Anna that she was cruel to her mother. And he knows that she is not likely to infer that she has been cruel from anything short of an explicit declaration from him. But he also knows that if he tells her that she was cruel to her mother, she will resent the act and he will not succeed in convincing her. So he searches for his most unfrontational, his most inner-ruminative tone of voice. "Didn't Martha look sad as we left?" he begins. "Some of those remarks did seem a bit cruel..."

It may not prove easy, but assume he pulls the indirection off. How would this look from Anna's perspective? From her perspective it will look as if Stephen intends to tell *someone* that Anna was cruel to her mother but perhaps not Anna herself. Since there is no one else present, it will look as if Stephen intends to tell at most only himself that Anna was cruel, that is, as if she has overheard an inner reflection of his that he did not intend for her ears.¹⁸ To put it neutrally, it will look to her as if Stephen is merely asserting that Anna was cruel to her mother. But one who merely hears, or overhears, a speaker assert that *p* is not thereby entitled to hold the speaker accountable for providing a reason to believe that *p*.¹⁹

There are thus two distinct kinds of act that Stephen might be performing in acting on an intention to make available to Anna an entitlement to believe she was cruel through her recognition of that intention. On the one hand, he might be telling her that she was cruel. On the other hand, he might be asserting that she was cruel, intending to give her evidence for the proposi-

¹⁸ Of course if there is someone else there the indirection will be in this one respect easier to pull off (though perhaps not in others).

¹⁹ On 'mere' assertion, see again note 13. One might counter that there is a norm governing our practices of (mere) assertion that entitles anyone to challenge anyone who asserts that *p* to produce a reason to believe that *p*. (Perhaps Brandom believes this: see the quotation in note 17 above.) I doubt that there is any such norm, but even if there were it would be a different norm from the one governing the practice of telling that I am describing. If *S* tells *A* that *p* but cannot when challenged produce a reason to believe that *p*, she counts as misleading *A*—an abuse of the illocution of telling. If, on the other hand, *S* merely asserts that *p*, then if she cannot produce a reason to believe that *p* she is not culpable for misleading anyone. If you say she has 'abused' the illocution of asserting, that's a different abuse.

tion that she was cruel.²⁰ In refraining from telling her that she was cruel, he refrains from granting her the special right to challenge him to justify his assertion that he would have granted her had he told her.²¹ By appearing not to address her, he makes available to her the entitlement he intends her to have without undertaking the justificatory burden he would have undertaken had he made the entitlement available by telling her. It may make a real difference to the quality of their conversation whether she views him as undertaking that burden. It may make a real difference whether she views him as offering her anything as intimate as an invitation to trust him.²²

²⁰ What evidence might he thereby give her? Simply the evidence that *he* believes she was cruel. If she has good evidence that his beliefs on this subject tend to be reliable, she acquires an entitlement to believe what he asserts.

The literature on Grice's analysis of speaker's meaning has been vitiated by failure to mark this distinction. See, for example, Strawson's counterexample to the sufficiency of Grice's analysis in section III of 'Intention and Convention in Speech Acts' (1964) in *Logico-Linguistic Papers*. The counterexample satisfies Grice's conditions on S's meaning that p but involves a kind of indirection that is not compatible with S's meaning anything at all by her act. What gives this result is that, as Strawson describes her, S is solely concerned to influence her audience's beliefs through the production of evidence. Strawson doesn't conceive Grice as aiming to characterize inherently communicative acts such as telling. And indeed that was not Grice's aim. But if we take not asserting but telling and other inherently communicative acts as the explananda of a neo-Gricean program—one that simply abandons the aim of analyzing speaker's meaning—we need no longer fear a counterexample such as Strawson's.

²¹ One might doubt that Stephen is even asserting that Anna was cruel in this scenario. Isn't he merely, as we say, wondering aloud whether she was cruel? One might view his strategy as that of getting Anna to believe that she was cruel through a kind of conversational implicature. First, Stephen gets Anna to understand him as wondering aloud whether she was cruel. He then relies on her to infer that he believes that she was cruel from the thought that unless he intends her to draw this inference he violates a conversational maxim, perhaps that requiring 'relevance.' (One might thus extend Grice's logic of conversation from its intended domain to the domain of illocutionary force.) Stephen thereby produces a linguistic manifestation of his belief that Anna was cruel without even having to assert that she was cruel. (We might conceive this as an indirection within an indirection.) No matter. The case, even so conceived, illustrates the contrast I am drawing with the case in which Stephen *tells* Anna that she was cruel.

²² We might say that Stephen *invokes* or even *exploits* his trust relation with Anna without inviting Anna to trust him on this issue. The contrast would then be between the case in which Stephen merely relies on the fact that Anna considers him trustworthy and the case in which he invites her to consider him trustworthy. (I owe this formulation to David Hills.)

We might place here the interesting psychological datum that overhearing a remark can be more epistemically compelling than having the remark addressed to you. Stephen might trade on this psychological fact. Or Anna might fear Stephen's trading on it and call him to accounts. How can we account for the datum itself? One explanation would note that the ground for believing a mere assertion is conceptually simpler than the ground for believing what a speaker tells you: in the former case you need merely draw upon evidence of trustworthiness, whereas in the latter you must acknowledge your place in a trust relation. Related to this is the equally interesting datum that it is often easier to believe that a speaker is lying to *you* than to a complete stranger. We usually imagine our own interpersonal relations to be more difficult—hence more interesting!—than those of

Speakers can, of course, perform such backings-off explicitly. Imagine your interlocutor manifests a belief that the Liar Paradox derives from a passage in St. Paul's *Epistle to Titus*. Contrast the following two backings-off. If alert, you say: (a) "In fact the Liar Paradox was discussed much earlier. But don't take my word for it; look under the entries for Eubulides of Miletus and Theophrastus in Diogenes Laertius's *Lives of the Philosophers*." If groggy, you say: (b) "In fact the Liar Paradox was discussed much earlier. Wait, let me think. No, I'm not sure. Perhaps I'm getting it confused with discussion of one of the Heraclitean paradoxes." If you say (b) you tell your interlocutor that the Liar Paradox was discussed earlier, and then retract the telling. You don't do quite the same if you say (a). Does this conflict with my account? No. The difference between (a) and (b) is not that you retract your telling in the latter but not in the former. You retract your telling in each case. The difference is that in (a) you retract only your act of telling your interlocutor that the Liar Paradox was discussed earlier but not your act of asserting it. In (b) you retract both your telling and your asserting. At the end of (b) you don't count as even asserting that the Liar Paradox was discussed earlier. At the end of (a) you do.

Now one might hold that in (a) you do not retract the telling but reinforce it with your allusion to Diogenes Laertius.²³ How could we resolve that dispute? I concede that you might have said, "The Liar Paradox was discussed much earlier, I tell you, but don't take my word for it..." Does that show that your act must be a telling, even after you instruct your interlocutor not to take your word for it? Hardly. 'I tell you' has roughly the same meaning as 'I'll have you know' (or 'by God'): it's a kind of exclamation. It is revealing that when you say (a) your body language will naturally convey the retraction of a kind of intimacy: you'll hold up your hands as if to push your interlocutor away from you, or as if to show him that you bear no concealed assumptions about your relationship. You do reinforce your act when you go on to refer him to Diogenes Laertius. You reinforce the assertion you just made, having backed off of your act of telling.

I do not, however, insist that it is wrong to call your speech act a telling. I insist merely that it is not a telling in the sense analyzed by (T). Of course, that will disappoint anyone who thought my principal aim was to explain what unifies our usage of the verb 'to tell.' But that is not my aim. My aim is to analyze an important sense of 'tell' and then on the basis of that analysis to explain how such acts create a distinctive species of epistemic warrant. Philosophers have not learned to distinguish telling A that p from asserting

people we do not know. (Thanks to an anonymous referee for reminding me of the first datum.)

²³ This was the thesis put forth by an anonymous referee, who framed the example as a counterexample to my account of telling.

that *p* in *A*'s presence. But I think they thereby overlook an important distinction. It is important that you can take responsibility for your interlocutor's entitlement to believe that *p* onto your own shoulders, rather than stepping back from that intimacy and referring him to others.

The possibility of such backings-off confirms the quasi-Gricean analysis of telling. It reveals that a speaker who tells her addressee that *p* intends that he thereby gain access to an entitlement to believe that *p* that, by contrast, doesn't rest on whatever evidence he may have either of her reliability or of the reliability of her sources. For his part, Stephen wouldn't be in the pickle he's in if an entitlement made available by telling rested on evidence of reliability. In that case, there would be no relevant difference between telling Anna that she was cruel and merely asserting that she was cruel. Whichever he did, her belief that she was cruel, should she form it, would derive from that evidence, and his role in the transaction would be merely to speak his mind. But there is a relevant difference between asserting that Anna was cruel and telling Anna that she was cruel. If he tells her that she was cruel, he offers her access to an entitlement to believe so that rests not on evidence but solely on his intention to make that entitlement available. In telling her, he would intend to make available to her an entitlement to believe that she was cruel *simply* through her recognition that that is what he intends.

III

It is important to distinguish Stephen's predicament from other common predicaments. Stephen believes he must adopt alternative means of convincing Anna that she was cruel because of the consequences that will ensue if she views him as telling her this. We'd have an importantly different case if Stephen's obstacle lay in his belief that Anna will not trust him and so will not provide the uptake required for an act of telling to be consummated. We'd have a third case if Stephen believed that, though Anna will trust him by giving his word due weight in forming a belief about her conduct, she will not come to believe that she has been cruel. In each case, Stephen can adopt the strategy of merely asserting what he wishes he could tell Anna.

Let me draw these distinctions schematically and provide some labels. In sincerely telling *A* that *p*, *S* intends to bring about two distinct things: she intends (a) that an entitlement become available to *A* through the quasi-Gricean mechanism, and she intends (b) that *A* acquire that entitlement by trusting her. We must accordingly distinguish two forms of uptake that *A* might contribute in bringing off the transaction. (i) If *A* recognizes *S*'s intention to make an entitlement available to him, he satisfies her (a)-intention and the entitlement is thereby made available to him. (ii) If *A* also trusts *S*, he satisfies her (b)-intention and thereby comes to acquire the entitlement. It is worth emphasizing that *A* might trust *S* on the question whether *p* without

actually coming to believe that p.²⁴ S will, of course, typically also intend (c) that A actually come to believe that p.

We can help ourselves to Austin's distinction and say that only S's (a)-intention is an intention to perform an *illocutionary* act. Her (c)-intention is an intention to perform a *perlocutionary* act. Her (b)-intention I've described as the intention that her illocutionary act of telling A that p be consummated. When S's illocutionary act comes off, A thereby gains access to an entitlement to believe that p. When A consummates S's act by trusting her, he gains the entitlement. Her perlocutionary act comes off, by contrast, only when he forms the belief that p.

You can tell someone what you're sure you'll prove unable to get him to believe. And you can tell someone certain not to trust you on the matter. You can put the invitation directly. 'You're a fool, take my word for it,' you can say. It sometimes makes good sense to issue an invitation certain to be refused.²⁵ The interlocutor you'll prove unable to tell is the interlocutor who will not recognize the complex intention informing your act. You can try to tell him, of course, but your telling will misfire as long as he fails to recognize your intention thereby to make available an entitlement to take your word.²⁶

²⁴ He might, for example, trust her to tell him the weather accurately but know that the 'window' into which she gazes as she tells him it's sunny is in fact a hyperrealistic painting. Here he comes to acquire a prima facie entitlement to believe that it's sunny that his ancillary knowledge defeats.

²⁵ My account thus avoids the fatal flaw of Grice's accounts of asserting/telling. (Grice didn't distinguish telling from asserting.) On Grice's accounts (see note 16 above), S tells A that p by intending that A believe that p, or—in the later version—that A believe that S believes that p. Grice combined this with the principle that "one cannot have intentions to achieve results which one sees no chance of achieving" ('Utterer's Meaning and Intentions,' 98). This yields the result that if you see no chance of getting your interlocutor to believe that p (or, to believe that you believe that p) then you cannot tell him that p. For some examples that bring out the absurdity of this result, see Ian Rumfitt, 'Truth Conditions and Communication,' *Mind* 104 (1995), 833-34. Even if the principle is too strong (for discussion, see J. David Velleman, *Practical Reflection* (Princeton: Princeton University Press, 1989), Ch. 4), it is intuitively implausible that success in telling should be tied—even weakly—to success in convincing (even of the proposition that you believe that p). The two forms of success seem entirely unrelated.

²⁶ I lack space to give this last claim the full defense it needs, in light especially of its bearing on debate over the idea that a violation of free speech could occur by virtue of an intended audience's inability to recognize the illocutionary intention. (Say, a woman tries to refuse sex, but the man cannot recognize what she's doing as manifesting an intention to get him to stop.) For objections to the claim, see Daniel Jacobson, 'Freedom of Speech Acts? A Response to Langton,' *Philosophy and Public Affairs* 24 (1995), 72-78. For a reply, see Jennifer Hornsby and Rae Langton, 'Free Speech and Illocution,' *Legal Theory* 4 (1998), 28-31; then see Jacobson's rejoinder, 'Speech and Action: Replies to Hornsby and Langton,' *Legal Theory* 7 (2001), 179-201.

IV

How could there be such an entitlement? How could you come to have an entitlement to believe that *p* simply by recognizing someone's intention to give you an entitlement to believe that *p*? How could you acquire this status merely by accepting an invitation to trust?

When I speak of an 'entitlement' (or 'warrant') I do not necessarily mean a proposition playing the role of premise in any process of reasoning in the subject. Some of a subject's epistemic entitlements are internally accessible—at whatever level of cognition—and some may not be. It may thus be possible for a subject to have an entitlement to an inferential transition without having a corresponding *justification*. As I'll use the terms, a 'justification' is something you could cite in order to justify an inferential transition. An entitlement is merely something you rely on, in relying on which you count as warranted. I mean to use the term 'entitlement' very broadly to cover both accessible and inaccessible species of warrant (should there be any of the latter).²⁷

When an addressee satisfies a speaker's intention to tell him that *p*, according to (T), he gains access to a *prima facie* entitlement to believe that *p*. Since his recognition of *S*'s intention to tell him that *p* suffices for his being told that *p*, *A* gains access to an entitlement to believe that *p* directly, not from evidence of the truth of '*p*' but from his mere understanding of what *S* is doing.²⁸ It is hardly a novel claim nowadays that the transmission of epistemic entitlements through testimony is thus direct. Tyler Burge, for example, argues for the following principle:

A person is entitled to accept as true something that is presented as true and that is intelligible to him, unless there are stronger reasons not to do so.²⁹

And Robert Brandom argues for the following account of assertion:

²⁷ My usage of 'warrant,' 'entitlement,' and 'justification' follows that of Tyler Burge 'Computer Proof, Apriori Knowledge, and Other Minds,' *Philosophical Perspectives* 12 (1998), 3. (Note that I'm for present purposes trying not to choose sides in the debate between internalists and externalists. I discuss an externalist interpretation of my view in sections VII and VIII.)

²⁸ Of course, in order to understand what *S* is doing here, on my quasi-Gricean view, *A* has to have a belief about *S*'s intentions. There is thus a sense in which whatever warrant *A* derives from that state of understanding must derive in part from whatever warrant *A* has for that belief. But this would not render the former warrant indirect, in any sense relevant to the current discussion, since, if my quasi-Gricean view is right, the belief about *S*'s intentions *just is* the state of understanding.

In the context of the present argument, I do not claim that the entitlement is apriori, merely that it is direct. I consider the question of apriority in 'Judging as Inviting Self-Trust,' in preparation. See also note 45.

²⁹ Burge 'Content Preservation,' *Philosophical Review* 102 (1993), 467; cf. the fuller statement of this principle in Burge, 'Interlocution, Perception, and Memory,' *Philosophical Studies* 86 (1997), n. 4. For a fuller treatment of the issue, see Burge, 'Computer Proof, Apriori Knowledge, and Other Minds.'

What is *done* in asserting ... consists in the way in which ... asserters alter the score interlocutors keep of the deontic statuses (commitments and entitlements) of their fellow practitioners.³⁰

Both Burge and Brandom thus claim that the mere act of asserting that *p* suffices to give any understanding hearer *prima facie* warrant or entitlement to believe that *p*.³¹ But this is stronger than the thesis I'm defending. I believe it is too strong.

In light of my argument so far, we can see that the stronger thesis is too strong in two respects. First, it is too strong to claim that any speech act of committedly putting it forth that *p*—whether telling or asserting—could (interestingly) *eo ipso* give a hearer an entitlement to believe that *p*.³² Such a speech act can at most *make an entitlement available* to a hearer, thereby causing that hearer not to acquire but to *gain access* to the entitlement. As indicated in section III, if *A* doesn't in fact trust *S* then he fails to acquire an entitlement to believe what she tells him. This is the difference between the uptake necessary and sufficient for an illocutionary act to transpire and the uptake necessary and sufficient for an illocutionary act to be consummated. Otherwise put, it is the difference between acknowledging an invitation to trust and actually entering into a trust relation. My thesis is that telling is inviting your interlocutor to enter into a species of relation that provides him with warrant. I do not claim that telling by itself provides him with warrant.³³

Second, and more fundamentally, these philosophers fail to distinguish within the class of hearers those to whom the speaker has addressed her

³⁰ Brandom *Making It Explicit*, 171; cf. his 'Asserting,' *Noûs* 17 (1983).

³¹ Burge and Brandom take very different routes, however, to these conclusions. For still other routes to similar conclusions, none of which involves distinguishing telling from asserting, see George Myro, 'Aspects of Acceptability,' *Pacific Philosophical Quarterly* 62 (1981); Michael Welbourne, *The Community of Knowledge* (Aberdeen: Aberdeen University Press, 1986) and 'Testimony, Knowledge and Belief,' in Bimal K. Matilal, and Ardinam Chakrabarti (eds.), *Knowing from Words* (Dordrecht: Kluwer Academic Publishers, 1994); C. A. J. Coady, *Testimony* (Oxford: Oxford University Press, 1992) and 'Testimony, Observation and "Autonomous Knowledge,"' in Matilal and Chakrabarti, *op. cit.*; Michael Dummett, 'Testimony and Memory,' in *ibid.*; P. F. Strawson, 'Knowing from Words,' in *ibid.*; John McDowell, 'Meaning, Communication, and Knowledge,' *op. cit.*, and 'Knowledge by Hearsay,' in Matilal and Chakrabarti, *op. cit.* (McDowell focuses on the speech act of letting know, not the same as asserting but not what I mean by telling); and, more recently, Peter J. Graham, 'Transferring Knowledge,' *Noûs* 34 (2000); Richard Foley, *Intellectual Trust in Oneself and Others* (Cambridge: Cambridge University Press, 2001), Ch. 4; and Gary Ebbs, 'Learning from Others,' *Noûs* 36 (2002).

³² The uninteresting case: *S* puts it forth that she is putting something forth, thereby giving any observer an entitlement to believe that she is putting something forth. (By 'committedly' putting it forth, I mean putting it forth with the commitment to truth characteristic of asserting or telling.)

³³ I'll further explain the basis of this claim in section V.

speech act—i.e. the hearers who count as *told* by the speaker.³⁴ As I've argued, only these hearers *eo ipso* gain access to an entitlement to believe what the speaker tells them; other hearers—overhearers—gain access to the same kind of warrant to believe what she asserts as they would if she'd manifested her belief in some other way.³⁵

Call the thesis that we *can* reduce testimony entitlements to more basic entitlements (such as those derived from perception and memory) *reductionism* about testimony.³⁶ The challenge for an anti-reductionist about testimony is not only to argue against reductionism but to *explain how* testimony is a *sui generis* source of epistemic warrant. The anti-reductionist must explain what it is about testimony that *makes* it so. One contribution of the present paper lies in its detailed explanation of how the status of testimony as an irreducible source of warrant is entailed by the nature of the illocutionary act of telling, an act whose nature in turn derives, like that of any illocutionary act, from the nature of the intentions that the speaker gets her addressee to recognize in performing it.

V

I'll now elaborate that explanation. Say S tells A that p. A thereby gains access to an entitlement to believe that p directly, claim we anti-reductionists.

³⁴ See the useful diagram in Herbert H. Clark, *Using Language* (Cambridge: Cambridge University Press, 1996), 14. For my purposes, what Clark calls a 'side participant' we should conceive as potentially an addressee. Side participants are participants in the conversation who are not being addressed at the moment. It seems that in the testimonial case a speaker needn't be distributing his attention evenly among all participants in the conversation in order to count as addressing them all in the way characteristic of telling. Unless there is some reason to think otherwise (e.g. the speaker winks at one participant while addressing another), the only hearers to whom the speaker is *merely* asserting are the ones Clark calls 'bystanders' and 'eavesdroppers'—in other words, those who are not participating in the conversation.

³⁵ For the case of merely hearing a speaker assert that p (without being told that p oneself), I agree with the reductionist account defended in Elizabeth Fricker, 'Against Gullibility,' in Matilal and Chakrabarti, *op cit.*: one's entitlement must derive from sources other than the testimony itself. But Fricker, like the anti-reductionists she criticizes, fails to distinguish asserting from telling.

³⁶ The reductionist tradition in the epistemology of testimony derives from Hume. (See *An Enquiry Concerning Human Understanding*, edited by L. A. Selby-Bigge and P. H. Niddich (Oxford: Oxford University Press, 1975), 113: "The reason why we place any credit in witnesses and historians, is not derived from any *connexion*, which we perceive *a priori*, between testimony and reality, but because we are accustomed to find a conformity between them.") The anti-reductionist tradition, dormant until recently, derives from Reid. (See *An Inquiry into the Human Mind on the Principles of Common Sense*, in *Inquiry and Essays*, edited by R. E. Beanblossom and K. Lehrer (Indianapolis: Hackett, 1983), 95: "It is evident that, in the matter of testimony, the balance of human judgment is by nature inclined to the side of belief; and turns to that side of itself, when there is nothing put into the opposite scale. If this was not so, no proposition that is uttered in discourse would be believed, until it was examined and tried by reason; and most men would be unable to find reasons for believing the thousandth part of what is told them.")

But exactly what entitlement? We can clarify the entitlement by clarifying the analogy between testimony and other basic sources of epistemic warrant such as perception and memory.

Other anti-reductionists claim that A's entitlement to believe that *p* derives from the fact that S asserted that *p*. This fails to distinguish asserting from telling. But say they amended it to: the fact that S told A that *p*. This would still be wrong. If this is the usual case, S doesn't intend that A should be entitled to believe whatever she tells him, merely because she told him. She intends that he gain access to an entitlement to believe that *p* on the basis of her telling him that *p*, but she doesn't intend that his entitlement derive merely from the fact that she told him that *p*. She intends that he be entitled not merely on the basis of being told but also on the basis of trust in her. She doesn't think that he would gain this entitlement should he fail to trust her. Should he fail to trust her, he would at best have whatever entitlement he may gain from the fact that she merely asserted that *p*. She intends that his entitlement to believe that *p* derive not merely from the fact that she spoke but from the fact that she addressed him in such a way as to engender a trust relation—that is, from the conjunctive fact that she told him that *p* and he trusts her.

Of course, all she can do to give him this entitlement is tell him that *p*; whether he trusts her is not up to her. She can attempt to 'inspire' his trust in her, but in the usual case she is not in the business of forcing him to trust her.³⁷ She makes her contribution to his gaining an entitlement to believe that *p* by telling him that *p*. But the entitlement she intends that he thereby come to have derives also from his trust in her, which is a contribution of his own. The entitlement derives, in other words, from the fact that they stand in a trust relation.

It is here that we may draw a useful analogy between testimony and memory. Say A forms the belief that he ate cereal for breakfast this morning by remembering that he did. It is an important point, much emphasized by anti-reductionists, that the epistemic basis of A's belief need not include evidence derived from memory.³⁸ His warrant for believing that he ate cereal derives from the fact that he remembers having done so, as long as his memory is reliable on such a matter. His memory must *be* reliable, but he needn't possess evidence of its reliability. He would need evidence of its reliability if the reliability of his memory figured as a premise from which he reasoned in forming his memory belief. But he needn't reason from any such premise. He needn't reason, 'My memory is reliable, and it tells me that I ate cereal, so I

³⁷ We *can*, I think, brainwash, frighten or flatter interlocutors into trusting us, but to view the usual case in these terms amounts to a skepticism about trust that I'll for present purposes set aside.

³⁸ See Burge, 'Content Preservation,' 462-5; and Dummett, 'Testimony and Memory,' 419.

must have eaten cereal.' In fact he needn't reason from premises at all in this case: he can be warranted in his memory belief simply by trusting his memory—as long as his memory is reliable on such a matter.³⁹ His warrant for believing that he ate cereal can derive from the conjunctive fact that he remembers having done so *and* he trusts his memory. His memory proposes, but he himself determines the warrant-generating self-trust relation. Of course, self-trust will typically be habitual. Even gripped by habit, however, the subject takes responsibility for his memory beliefs by remaining answerable for the reasonability of the self-trust they manifest.

The same point applies to other basic sources of epistemic warrant such as perception and inference. If your eyes or your inferential reasoning were not generally reliable guides to the truth, your trust in them wouldn't be reasonable and so wouldn't provide you with any warrant. But given that they are generally reliable, you are warranted when guided by the disposition to believe what they tell you.

Trust gives rise to epistemic warrant, I claim, but it is importantly not the case that trust does so *on the basis* of epistemic warrant. Trust is not belief, although it may give rise to belief. Trust, we can say generally, is a species of willed dependence, where the dependence is under appropriate guidance of a counterfactual sensitivity to evidence of untrustworthiness in the trusted. 'Appropriate guidance' means: you would not trust if you had evidence that the trusted is not worthy of your trust. Trust need therefore involve no active assessing at all, if the trusted is trustworthy; it is enough merely to be (and, counterfactually, have been) disposed to notice and respond appropriately to evidence of untrustworthiness should there be (or have been) any. Though you may have practical reasons for or against trusting, I don't think the epistemic reasonability of trust derives from your possession of epistemic reasons or warrant *for* the trust.

Better, I think, to put it this way. Trust is a source of epistemic warrant just when it is epistemically reasonable. Trust is epistemically reasonable when the thing trusted is worthy of the trust—as long as there is no evidence available that it is untrustworthy. Assuming satisfaction of this negative evi-

³⁹ His memory may of course be reliable on some matters but not on others. When I speak of the reliability of an epistemic faculty, I'll always mean: its reliability on the subject matter at hand. This does raise the question how broadly to construe the 'subject matter.' Is the subject matter of the present case *what I had for breakfast this morning* or *what I did within the past twenty-four hours* (to mention two possibilities among many)? I don't see how we could settle this apart from the full context of justification, or why it should matter if we cannot. The fact that you are poor at remembering names doesn't impugn the reasonability of trusting your memory that you ate cereal this morning. Your myopia doesn't impugn the credentials of your visual beliefs about what's right under your nose. (In Goldman's terms, I am appealing to more-or-less *local* (rather than *global*) reliability. See Alvin Goldman, *Epistemology and Cognition* (Cambridge: Harvard University Press, 1986), 44-45.)

dential condition (which I'll discuss in the next section), when an epistemic faculty is trustworthy by serving as a reliable guide to the truth, it makes available an entitlement to believe what it tells you whose basis lies simply in the fact that you trust it.

Memory, perception and inferential reason do not of course tell us things in the sense of 'tell' that I'm analyzing in this paper. Our epistemic faculties do not perform illocutionary acts; they merely deliver to us propositions which we then believe or not on the basis of trust. I wish to argue not that these faculties are like speakers in telling us things but that speakers' acts of telling us things are like these faculties in generating entitlements for us to believe on the basis of trust. A speaker's telling you that *p* will often deliver to you a proposition which your trust in her disposes you to believe. Of course, the mechanism of delivery is not the same as in the other cases: there it was simply a process, here she *invites* your trust.

Such acts of telling are nonetheless a source of warrant on a par with the deliverances of those other epistemic faculties. You can have an entitlement to believe that *p* derived directly from being told that *p* in a context of reasonable trust. This is true regardless of the kind of 'teller'—whether memory, perception, inference, or other people.⁴⁰

VI

This approach preserves what is worth preserving in reductionist and anti-reductionist accounts. The reductionist is right that it is reasonable to trust only those who are worthy of your trust. And the anti-reductionist is right that reasonable trust does not rest on an assessment of your interlocutor as worthy of it.⁴¹

⁴⁰ One may be tempted to say that A's trust in S creates a context in which A's *belief* that S has told him that *p* is what entitles him to believe that *p*. The temptation derives in part from the fact that only in the testimonial case does the fact that you were told that *p* come apart from the fact that you believe you were told that *p*. You don't misremember what you had for breakfast by being mistaken about what your memory tells you you had for breakfast, in a way that doesn't reduce to your memory's being mistaken about what you had for breakfast. You don't misperceive the lengths of the lines in the Müller-Lyer illusion by being mistaken about what your eyes tell you about those lengths, in a way that doesn't reduce to your eyes being mistaken about the lengths. But you do sometimes mistakenly believe that it's noon by mistakenly believing that your interlocutor told you that it's noon, in a way that doesn't reduce to your interlocutor's being mistaken that it's noon. (What she said was 'Pass me a spoon.')

⁴¹ Remember, by 'trustworthy' or 'reliable' I always mean: trustworthy or reliable on the subject matter at hand. See note 39 above.

Reductionist accounts suffer from the problem that they would make taking a speaker at her word reasonable only when you don't need to. When you have evidence of a speaker's reliability you don't *need* to trust her: you can treat her speech act as a mere assertion and believe what she says on the basis of the evidence you have of its truth. You can ignore the fact that she's addressing you, inviting you. You can treat her as a truth-gauge. Anti-reductionist accounts, on the other hand, have suffered from the problem that they would make taking a speaker at her word reasonable when we think you'd be out of your mind to. If your interlocutor is utterly untrustworthy, you've no epistemic right, not even a *prima facie* right, to take her at her word. My account threads a path between these difficulties. Trust in a speaker can be a source of warrant for believing what she tells you even if you have no evidence of her reliability. But not, we must add, if you have positive evidence of her unreliability.

That yields two defeating conditions on the presumption that A's trust in S is reasonable and thereby a source of epistemic warrant. We might say that A's trust in S is *not* reasonable if either (i) S is untrustworthy on this subject, or (ii) A has good evidence of S's untrustworthiness on this subject. This suggests that A's trust in S would be reasonable in the case in which S is trustworthy but there is good evidence that she is untrustworthy, although A happens, from epistemic sloth, not to possess that evidence. Here A's trust in S would not be reasonable, we think, because it would be undermined by evidence which A ought to possess, where the 'ought' is epistemic. We must therefore reformulate the second defeating condition to read: (ii) *there is* good evidence of S's untrustworthiness on this subject, evidence moreover which A *ought* to possess, whether A is aware of that evidence or not.⁴²

Why not drop the first defeating condition, we might wonder, and let the reformulated second condition take up the slack? When externalists about epistemic warrant make the move I've just made in response to counterexamples, it seems to amount to a concession that is difficult to square with the spirit of externalism. For the view now seems to be that the subject's warrant in trusting derives entirely from the existence or nonexistence of evidence. Given that warrant is undermined by evidence of untrustworthiness, shouldn't warrant be sustained by evidence of trustworthiness? The putatively externalist account seems poised to qualify itself into a form of internalist account.⁴³

The distinction I have defended between asserting and telling motivates an insistence that the account not be so qualified. When A merely hears S assert

⁴² The 'ought'-clause leaves open the possibility that in a given case where there is evidence of S's untrustworthiness we nonetheless *don't* think A ought to possess that evidence—since an agent is not responsible for all his evidentiary failures. (Thanks to an anonymous referee for making me feel the need for this qualification.)

⁴³ For this objection, see Laurence Bonjour, *The Structure of Empirical Knowledge* (Cambridge: Harvard University Press, 1985), Chapter 3.

that *p*, his warrant for believing that *p* does indeed rest on whatever evidence he might gather of *S*'s trustworthiness. But when *S* tells *A* that *p*, the status of *A*'s trust in *S* as a source of warrant does not rest in any way on positive evidence of *S*'s trustworthiness. If *A* came to believe that *p* on the sole basis of positive evidence of *S*'s trustworthiness, he would not be treating *S*'s act as a telling addressed to him. Her act would thus go unconsummated.

Still, my account leaves it open that the negative defeating conditions should be understood along internalist lines in terms of the availability of evidence. As long as this stops short of smuggling in a positive evidential condition—that there must be evidence of trustworthiness in order for trust to be reasonable—I see no reason, in light of anything I've argued thus far, to exclude this interpretation. It is not inconsistent to hold that reasonable trust both requires that there not be evidence of untrustworthiness and does not require evidence of trustworthiness. An internalist reading of the defeating conditions on which the first condition, that *S* is untrustworthy, collapses into the second, that there is evidence of her untrustworthiness, need not violate my thesis of the presumptive reasonability of trust.

We create space for the internalist reading by formulating the defeating conditions in terms of a generic notion of trustworthiness, rather than the more specific notion of reliability. We can thus say that *S* is untrustworthy to the extent that there is evidence that she is untrustworthy. 'X is worthy of Y's trust' says that Y is (pro tanto) entitled to trust X. We can leave it open that what *disentitles* you to trust is the existence of relevant evidence, without thereby holding the same of what entitles you to trust. An internalist reading of the first defeating condition would not amount to an internalist view of the positive reasonability of trust, since this reading is compatible with trust's being reasonable in the absence of evidence that the trusted is worthy of it. In the next section I'll consider an externalist reading.

Whichever way we read the first defeating condition, my argument yields the following view of how telling and asserting can give warrant. When *S* asserts that *p* in *A*'s presence, but does not tell *A* that *p*, *A* cannot derive an entitlement to believe that *p* unless he reasons as follows: *S* is trustworthy on the question whether *p*, *S* has manifested her belief that *p*, so it must be the case that *p*. Here, *A*'s belief that *p*, should he form it, would be based on evidence of whether *p* is likely to be true. When *S* tells *A* that *p*, by contrast, *A* may derive an entitlement to believe that *p* without any such reasoning. In the case of telling, but not of mere asserting, *A* may derive an entitlement to believe that *p* simply from recognizing *S*'s intention that he trust her.⁴⁴ *A*'s

⁴⁴ There will be cases in which it is unclear whether a hearer should count as part of a speaker's audience—that is, one of those whom the speaker is inviting to trust her—because it is unclear whom the speaker intends to be addressing. There will be other cases in which, though her intentions are definite, the speaker is nonetheless not entitled to assume that her audience can recognize these intentions. In such a case, she shouldn't

belief that *p*, should he form it as *S* in telling intends him to form it, would be based not on reasoning from evidence of *p*'s truth but simply on his trust in her.⁴⁵

VII

It is a problem for an internalist reading of the defeating conditions on reasonable trust that trust can be reasonable even if there is no evidence of the speaker's trustworthiness. Indeed, even when there is evidence of the speaker's trustworthiness available to the addressee, it would, as I've argued, constitute an illocutionary 'abuse' to form your belief self-consciously on the basis of that evidence. One who tells offers an invitation to trust, not an invitation to consider evidence of her worthiness of trust.

The problem, specifically, is that it's hard to see how the mere fact that there is no evidence of a speaker's untrustworthiness could *ground* a hearer's entitlement to believe what the speaker tells him. On an externalist reading of the first defeating condition, the speaker must be in fact reliable, and the externalist could appeal to this fact to ground the hearer's entitlement. The internalist will in turn complain that it can be reasonable to trust a speaker

be surprised if some of her hearers take on trust what she hadn't invited them to take on trust. We can view these cases as similar to a sloppily tendered invitation: you mean to invite only people on your block to your party, but you misjudge who will see invitations posted on your block and lots of vagabonds from adjacent blocks come to believe they're invited. Are they? You could reasonably say that they weren't, though you'd have to take responsibility for leading them to believe they were.

This mirrors a debate within speech act theory about the nature of Austinian 'uptake,' an application of which informs the debate mentioned in note 26 above.

⁴⁵ One might codify this contrast by claiming that while in the former case *A*'s entitlement to believe that *p* would, in deriving from evidence, be empirical, *A*'s entitlement in the latter case would be apriori. While I am not unsympathetic with this claim, the issue is complicated. I cannot give a full treatment of the question, but let me briefly defend the claim of apriority against an objection suggested by my defense of the thesis that epistemic warrant could be grounded directly in trust.

I defended that thesis by comparing testimony with other sources of epistemic warrant such as perception and memory, and the comparison with other epistemic faculties may make it look dubious that testimony could give us apriori warrant. Perception doesn't give us apriori warrant. Why should testimony? If *S* looks out the window, sees that it's sunny, then tells *A*, how could *A*'s entitlement to believe that it's sunny count as apriori, given that *S*'s does not? We must formulate the entitlements more carefully, however. *A*'s entitlement to believe that *S* told him it's sunny is not apriori; he believes that on the basis of perception and memory. But we can ask: given that *S* told *A* it's sunny, what entitlement does he have to believe that it's sunny? His entitlement to believe that she told him it's sunny is not apriori, but one might claim that his entitlement to believe what she told him nonetheless is apriori. *S* and *A* are in different epistemic predicaments, and we must mark these differences in how we formulate their epistemic entitlements. That we can formulate *S*'s and *A*'s entitlements as 'an entitlement to believe what *S* told *A*' does not show that *S* and *A* must have the same entitlement to believe what *S* told *A*, or even the same kind of entitlement. *S* derives her entitlement from trust in her perception when she looked out the window and trust in her memory as she remembers the incident. *A* derives his entitlement from trust in *S* when she tells him what she saw.

who is not in fact reliable—provided there is evidence that she is reliable. I don't want to try to resolve the debate here.⁴⁶ Instead I'll simply mention this rationale for treating my argument in this paper as the first step of a larger argument for externalism, and move on to sketch some features of the externalist view of testimonial warrant that this larger argument, if successful, would vindicate.

If we interpret the first defeating condition in externalist fashion, we get a form of externalism strikingly different from reliabilism. A reliabilist view of testimony would view it as giving rise to warrant by virtue of the fact that the process by which subjects base their beliefs on testimony is a reliable process, that is, one likely to produce true beliefs. This leads immediately to the problem that there are as many different processes potentially in play as there are principles on whose basis we might view our subject A as forming his belief. Does his belief rest on the principle 'Believe S when she tells you something,' or the principle 'Believe someone in the circumstances that S is in when that person tells you something,' or the principle 'Believe anyone who tells you something'? These principles pick out different processes of belief formation, yet the question which process governs A's belief formation seems viciously underdetermined.⁴⁷ The question is underdetermined, on a reliabilist view, because the strong form of externalism entailed by reliabilism holds that we need not settle the question of a subject's epistemic entitlement to his belief by appeal to the subjective standpoint manifested by his belief formation.

Now I agree that the question of the subject's epistemic entitlement to his beliefs need not be settled by appeal to the *justification* he has for them: it need not depend on anything he is in position to cite. I join the strong externalist in denying this not only in the case of testimony but in the cases of memory, perception and inference I considered earlier. But my emphasis on trust amounts to a denial of the reliabilist's strong externalism. The subjective standpoint manifested in belief formation includes the subject's exercise of his capacity for reasonable trust. And in all the cases I considered the subject's warrant derives from his exercise of that capacity.

The reliabilist poses the question of warrant in terms of the reliability of a process of belief formation, I in terms of the reliability of the thing trusted. If the reliability of a process of belief formation determined warrant, then if S is

⁴⁶ I discuss the internalism/externalism debate at length in 'Trust and Reasons,' in preparation.

⁴⁷ This looks like an instance of the 'Problem of Generality' for reliabilism mentioned in Alvin Goldman, 'What is Justified Belief?' in George Pappas (ed.), *Justification and Knowledge: New Studies in Epistemology* (Dordrecht: D. Reidel, 1979), 12, and developed in Richard Feldman, 'Reliability and Justification,' *Monist* 68 (1985). For a reply to Feldman, see William P. Alston, 'How to Think about Reliability,' *Philosophical Topics* 23 (1995). Then see Earl Conee and Richard Feldman, 'The Generality Problem for Reliabilism,' *Philosophical Studies* 89 (1998).

reliable A will thereby acquire an entitlement to believe what she tells him, since forming beliefs on that basis is reliable. I agree with internalists that there can be no such entitlement, because there can be no entitlement external to a subject's capacity for cognitive self-governance. Yet it does not follow that all entitlements must be internally accessible to the subject. Trust is not a mere disposition to believe. Trust is not a 'mechanism' of belief formation; it does not merely cause beliefs but manifests the subject's cognitive self-governance. On an externalist view, the way a trusting subject governs himself is by letting himself be governed by the word of a reliable other. This is not internalist cognitive self-governance; it specifically abandons the project of apportioning belief to positive evidence, even to evidence of reliability. Still, one can hold this view without subscribing to the strong externalism entailed by reliabilist views, which treat the question of warrant in substantial independence from the subject's personal responsiveness to epistemic norms.⁴⁸

The externalist element would codify a dimension in which a subject can actively take responsibility and thereby become accountable for the course of his cognitive life even when passively acquiescing in the word of another. He takes responsibility by remaining answerable for the reasonability of his trust.

VIII

The externalism would mark a principal risk in trusting. It is difficult, at times, to know if trust is reasonable. A platitude has it that bestowing trust can itself make a person reliable. Perhaps your trust in a speaker, in making her reliable, will causally determine its own status as reasonable. Or perhaps you'll hope for this outcome in vain. Determining whether someone is reli-

⁴⁸ My point is not that reliabilist theories are *incompatible* with viewing subjects as responsive to epistemic norms. My point is that this responsiveness as such is not central to the theory. On a reliabilist theory, the subject's responsiveness to norms is manifested by a reliable process to which the subject is party by being disposed in a specified way. There is no question of his taking responsibility for himself or for his attitudes *as* responsive to these norms.

Although I cannot pursue the issue here, I find plausible the generalization of these remarks to the other forms of trust—in memory, in perception, in inference—discussed in section V. The trust is perhaps more habitual there than in the testimonial case (or does that underestimate the depths of our testimonial habits?), but in each case the subject can be viewed as taking responsibility for his trust by remaining answerable for its reasonability. I would argue that the cogency of sceptical challenges reveals an epistemic requirement that subjects take a kind of responsibility for their belief incompatible with forming it through the exercise of a mere disposition. (I explore this further in 'Judging as Inviting Self-Trust,' in preparation.)

able, in the relevant sense, is no easier than determining whether you can or could count on her. Sometimes only retrospection will tell you.⁴⁹

The risk you sometimes run in trusting has the effect of making it difficult, at those times, to know whether the entitlement you take yourself to have to believe what your interlocutor tells you is really an entitlement—that is, whether it really warrants your belief. In many ordinary cases, on the other hand, there is no such uncertainty. When you ask the fan walking back from the stadium who won the game, or a colleague what went on in the meeting, or your lover what he or she did on that trip to Paris, you trust your interlocutor, your trust is reasonable, and you know that your trust is reasonable. Given that most people are reliable most of the time, your trust in others is in another respect like your trust in endogenous sources of warrant. Whether it is your memory testifying to the contents of your breakfast or your uncle testifying to the contents of his stamp collection, unless you have positive reason to do otherwise it makes sense to trust first and ask questions later.

One might object that while trust may be presumptively reasonable with respect to some subject matters (the time, the weather, what happened when), it is not presumptively reasonable with respect to others. If a colleague walked up to you in the hall and simply *told* you what consciousness is, for example, you would not be reasonable if you took her word for it—whereas you would be if she'd told you it has started to rain. I can accommodate this difference, however, without altering my account. By 'reliable' I have all along meant *reliable on the given subject matter*. And there are some subject matters on which simply telling someone what's so immediately gives him evidence of your unreliability. Simply telling someone what consciousness is, instead of telling him that you have a theory of consciousness that goes thus-and-so, immediately gives him evidence that you have no idea how hard it is to come up with an adequate theory of consciousness and therefore that you are an unreliable philosopher of mind. Trust is presumptively reasonable with respect to all subject matters, but there are some subject matters on which an unhedged telling would immediately trigger the evidential defeating condition on the reasonability of trust.⁵⁰

⁴⁹ The notion of reliability in question is, it seems, a character concept. The concept of a person's character functions in part to codify where it makes sense and where it doesn't make sense to count on her. What makes trust in someone reasonable is not that person's history of truth-telling but her character here and now. When we say that a reformed liar must earn her interlocutors' trust, we don't mean that she must tell as many truths as she formerly told falsehoods before anyone is entitled to believe her. Perhaps you look her in the eye and see she has changed.

⁵⁰ Thanks to Steven Everson and David Velleman for help in formulating this objection and reply. We could complicate matters further, within the confines of my account, if we linked the triggering of the defeating condition with the way the telling is hedged. When you tell someone what consciousness is, there may be no amount of hedging that would suffice to avoid the defeating condition—since there is no such thing as 'expertise' on that subject. But there may be subject matters on which giving an impression that you

Admitting there's no guarantee she's reliable, what would tend to show that you can count on your interlocutor? One thing that would show her reliable is her ability and willingness to cite good reasons to believe what she tells you. Her citation of reasons is not what entitles you to believe what she tells you: your (reasonable) trust in her does that. But an ability and willingness to cite good reasons in support of what she says attests to her reliability as a testifier, which on an externalist view is what grounds the reasonability of your trust in her. Here at last we have an explanation of the phenomenon I described in section II. It constitutes illocutionary 'abuse' if you tell someone that *p* without being able or willing to cite reasons to believe that *p*. As I argued in section I, your addressee equally abuses the illocution if he fails to take your word for it that *p*, and challenging you to produce reasons to believe that *p* is incompatible with taking your word for it. So you probably won't have to cite those reasons. But if someone challenges your addressee to justify his belief that *p*, and he formed the belief on your word, then he is entitled to turn to you to satisfy his challenger. Your inability or unwillingness to produce the justification tends to undermine the thought that his belief is warranted by tending to undermine the assumption that trusting you on the matter was reasonable.⁵¹

Speakers who undertake such commitments are frequently unable to discharge them. But that doesn't vitiate my account of telling—not, at least, on the externalist reading. It merely shows that a speaker may succeed in making available to her interlocutor an entitlement for him to believe what she tells him without herself having a justification for believing it. Indeed, her belief may be entirely unwarranted. The externalist reading yields the interesting result that a speaker who is not herself entitled to the belief that *p* may, if she is reliable, make that very entitlement available to her interlocutor. The result is interesting but not odd.⁵² The interlocutor relies, after all, on the speaker's word—that is, on her status as trustworthy in testifying—not on whatever entitlement she may have to believe what she tells him, an entitlement that he needn't ever come to know.

IX

I have defended this analysis of testimonial telling (now expressed in traditional Gricean form):

appreciate how hard it is to discover the truth counteracts the evidence of unreliability you would otherwise give by simply inviting someone to trust you.

⁵¹ How much must you be able to say? As much as you would have to say to justify your own belief that *p*. In many cases, that may be pointedly little. If what you've told your addressee is that it's raining out, it usually suffices to answer his request for a reason with 'I just looked.'

⁵² I investigate the result further in 'Judging in Context,' in preparation.

- (T) S tells A that p (sincerely) iff A recognizes that S, in asserting that p, intends:
- (i) that A gain access to a prima facie entitlement to believe that p,
 - (ii) that A recognize S's (i)-intention, and
 - (iii) that A gain access to the entitlement to believe that p as a direct result of A's recognition of S's (i)-intention.

I've argued for this analysis by appeal to cases that reveal what's at stake for interlocutors when one sets out to tell another something. What's at stake, I've argued, is the establishment and maintenance—however briefly—of a trust relation.

I began with Austin's thought that epistemic warrant is transmissible and thus that if I tell you something I may be epistemically responsible for getting you into trouble. We've seen that this need *not* mean that my entitlement to believe what I tell you can thereby become your entitlement to believe it. It can mean that in telling you I make available to you a new entitlement. I can be responsible for your belief because in telling you I invite you to trust me, and in inviting you to trust me I take a degree of responsibility for your trust by certifying its status as reasonable.

I have not argued against scepticism about the possibility of telling anyone anything. I have argued merely that telling, where it occurs, constitutively aims to create a trust-based species of warrant.

Though one may be sceptical of either act, Austin turns out to be right: telling is indeed like promising. When I sincerely promise you I'll ϕ , I intend to make available to you an entitlement to perform acts that rely on taking me at my word. When I sincerely tell you that p, I intend to make available to you an entitlement to believe that p. If each case goes as I intend, in giving you my word I entitle you to take it.