

**CAN TRUST ITSELF GROUND A REASON TO BELIEVE THE
TRUSTED?**

Edward S. Hinchman

Trust in testimony is curiously self-reflexive. As Paul Faulkner emphasizes in *Knowledge on Trust*, when we trust a testifier, we rely not simply on her, as we might if we merely overheard her, but specifically on her attitude toward the fact that we are relying on her. This distinctive form of reliance explains the second-personal nature of the testimonial relationship: when you trust someone's testimony, you rely not merely on her attitude toward the proposition that she asserts but on her second-personally directed attitude toward you yourself as the one whom she is addressing. It creates space for the reasonability of trusting a speaker who lacks a track record of reliability, since you may be in position to gauge how her present conscientiousness and responsiveness to your epistemic needs outweighs any history of error or insincerity. But it makes explaining testimonial warrant trickier than explaining the warrant that we may derive from observing forms of behavior in which we are not so intimately implicated.

Thus far, I completely agree with Faulkner. Though the claim needs a fuller vindication than I'll attempt to provide, I'll take for granted that you can acquire an epistemic reason to believe that *p* by accepting a speaker *S*'s testimony that *p* even when *S* has a track record of error or insincerity. How might you acquire such a reason? Consider how a different reason might arise on *S*'s side of the transaction. Just as you can make a clock reliable by fixing its gears, so you can make a speaker reliable by giving her a reason to tell you the truth.¹ How can you give someone a reason to tell you the truth? I agree

¹ It would amount to an odd view of testimonial trustworthiness to insist that *S* could not become trustworthy until she had established an appropriate history. Though we cannot inspect the mechanism that makes a speaker trustworthy in the way that we can open the back of a clock, to insist that testimonial trustworthiness must be historically grounded is like insisting that a clock could not count as repaired until it had run properly for a while. Just as an unreliable clock can be made reliable more or less at once by undergoing a repair of its

with Faulkner that one way is by manifesting trust in her. But what is the relation between S's reason to tell you the truth and the reason that you have, on your side, to believe what S tells you?

Say the intervention succeeds. Imagine that your trust succeeds in giving S a reason to tell you the truth and that S acts on that reason in telling you that p. Imagine further that S's testimony gives you an epistemic reason to believe that p. I agree with Faulkner that your trust can give S a reason to tell you the truth, but can it also inform the reason that you have to believe what S tells you? Here Faulkner and I disagree. Faulkner argues that because the way in which your trust gives S a reason to tell you the truth makes it more likely that S is telling you the truth, your trust can count as evidence that S is telling the truth. I think that this argument proves too much, since it would let you bootstrap your way into possessing reasons that you clearly do not possess. What Faulkner's account leaves out is what you hope your trust will cause: here-and-now reliability in the speaker. Of course, Faulkner does not altogether overlook how the speaker's perhaps new-found reliability figures in your reason to believe what she tells you, since he emphasizes that your trust can make the speaker more likely than she would otherwise have been to tell you the truth. But he views the core of your reason to believe that p as deriving from the *reason* that you give S to be reliable in telling you that p, rather than from S's actual reliability.

It may at first seem that my objection addresses only a narrow issue in the epistemology of testimony: whether a reason to believe testimony can derive from the addressee's trust itself or only from reliability in the speaker that the trust perhaps causes. But beyond my narrow disagreement with Faulkner lie two broader issues. In section II, I'll argue that Faulkner misappropriates Bernard Williams's genealogy of testimony when he makes use of Williams's genealogical argument in his own preferred assurance view of testimony. Though Williams doesn't clearly articulate it, there is a deep reason why Williams's genealogy cannot underwrite an argument for trust-based testimonial reasons.

time-telling mechanism, so an untrustworthy speaker can become trustworthy more or less at once by undergoing a reform of her truth-telling mechanism – that is, of her dispositions to be accurate and sincere. The latter process is usually not as straightforward as the former, and determining whether it has occurred is not as simple as inspecting some gears. But there is no reason to doubt that such a process can occur and, as Faulkner emphasizes, some positive reason to be confident that it can occur through the addressee's own intervention. (I'll press this analogy further in section I.)

This raises a second issue: can a genealogical argument underwrite any version of the assurance view of testimonial reasons – that is, any view that like Faulkner’s emphasizes the second-personal trust relation between speaker and addressee? In section III, I’ll argue that the answer is yes. I’ll sketch an assurance view of testimonial reasons that rejects Faulkner’s thesis that such reasons could be grounded in trust.

On my alternative assurance view, a testimonial trust relation derives from the addressee’s eliciting and from the speaker’s exercising a species of reliability that brings the two of them into a second-personal relation of mutual recognition. This is not the truth-conducive reliability on which reliabilist arguments in epistemology typically focus. It is instead a knowledge-conducive reliability that pertains to the addressee’s epistemic entitlement to treat the evidence at hand as sufficient, in his particular context, for him to close doxastic deliberation by forming a belief. I’ll thus call it ‘closure-conducive’ reliability. I’ll argue that when we see how warranted testimonial belief requires that the speaker be both truth-conducively and closure-conducively reliable – where the latter requires that she be responsive to the addressee’s context-sensitive epistemic needs, beyond his broad need to believe the truth – we’ll see how there is an ineliminable role for assurance in the epistemology of testimony. And we’ll see how filling that role requires rethinking what would count as a vindicating genealogy of testimony.

I. Is it trust that grounds a testimonial reason, or the reliability that trust causes?

The core of Faulkner’s argument lies in Chapter 6, where he explains how the presumption of trustworthiness informing your affective trust in S provides what I’ll call a testimonial reason: an epistemic reason to believe what S tells you.² This argument is not quite complete in Chapter 6, since a key part of it rests on a claim that he defends in Chapter 7. In the genealogical treatment of testimony offered in Chapter 7, Faulkner grounds his account of testimonial reasons in an observation about our contingent social condition: that

² Faulkner aims to explain not testimonial reasons in general but only the testimonial reasons that derive from assertions considered as tellings – that is, as addressed to the one who gets the reason. He argues that there can be testimonial reasons that do not derive from assertions considered as tellings – for example, the reason an overhearer gets to believe what S tells someone else. This distinction won’t matter for my purposes in this paper.

we have left the State of Nature via the intrinsic value that we place on sincerity and accuracy, the twin virtues of truthfulness. This yields an account of testimonial reasons that is neither reductive nor straightforwardly non-reductive. On Faulkner's account, testimonial reasons cannot be reduced to reasons derived from our predictive faculties, because the affective trust at the core of testimonial relations is crucially different from predictive species of trust. But his account is not straightforwardly non-reductive either, since it denies the principle of credulity (on which we can get a testimonial reason by accepting what the speaker asserts entirely credulously).

Faulkner thus aims to sidestep the dialectic that informs much recent work in the epistemology of testimony. But how does Faulkner intend his account to work within the dialectical context that motivates it? That dialectical context emphasizes what Faulkner calls *the problem of trust*, a problem that I'll examine in detail as we proceed. On Faulkner's most general formulation, it is the problem that testimony is a product of the speaker's communicative intentions and as such should not be treated on the model of a regularity in nature. On the assumptions (a) that the audience is specifically ignorant of what moves the speaker to communicate in a given case and (b) that the speaker may be motivated by self-interest, there can be no *general* norm entitling audiences to believe speakers. Faulkner's explanandum is the entitlement to believe not merely what the speaker says but specifically to *believe the speaker* – that is, to believe what she says on her say-so. When, then, is a given audience entitled to believe the speaker? How, that is, could the audience acquire a reason to believe on the speaker's say-so? Only, Faulkner argues, in a context of affective trust.

Faulkner's argument turns on a distinction between two species of trust, predictive and affective, which he defines as follows³:

(PT) A trusts S to ϕ (in the predictive sense) iff (1) A depends on S ϕ -ing, and (2) A expects S to ϕ (where A expects this in the sense that A predicts that S will ϕ). (145)

³ I've added the labels – '(PT)' and '(AT)' – but the definitions quote Faulkner's text verbatim.

- (AT) A trusts S to ϕ (in the affective sense) iff (1) A depends on S ϕ -ing, and (2) A expects (1) to motivate S to ϕ (where A expects this in the sense that A expects it *of* S that S be moved by the reason to ϕ given by (1)). (146)

These definitions yield two key differences between predictive and affective trust. The first difference concerns the nature of the expectation in clause (2) of each definition: the expectation in (PT) is purely predictive, whereas the expectation in (AT) is also normative – a matter of what A presumes that S *ought* to do. The second difference concerns the content of these expectations: in (PT) A expects that S will ϕ , whereas in (AT) A expects that S will ϕ for the reason that A is depending on her to ϕ .

One might try to argue that the second difference entails the first. When you expect that S will ϕ specifically for the reason that you are depending on her to ϕ , does not your expectation presume that your dependence gives you a claim right over S, a presumption that in turn makes your expectation that S will ϕ normative rather than merely predictive? No, your expectation that S will ϕ for the reason that you are depending on her to ϕ does not necessarily include a presumption that your dependence gives you a claim right over S. Sometimes it does, and those are the cases on which Faulkner is focusing. When you depend on S for the truth, it is plausible that your dependence gives you a claim right to the truth from S. But you might depend on S to ϕ , while acknowledging that your dependence gives S a reason to ϕ , given background conditions, without presuming that you have a claim right over S to ϕ – that is, that you could rightly criticize or resent S for failing to ϕ . Perhaps S is your student or research assistant, and you've asked her to do you a small favor in the context of your research. You might expect that your dependence on S gives S a reason to do you the favor, partly because one thing that S is in the business of doing is proving her research potential and will in turn depend on you to attest to that potential in a letter of recommendation. So you do expect S to do you this favor because you expect that your dependence on her to do it gives her a reason to do it. But you don't for a moment presume that you have a claim right over S for this favor. In fact, it's quite the opposite:

you're hesitant to impose upon S precisely because you expect her to be especially responsive to the imposition, given the institutional context in which you make it. Faulkner does not explicitly link the two components of (AT) in the way that I'm criticizing, so the point isn't an objection to his account. But the possibility of such a link will matter to an issue that I'm going to raise in section IIIb below.

How does Faulkner's appeal to affective trust solve the problem of trust, given that it does not falsify the two assumptions that generate the problem? Here is his core move:

[I]n trusting S to ϕ , the grounds of A's attitude of trust are the belief that S can recognize his, A's, depending on S ϕ -ing, and the presumption that this will move S to ϕ . Thus, A will perceive the situation defined by this act of trust as one wherein S has a reason to ϕ . So other things being equal A will presume that S will ϕ . If this turns out to be true and S acts as A expects, S will have proved trustworthy. So in affectively trusting S to ϕ , A presumes that S will prove trustworthy just as in predictively trusting S to ϕ , A would believe this. This is not to suggest that trust involves A reasoning to this conclusion but is rather to claim that in trusting S to ϕ , A makes this presumption. However, the presumption that S will ϕ rationalizes A's act of trust in the same way that the belief that S will ϕ would do so. Consequently, the act of trust is rationally *self-supporting* in that it is based on an attitude of trust, which through implying the presumption that the trusted is trustworthy, gives a reason for trusting. (151)

And here is the core move applied to the case of testimonial trust:

[T]he attitude of affectively trusting a speaker for the truth provides an epistemic reason for believing the speaker's testimony. For suppose A trusts S for the truth as to whether p and S tells A that p. Then A's attitude of trust, I argued [above], involves A accepting various propositions about S and the trust situation, where the acceptance of these propositions defines what it is to see depending on S for information as to whether p in the positive light of trust. So in affectively trusting S for the truth, A accepts that S will see his, A's, depending on S for information as to whether p as a reason to tell A the truth on this matter. So trust involves A accepting that S has a reason to tell him the truth, and accepting that S will act on this reason, other things being equal.... In accepting these things about S and the trust situation, A thereby presumes that S is trustworthy, or that S will tell him the truth and will do so for the reason that he, A, depends on S for this.

This presumption need not amount to the belief that S is trustworthy since its ground is things which need be merely accepted in the trust situation. However, this presumption, like the belief with the same content, makes it probable for A that *p* is true given that this is what S tells him. So A's attitude of trust raises the probability of *p*, when this is what S tells him. So A's trusting S for the truth, in a situation where S tells A that *p*, provides A with an epistemic reason to believe that *p*. (154)

Let's consider the core move in light of two worries. First, why are we reasoning from A's trusting presumption that S will see A's dependence as a reason to tell him the truth? The problem case is the one where S tells A that *p* *without* being motivated by any recognition of A's dependence – that is, where the presumption is false. Second, why not suppose that what needs to be shown is not that A's attitude of trust has given S a reason to be trustworthy but that A's attitude of trust *has succeeded in making* S trustworthy in this specific instance? In the kind of case we're considering, we're entitled to assume that S knows that *p*, and (if this is not entailed by the first assumption) that it is indeed the case that *p*. But we're not entitled to assume that S is relevantly trustworthy – since the problem of trust is precisely that S may be untrustworthy in a given case.⁴

Faulkner replies to the second worry that if A's attitude of trust can explain why S's telling is likely to be true, then that attitude of trust can itself figure as evidence. But, elaborating the second worry, we may note that the trust itself doesn't *directly* explain why the telling is likely to be true. What directly explains why the telling is likely to be true is the speaker's truth-conducive reliability – that is, her disposition to assert the truth – in this interlocutory context. The trust may explain why the speaker is reliable. But it's the reliability that explains why the telling is likely to be true. Compare: you can cause your clock to be a reliable indicator of the time by repairing it, but the repair job – even if it is ongoing: say you have to hold the cord at a precise angle to retain the electrical connection

⁴ The parallel formulations in an earlier paper – “What is Wrong with Lying?” (*Philosophy and Phenomenological Research* 75:3 (2007), 553) – made it seem that Faulkner holds that the epistemic reason is provided not by the audience's affective trust but by the *de facto* affective trust relation between speaker and audience. The reason is created “when things go right,” “when this presumption is fulfilled.” I take it that Faulkner would no longer put his view like that. If there no reason unless the speaker is (or shows herself to be) as the audience's presumption represents her as being, then it is not the *attitude* of affective trust that provides the reason but the *relation* of affective trust, which is something else entirely: the attitude plus appropriate responsiveness to the attitude on the speaker's side.

– isn't itself evidence of the time. It causes evidence of the current time to be produced, but it itself is not that evidence. The evidence is what the clock says, given that the clock is reliable.

Faulkner would reject this comparison because trusting a clock for the time manifests predictive trust, whereas trusting a speaker for the time typically manifests affective trust, where affective trust differs from predictive trust insofar as its expectation of performance is normative – not merely an expectation-*that*, but an expectation-*of*. But why should this difference matter to the point at issue? With a little science fiction, we could imagine a multi-step process of creature construction moving from merely predictive to fully normative reliance: (a) a regular clock that works only if you hold the cord just so; (b) a speech-interpreting clock that works only if you keep saying 'Hold that connection'; (c) a mind-interpreting clock that works only if it interprets you as wanting it to hold that connection; (d) a norm-sensitive clock that works only insofar as it is responsive to your stance of normatively expecting it to hold that connection.⁵ Your intervention or influence clearly does not count as evidence of the time in (a) or (b). But if your intervention or influence counts as evidence in (d), as it seems Faulkner would have to say (consistently with his theory, assuming that relevant background conditions are met), then why not say the same of (c)? But (c) seems a mere extension of (a) and (b). Why should the introduction of the normative element in (d) make this difference?

Faulkner's full answer to this question will rest on the genealogical argument that he develops in Chapter 7 and that we'll consider in section II below. But we can anticipate one issue for that argument by considering norm-sensitivity that is purely instrumental. You can 'give S a reason' to tell you the truth by betting her or by threatening her. In such a case, it may be that the bet or threat 'explains why' S is likely to be asserting the truth – because without the bet or threat S would have lied or been less careful. While you might cite the bet or threat in explaining why you're entitled to believe what S asserted, it seems very odd to say that the bet or threat itself figures as evidence or as the basis of an epistemic reason – *unless*, of course, we regard the evidence or reason as lying in S's status

⁵ For the idea of a 'creature construction' see Paul Grice, "Method in Philosophical Psychology," in his *The Conception of Value* (Oxford: Oxford University Press, 1991), especially section V.

as a reliable testifier and regard that status as simply including the bet or threat. So what would be reliable is: S-having-bet or S-under-threat. And we can now perhaps make a move in the direction of Faulkner's conclusion. We can say that there's no need for a hyphenated condition when the explanatory factor gives S an intrinsic reason: we don't need to say 'S-insofar-as-she-cares-about-the-norms-of-trust' because her care about the norms of trust somehow figures in her practical identity. But this move presupposes a reliabilist framework that Faulkner wants to reject.

What Faulkner wants to say is that the bet or threat is *the* explanation why S's assertion is likely to be true, just as trust is *the* explanation in the central trust cases. But without an appeal to S's reliability, an account of A's testimonial reason would admit the possibility that A has bootstrapped his way into possession of a reason through his mere affective trust in S. The problem for Faulkner is that we do not think such bootstrapping is actually possible.⁶ Trusting testimony no more generates a reason to believe what the speaker asserts than trusting a promise on its own generates a reason to perform acts that depend on the assumption that the promisor will keep her promise. In each case, I would argue, the trust can at best cause the speaker to be relevantly reliable, where the reliability would in turn provide the reason.⁷ We can cast the threat of illicit bootstrapping in terms that parallel Faulkner's own problem of trust. The problem is this: that even though the

⁶ Faulkner raises and responds to a worry about bootstrapping in the following passage, but this is not the worry that I'm pressing:

It might seem odd that trust can bootstrap itself into reasonableness in this way. However, this oddness should be lessened once it is clear that trust is both an attitude and an action and that what is being offered is an account of the interaction between these two aspects of trust. The dynamic by means of which reasons for trusting are generated can then be clarified by separating out the temporal stages wherein an act of trust follows a decision to trust. (151)

Here and in his discussion through 153, Faulkner seems to assume that the only worry about bootstrapping to which his account might be susceptible is the worry that an attitude of trust commits the trusting to act in a way that is insensitive to evidence that the trusted is not worthy of the trust. That is not my worry. I take for granted Faulkner's point that "in deciding to trust S to ϕ , A does not decide to trust come what may" (152). My worry applies most sharply when there is no evidence of S's untrustworthiness available to A, though S is nonetheless unreliable. Faulkner claims that in such a case A may have an epistemic reason to believe what S tells him. I'm arguing that that would involve illicit bootstrapping.

⁷ I lack space to defend the claim about promising here. For a full defense, see my "'You May Rest Assured': A Theory of Normative Powers," in preparation. Of course, promissory reliability is not the same as testimonial reliability.

presumption at the core of A's affective trust *can* make S trustworthy, *for all A knows* S is not in fact trustworthy in this instance – indeed, for all A knows, S is merely *exploiting* the presumption at the core of the attitude that she invites A to take toward her when she tells him that p. If we say that A's affective trust gives A a reason to believe what S tells him, we appear to be merely wishing this problem away. A's trust may include a *belief* that A has a reason to believe what S tells him, but that belief merely sets the stage for the problem of trust. If your attitude of trust itself gave you a reason to believe what the trusted speaker tells you, there simply would be no such problem as the problem of trust that Faulkner describes. The problem of trust is that the attitude of trust cannot itself provide a reason to believe the trusted, independently of the speaker's status as reliable.

Why does Faulkner eschew such an appeal to reliability? He regards a reliabilist framework as failing to do justice to the interpersonal element in testimonial trust relations. But he thereby overlooks the possibility of arguing as I will, with a emphasis on closure-conducive reliability. In order to make available a reason to believe what she says, a speaker must be both truth-conducively and closure-conducively reliable. The dimension of closure-conducive reliability gets the second-personally normative element fully in play, I'll argue, by ensuring that S's status as fully reliable includes her responsiveness to A's doxastic predicament. Since that's what serves to distinguish the addressee's stance from an overhearer's stance – that the former but not the latter trusts in a way that presupposes that he is the recipient of such second-personal responsiveness – I'll conclude that testimonial reasons are irreducibly second-personal, providing the epistemic upshot that we would expect to derive from an assurance view of testimony, on which a testimonial telling is an invitation to trust.

II. Can a genealogy vindicate trust-based reasons?

I haven't thus far considered what Faulkner regards as the key piece of his argument for trust-based testimonial reasons: the genealogical argument that he develops in Chapter 7. Building on Bernard Williams's genealogical argument in *Truth and Truthfulness*, as well as on his own argument in Chapter 6, Faulkner now argues that the problem of trust can be disarmed when certain social conditions are met. Again, the problem of trust is that a given

speaker could prove unmoved by the reason that your trust presumes she has to tell you the truth. Faulkner argues that this possibility does not constitute a deep threat to the possibility of trust-based testimonial reasons because even a speaker unmoved by it will – assuming the social conditions are met – nonetheless have this reason to be truthful.

In the present section I'll argue that there are two problems with Faulkner's genealogical argument. First, Faulkner's argument seems to presuppose without argument the denial of a plausible and widely accepted internalist thesis about reasons. Second, his account cannot distinguish between the reason-giving force of affective trust and the reason-giving force of what I'll call 'institutional' trust, despite the fact that institutional trust cannot serve as the ground of a testimonial reason.

IIa. The problem of internalism

While a speaker could prove unmoved by your trust in a given case, Faulkner's genealogical argument aims to ensure that the speaker does nonetheless have a reason to be moved by it. Let's call that reason her *aretaic* reason, to contrast it with the *testimonial* reason that your trust presumes that you have to believe what she tells you: 'aretaic' because it is a reason for her to speak with sincerity and accuracy, the virtues of truthfulness. For Faulkner's account of testimonial reasons to work, something in the broader social practice must ensure that the speaker does actually have the aretaic reason. Faulkner's strategy is to turn that necessary condition on trust-based testimonial reasons into a sufficient condition, arguing that if we have escaped the state of nature to the extent that we are entitled to take for granted that our affective trust gives any speaker who addresses us an aretaic reason to tell us the truth, then we are equally entitled – epistemically entitled – to believe what the speaker asserts simply on her say-so. If the argument works then the problem of trust is not the problem that it appeared to be – at least, not for us. More exactly: the problem of trust becomes the problem that a speaker may prove unmoved by a reason to tell you the truth that she does nonetheless continue to have. As we'll see, this domesticates the problem of trust by delinking it from a natural application of an internalist thesis about reasons: that in order to have a reason to ϕ , S must

have at least a sound deliberative route to a motive to ϕ .⁸ Faulkner argues that that element of risk is compatible with our nonetheless having a testimonial reason. After all, even accounts that put S's reliability at the core of A's reason to believe what S asserts have to live with the possibility that S's assertion will prove false though S herself is reliable. Faulkner aims to relocate the risk from A's reliance on a belief-forming mechanism presumed reliable to A's trust in a intentional agent presumed to be moved by that trust. Doesn't the latter formulation better capture the riskiness of taking someone at her word? Indeed it does. I'll endorse an alternative formulation of that riskiness when I sketch my alternative account of testimonial reasons in section III.

We're now ready to see the first problem with Faulkner's genealogical argument. Why should we believe that the aretaic reason at issue – that is, the speaker's reason to be truthful, grounded in your trusting dependence on her to be truthful – is a reason that just any speaker in our practice will have? Perhaps if the speaker is 'one of us' in relevant respects she must value sincerity – construed as involving the dependence-responsiveness posited by affective trust – *in general*. But it would be an absurdly strong claim to say that a given speaker must, to count as a participant in our practice, prove responsive to *your* dependence on her. The protagonist in the problem case is someone who may value sincerity in general but who in a given case *makes an exception of herself*. She wouldn't lie to her friends but she will to you, her mere business partner. Or she wouldn't lie to people whom she believes likely to smoke her out, but she will to you, whom she believes naive. Perhaps she wouldn't lie to anyone in her quotidian sunny mood, but that changes with the descent of crepuscular anger. And so on. Can genealogical reflections show that affective trust gives rise on its own to any reason for an audience to believe what a speaker tells them – even in social contexts where affective trust is generally valued? The problem of trust seems undomesticated.

Faulkner claims that the normative expectation at the core of affective trust is reason-giving at least in communities bound together by norms of trust. In such a

⁸ For this use of 'sound deliberative route' see Williams's "Internal Reasons and the Obscurity of Blame," in *Making Sense of Humanity* (Cambridge: Cambridge University Press, 1995), 35; which extends the argument in "Internal and External Reasons," in *Moral Luck* (Cambridge: Cambridge University Press, 1981).

community, failure to live up to this expectation brings a sanction administered through reactive attitudes. This gives anyone who is committed to the norms governing community interactions – anyone who thereby gives standing to the reactive attitudes – a reason to live up to the expectations. And perhaps (though Faulkner doesn't argue this way) only an insider, someone who shares the identity, can understand how the reactive attitudes are to be administered and thus what motives they inculcate. So perhaps – if the identity is somehow constructed or sustained by that encounter – we can get round the worry that the reason-giving aspect of the trust is merely predictive. But now the question, the whole problem of trust, simply shifts ground: why think this speaker is committed to these norms *in application to the present case*? To say that she is committed to the norms in application to this case is to say she is relevantly trustworthy. So in wondering whether she is committed to the norms we're back wondering how to solve the problem of trust. The loop through local identities seems not to have helped.

The problem again appears to derive from an assumption that Faulkner does not make explicit: the falsity of a practical internalism about such aretaic reasons. On this internalism, there's a link between the speaker's possessing the reason that your presumption, when you trust, presents her as having and some fact about her motivations – let's call it, generically, the fact that she *could care* about not misleading you, whether for your own sake or in order to avoid others' disapprobation. The problem of trust is that *for all you know* the speaker not only does not care about not misleading you but could not be brought to care. I don't think we need to appeal to the concept of a psychopath to make this possibility clear. There may be people who simply dislike you or are angry with you to such a degree that they are incapable of caring not to mislead you, and there may be predicaments that make people desperate or despairing to such a degree that they are incapable of caring not to mislead any addressee, however they may feel about him. But it's easier to focus on psychopaths, and to note that a psychopath is defined as someone who could not care – someone who has a deficit in their capacity to care – about the normative pressures that Williams's and Faulkner's genealogical stories emphasize. If internalism is true, the problem of trust is that for all you know the speaker addressing you is a psychopath and as such does not have a trust-based reason not to mislead you.

Of course, Faulkner can avoid the problem by simply taking for granted that internalism is false. Indeed, that seems to be his strategy. But the denial of internalism – that is, externalism – is a highly controversial position about the nature of aretaic reasons, just as it would be about the nature of practical reasons in general.⁹ And the problem for Faulkner is actually worse than that, since his problem of trust would appear to survive any such assumption. If we assume that even a psychopath has a reason not to mislead you, however incapable she is of being motivated to act on the reason, then the problem of trust becomes the problem that for all you know the speaker is incapable of being thus motivated. The fact that any speaker does, assuming externalism, have an aretaic reason to be truthful would not at all tend to show that you have a testimonial reason to believe what she says.

These reflections on internalism and externalism reveal an incoherence in the dialectic framing Faulkner's argument for trust-based reasons. His solution to the problem of trust – the solution that works by positing a link between the speaker's aretaic reason and the addressee's testimonial reason – actually depends on rejecting an externalist view of aretaic reasons. If we adopt an externalist view of aretaic reasons, then showing that the speaker possesses a reason to be truthful does not solve the problem of trust, since on an externalist view of reasons a speaker can have a reason to ϕ while lacking a sound deliberative route to a motive to ϕ . On an externalist view of reasons, a speaker can thus possess a reason to be truthful while being, as we might now put it, *deeply unmotivated* to tell her interlocutor the truth. This possibility shows that if he adopts or assumes an externalist view of aretaic reasons then Faulkner cannot solve his own problem of trust. If Faulkner instead embraces an internalist view of aretaic reasons, he thereby gets round this problem, since an internalist view ensures that a speaker who has a reason to tell her interlocutor the truth cannot be deeply unmotivated to tell him the truth. But now Faulkner confronts directly what I've been calling the problem of internalism: an internalist link between having a reason to be truthful and being motivated (or having a sound deliberative route to a motive) to be truthful ensures that speakers with psychopathic tendencies do not

⁹ For a treatment of the debate over internal and external reasons that emphasizes interpersonal trust in practical reason-giving, see my "Trust and Reasons," in preparation.

have a reason to be truthful. Whichever view of aretaic reasons Faulkner adopts, his argument confronts a serious problem. The dialectic within which he is pursuing his argument for trust-based testimonial reasons appears to have painted him into a corner.

IIIb. The problem of institutional trust

Let's now put that first problem aside. Even if we imagine that Faulkner's genealogical argument could somehow ensure that you'd never receive testimony from a psychopathic or otherwise similarly incapacitated speaker, it nonetheless confronts a second problem. Recall that by Faulkner's definitions (PT) and (AT) – given in section I above – affective trust differs from predictive trust in two respects: in its nature, affective trust involves an *expectation of* the trusted, not merely an *expectation that* the trusted will act as she is trusted to act; and in its content, affective trust specifically involves an expectation of the trusted to act on a motive generated by recognition of the trusting's dependence on her. The problem is that there are cases that satisfy Faulkner's definition of affective trust, (AT), without thereby amounting to cases that are *second-personal* in the way that Faulkner assumes any case of affective trust must be. The problem cases are cases of institutional trust, wherein being trustworthy in relevant respects is part of the trusted's role or job within the institution. The institutional nature of the trust ensures that it differs from predictive trust in both of the respects described by (AT). But the institutional nature of the trust equally ensures that it is not second-personal and therefore that it does not involve trust-based reasons.

Cases of institutional trust are common. I trust the bank teller to give me prompt and competent service concerning my banking needs, and in doing so I expect her to be moved by my dependence on her but not solely or even primarily by my *particular* dependence on her. I merely trust her to do her job. Part of doing her job involves knowing when and how customers are depending on her in ways to which her job requires her to be responsive. Say I know that this teller hates me and would probably mess with me in another context. Say I also know that, like most people, she takes her job seriously and acts with the integrity that is the natural expression of that attitude. When I trust her to transfer or deposit funds, my trust manifests a normative presumption not about her relation

to me but about her relation to her job and more broadly to the institutional context in which she performs it. That presumption is indeed fully normative: I expect no less of her. And it does involve the thought that she is motivated by my dependence on her. If she failed to be responsive to a customer's banking needs, I would judge her harshly, and expect others to judge her harshly, with the full reactive-attitudinal wallop characteristic of such judgments. (I expect no less of them, and I expect no less of myself than to accord others the respect that such normative expectations express.) I am not merely predicting that she's reliable. I may have no basis for such a prediction. If the little that I know about her suggests that I must have some predictive expectations, switch to a case in which I trust a store clerk about whom I know nothing. Barring a reason for mistrust, I'll trust the clerk to meet my pertinent shopping needs, on no other basis than that it is part of his job. Again, it isn't about his attitude specifically toward *me*. He may barely have seen me or may address me as part of an amorphous crowd of shoppers. Though these details do not entail that testimonial reliance on the teller or clerk could not be predictive – with added details, it appropriately could be predictive – these do not sound like cases in which trust in the speaker is affective in Faulkner's sense, with emphasis falling on a presumption about the speaker's attitude toward this specific instance of trusting dependence. In each case the trust seems fundamentally institutional, with emphasis falling on the relation not between the trusted and the trusting but between the trusted and her job or other institutional affiliation.

I agree with Faulkner that an assurance view of testimony ought to emphasize the second-personal trust relation specifically between the speaker and the addressee, not a species of trustworthiness grounded in the speaker's relation to her job or other institutional affiliation. But how in general should we think of the difference between institutional trust and a genuinely second-personal species of trust? Here again I think it helps to use the fiction of a creature construction; we might thus imagine a transition from a merely predictive reliance on a bridge through institutional trust in someone doing his best to replicate a bridge to, finally, genuinely second-personal trust in such a bridge replicator. Consider four cases (and forgive the cartoonish nature of the final three; I'm aiming at simplicity): (a) you rely on a short but rickety bridge to carry you across a dangerous chasm (just a bit too far to jump), (b) you rely in the same way on a rigid-looking man whom you

encounter bizarrely but apparently securely spanning the chasm (his eyes staring vacantly, his toes and fingers gripping gnarly roots), (c) you rely on a man in that pose whom you know to have been hired by the forest service to perform this important job but who is otherwise inattentive (reading a magazine, humming to himself, etc.), (d) you rely on a man who may or may not be making any money in the chasm-spanning racket but who looks you in the eye and asks you to trust him.

Since there are no expectations-of in play in (b), I expect Faulkner to share my sense that that case is no more an instance of second-personal trust than (a). And I expect Faulkner to share my sense that (d) clearly is an instance of second-personal trust. But what of (c)? It seems clear to me that if (b) isn't second-personal trust – that is, the kind that figures in the species of trust relation that Faulkner is trying to theorize – then neither is (c), even though (c) generates a expectation of performance that seems to meet clause (2) of Faulkner's definition of affective trust, (AT). What generates the expectation in (c) are the considerations (i) that serving as a reliable bridge is this man's job (with all that that involves) and (ii) that this particular job has moral implications. (Note that I'm not asking whether the man has a moral obligation to serve as a bridge. I'm asking whether your reliance on him as a bridge has moral content. In case (b) it does not, but in case (c) it does: you're counting on him to take seriously his important job.) If the man in (c) fails to do his job, that may naturally trigger a reactive-attitudinal response.¹⁰ And a crucial part of what it is for him to do his job is to be motivated appropriately by the recognition that a traveler is depending on him to avoid plunging to his or her death.

As far as we've described case (c), there is nothing to show that you are not just predictively trusting this man, albeit in a way that has moral, and *thereby* reactive-attitudinal, content. For one thing, your moral expectation has nothing specifically to do with his relation *to you* – toward whom, after all, he is being quite inattentive. (Imagine again a rude clerk: you suspect he would gladly steer you in specifically the wrong

¹⁰ I'm not sure that *resentment* is the right term for your reactive attitude here. Even though Strawson used 'resentment' as the catch-all term in his famous paper, it seems likely that you'd not be so much resentful as *disappointed* in the man, a disappointment that has nothing specifically to do with his relation *to you* (apart, of course, from the fact that he has let you fall into the chasm! – but that would be true in cases (a) and (b) as well, should they go wrong).

direction if his job didn't demand otherwise. You have expectations of him, but this is not second-personal trust.) From your perspective, you assess his bridge-relevant reliability in generically the same way as you do in (b). The specific difference is that in (b) you're interested only in organismic qualities of the man (is he asleep? in a trance? stiff-enough?), whereas in (c) you're also interested in some social qualities (does he take his job seriously? is he well enough paid?). But in each case you're merely looking for evidence of an incipient unfortunate collapse. What makes (d) crucially different in this dimension, I would say, is that your reliance on him naturally rests on a presumption of trustworthiness that projects a relation of mutual recognition. But if that's right, then the key to distinguishing non-second-personal trust – whether predictive or institutional – from genuinely second-personal trust lies in grasping the force of a species of mutual recognition that can't be defined along the lines of Faulkner's (AT). We'll return to this idea in section III.

Let's see how the problem arises for Faulkner's (AT). In section I, I suggested a link between the two components of affective trust, as Faulkner defines it in (AT). I suggested that when you expect that a speaker will ϕ specifically for the reason that you are depending on her to ϕ , your expectation presumes that you have a claim right over her, which in turn makes your expectation normative rather than merely predictive. I went on to criticize the suggestion, arguing that what amount to cases of institutional trust – in my example, trusting a research assistant to help you with your research – may manifest an expectation that S will ϕ for the reason that you're depending on her to ϕ without thereby manifesting any presumption that you have a claim right over S to ϕ . But suppose the link nonetheless holds. We can now see that even such a link would not distinguish institutional from properly second personal trust. Even if your expectation that S will have a reason presumes that you have a claim right over S, that claim right may be mediated by an institutional context in a way that deprives it of second-personal content of a sort that could be articulated in terms of a trust-based reason. Even if your dependence on the teller gives you a claim right to prompt service from her, or if your dependence on the clerk gives you a claim right to conscientious advice about some product that she is trying to sell you, these claim rights are not themselves grounded in your trust but in the institution structuring the

exchange. I don't deny that there is a genuine phenomenon toward which Faulkner is gesturing with his definition of affective trust, (AT). But (AT) has not captured that phenomenon. The phenomenon in question is a trust relation that is irreducibly second-personal, involving a claim right that pertains specifically to the parties to this particular trust relation. As we'll see in section III, what (AT) leaves out is the way in which second-personal trust is mediated by a relation of mutual recognition.

One might at this point wonder if there is such a thing as a genuinely second-personal trust relation. If institutional trust can fail to be second-personal even when the trusted is motivated by recognition of the trusting's dependence on her, why think testimonial trust is ever second-personal? Here the dialectic gets more complex than I can cope with in this paper. The dialectic that we're pursuing begins from a point of agreement with Faulkner's argument: that, even if the epistemology of testimony is not simply the epistemology of testimonial tellings,¹¹ the latter is (a) a crucial part of the former and (b) a matter of irreducibly second-personal trust relations. My argument in the present subsection is that Faulkner is not in position to offer a compelling explanation why this is so, for the simple reason that his attempt to define the distinctively second-personal element in testimonial trust falls short of its aim, since it is compatible with cases in which the trust is not second-personal. One way to put my point is to note that Williams's treatment of testimonial trust emphasizes what we might naturally understand as its institutional nature – not that Williams conceives everyone as having a 'job' such that every addressee is in some respect a 'customer' or 'client' but that Williams conceives speakers as having been inculcated into a practice of what he calls 'normal trust,' a practice that has a fundamentally institutional nature insofar as it is upheld by norms of shaming, shunning, and the like. In this respect, as Faulkner acknowledges, Williams does not conceive of testimonial trust as purely predictive.¹² I'll discuss Williams's approach more fully in section III. My present point is merely that I agree with Faulkner that Williams's approach leaves out what

¹¹ See again note 2.

¹² On 174, Faulkner notes that Williams's solution to the problem of cooperation is not a reductive solution. I think this effectively concedes that Williams's conception of trust is not merely predictive. Though his account of testimonial reasons is reductive (as Faulkner rightly notes), Williams's conception of the assessment that that hearers must make of the speaker is not purely historical or predictive but typically refers to what I'm here characterizing as institutional elements.

Faulkner wants to emphasize: the second-personal nature of trust. So there is a burden on Faulkner's shoulders to explain how his approach differs from Williams's approach. But it does not appear that he can do that, given that his attempt to characterize the second-personal element in testimonial trust, (AT), admits cases that are purely institutional and as such lack the second-personal element.

One way to put this is to say that these cases of *merely* institutional trust do not involve an assurance. Of course, many cases of institutional trust do involve an assurance. Tellers and clerks often address you in the genuinely second-personal manner characteristic of a testimonial assurance. But sometimes they do not. Sometimes a teller or clerk merely asserts that p within your earshot, acknowledging that you are as a customer or client depending on her for the information whether p but without representing herself as aiming to do justice to your epistemic needs. I'll say much more about the relevant concept of an epistemic need in sub-section IIIb. My present point does not, however, depend on those details. It is easy to imagine (or remember) how a teller or clerk may manifest an institutional trustworthiness when she asserts that p without inviting your trust by manifesting appropriate responsiveness to your epistemic needs as addressee. Perhaps she isn't looking at you. Perhaps she looks at you but with a bored expression that conveys no interest whatsoever in any doxastic-deliberative circumstance that might inform your need for the information whether p. Still, it's her job to tell you the truth whether p, and there's no reason to doubt that she is incompetent in her job or that she fails to take it seriously. Though I cannot offer a full defense of the claim here, I find it plausible that such a speaker is fundamentally violating a constitutive illocutionary norm on testimonial tellings.¹³ By a 'fundamental' violation I mean what J. L. Austin called a 'misfire,' not merely an 'abuse.'¹⁴ It is an illocutionary abuse when S fails to do justice to A's epistemic needs in ways relevant to her testimonial telling. But there are cases in which S does not present herself as even attending to, or as trying to do justice to, A's epistemic needs. These are cases in which S's assertion manifests an aim of ignoring or dismissing A: she asserts that p within A's earshot but without addressing A in the way that Faulkner and I would agree is required

¹³ For a full defense of this claim, see my "Assurance and Warrant" (*Philosophers' Imprint*, forthcoming).

¹⁴ For the distinction between these two forms of illocutionary 'unhappiness,' see J. L. Austin, *How To Do Things With Words*, second edition (Cambridge: Harvard University Press, 1975), 13-17.

for her speech act to count as telling him that *p*. My argument in this sub-section is that the institutional context may ensure that the terms of (AT) are satisfied, simply because *S*'s failure to address *A* in a genuinely second-personal way is compatible with *S*'s being aware that *A* is depending on her for the truth, where providing such an auditor with the truth is required by *S*'s job, which she takes seriously. By (AT), such a case should – implausibly – count as a testimonial telling. And by Faulkner's broader argument, *A* may – absurdly – count as thereby acquiring a trust-based reason. These results are unacceptable by the lights of what I agree with Faulkner is the core insight of the assurance-theoretic approach to testimony.

Note well that neither the problem of internalism nor the problem of institutional trust is a problem for Williams's genealogical argument in the context of his reductionist approach to testimony. Williams developed and defended an influential version of internalism about practical reasons, and the species of trust that his genealogy vindicates as generally reasonable is primarily institutional, emphasizing not one-to-one trust relations but more general practices of inculcating and valuing the virtues of truthfulness. Each problem arises for Faulkner because of a key respect in which his position or argument differs from Williams's. Faulkner rejects Williams's reductionism in favor of a version of the assurance view on which affective trust can give reasons that are not grounded in the speaker's status as relevantly reliable. But (a) he assumes without defense an externalist hypothesis about aretaic reasons required to avoid the problem presented by psychopathic (or otherwise deeply uncaring) speakers, and (b) he does not seem to appreciate the possibility of a normatively robust institutional species of trust that is not yet affective, in his sense. These differences indicate key points of contrast between Williams's genealogical argument and Faulkner's redeployment of it. The argument of Chapter 6 for trust-based reasons has yet to receive its genealogical vindication. Faulkner's defense of an assurance view of testimonial reasons remains incomplete.

III. Can a genealogy vindicate second-personal reasons?

Does an assurance view of testimonial reasons need genealogical vindication? I'll now argue that an assurance view needs only as much genealogical vindication as Williams

gives his reductionist view of testimonial reasons. Faulkner argues that Williams's genealogical argument is in a key respect incompatible with Williams's own emphasis on truthfulness as a species of trustworthiness. I think that misses the point of a genealogical vindication of truthfulness as a species of trustworthiness. Faulkner is wrong to claim that the point of a genealogy is to vindicate the possibility of trust-based reasons. The point of a genealogy is to vindicate, if we can, the intelligibility and real-world applicability of the concepts of reliability and trustworthiness that we use to explain the possibility of reasons that employ those concepts. The key to defending an assurance view of testimonial reasons therefore lies in showing how the concept of a testimonial assurance is informed by concepts of reliability and trustworthiness that can receive such a vindication. We can defend an assurance view without positing trust-based reasons but not without a vindication of relevant concepts of reliability and trustworthiness.

IIIa. How an assurance view needs genealogical vindication

To see this, let's step back and ask what a genealogical argument might do for a theory of testimony. A genealogical argument addresses a question about obligation, in a way designed to engage skepticism about obligation. Does one really have an obligation to keep one's promises? Does one really have an obligation to tell the truth? Here are three answers: (a) yes, in all cases; (b) no, in no cases; (c) yes, in most (but perhaps not all) cases. A debunking genealogical argument explains why the fantasy of (a) obscures the truth of (b). A vindicating genealogical argument explains why the fantasy that the answer must be all-or-nothing – that is, (a) or (b) – obscures the truth of (c). The proponent of a vindicating genealogy, such as Williams, rejects (a) but aims to avoid (b) by vindicating (c).¹⁵

Why despair of (a)? Here again lies the problem of internalism. It's fine to say that one always has an obligation to keep a promise or to tell the truth, if what we mean is that not doing so amounts to a moral or illocutionary misstep. The question is whether one always has a reason to avoid such missteps. What if someone really did not care to step in line with the rest of us, regarding moral or illocutionary 'obligations' as fictions imposed

¹⁵ The debunking versus vindicating contrast doesn't map neatly onto the contrast between Nietzsche and Hume. As Williams argues (*Truth and Truthfulness*, 12-19, 37-38), there are vindicating elements in Nietzsche – though he aims to debunk morality as we understand it.

by the powerful? What if someone were determined to avoid ‘buying into’ these fictions? Well if the person really lacks any sound deliberative route to a motive to uphold her moral or illocutionary obligations, then it is not absurd to wonder if she really has any reason to uphold them. If we conclude no, then she would have the obligations in one sense but not in another. Her conduct would be correctly criticized if she failed to keep a promise or to tell the truth, but not because we must view her as acting contrary to one of her reasons. A vindicating genealogy aims to show that we can be confident that such cases are rare because the normal conditions in which we make promises or give testimony are shaped by obligation-supporting institutions that we all – even the above-described renegades – have reasons to promote. We all – even the renegades – have instrumental reasons to promote these institutions, and the institutions in turn give anyone who is not such a renegade obligating reasons to keep a promise or to tell the truth.

Can Faulkner’s assurance view make use of such an argument? Faulkner plausibly criticizes Williams for leaving out the second-personal element in trust. But in endorsing Williams’s account of how we have escaped the state of nature in the respect relevant to testimonial reasons, Faulkner overlooks the deep connection between Williams’s vindicating genealogy and his implicit reductionism about testimonial reasons. As Faulkner observes, Williams treats testimonial warrant reductively, so Williams doesn’t need a genealogical vindication to do what Faulkner needs it to do. Williams needs to make sense of what it would be for a speaker to be generally reliable in testifying; appeal to the intrinsic values of sincerity and accuracy explains the nature of this reliability and thus what one would have to ascertain in assessing a speaker as reliable. Without the vindicating genealogy, Williams could not explain the terms in which we assess each other, and in which we expect to be assessed, when we give and receive testimony. But Faulkner’s argument imposes a heavier explanatory burden. Anyone who views testimonial reasons as grounded directly in trust must explain how it is that every speaker whom one might encounter, including the renegades, has a reason to tell one the truth.

As we saw in section II, Faulkner’s genealogical argument does not deliver this result. We can now explain that lapse by noting that that is not a result that a genealogical argument is designed to deliver. Faulkner appears to assume that the genealogical

vindication works for an entire practice – vindicating an obligation to keep one’s promise or to tell the truth for ‘our’ practice though perhaps not for other practices. But, as we’ve seen, a genealogy such as Williams’s does not aim to vindicate promissory or illocutionary obligations for the arbitrary speaker in a given practice. Williams’s genealogical argument explains why a typical speaker – not the same as an arbitrary speaker – has an obligation to be truthful. His genealogy does not show that every person has a reason to internalize the virtues, merely that there is a general reason to. Any given person can be an exception – gaining the cooperative goods without pulling her weight. Therein, of course, we confront the problem of trust. It is in recognition of this problem that Williams insists on his reductive approach to testimony. It is the same recognition that leads me to insist on grounding testimonial reasons in relevant reliability.

Williams does not explicitly appeal to an internalist thesis in the course of making his genealogical argument, so one might question my claim as an interpretation of Williams’s text. I do think that an internalist challenge is clearly enough figuring in the background of his argument.¹⁶ But if one doubts that interpretive claim, there are explicit parts of Williams’s argument that entail the crucial difference from Faulkner that I’m emphasizing. Even if we set aside both of the problems discussed in section II – the problem of internalism, with its worry about psychopaths and the like, and the problem of institutional trust – we still get the problem that the addressee in a given case may not, as Williams puts it, *deserve* the truth. If you do not deserve the truth from a speaker, he argues, then that speaker has no obligation to tell you the truth. It is an important part of Williams’s genealogical argument that we not make a ‘fetish’ of assertion by regarding the illocutionary norm governing assertion as directly imposing an obligation of truthfulness – independently of whether the speaker’s addressee deserves the truth from her.¹⁷ It directly follows that Williams is committed to denying that testimony *eo ipso* imposes an obligation

¹⁶ Compare, for example, Williams’s endorsement of Nietzsche’s stance toward genealogy in *Truth and Truthfulness*, 12-19, 37-38, with his discussion of Nietzsche in “Nietzsche’s Minimal Moral Psychology,” (in *Making Sense of Humanity*, *op. cit.*); then compare the latter discussion with his treatment of internalism in other papers in that volume (cf. note 8 above).

¹⁷ *Truth and Truthfulness*, 100-110. Faulkner endorses Williams’s underlying point (without the metaphor of a ‘fetish’) at *Knowledge on Trust*, 180-1. I discuss what Faulkner’s endorsement shows about how he conceives epistemic needs in section IIIb below.

of truthfulness. From the addressee's perspective, this is not merely an epistemic issue, whether you can know that the speaker addressing you is obligated to tell you the truth. The issue is whether you can have a trust-based reason to believe what she tells you – that is, a reason grounded in a reason that she might have to tell you the truth simply because your trust manifests dependence on her for the truth. You can have no such reason because she can have no such reason. If she does have a reason to tell you the truth, that's in part because you deserve the truth from her, and your status as deserving the truth from her is not determined by, and in fact typically has little to do with, your attitude of trusting her.

Williams does not aim, then, to show that every speaker has a reason to be truthful. This is why he appeals to the addressee's assessment to explain how the addressee gets a testimonial reason. But Faulkner is correct to note that such an appeal to assessment is incompatible with an assurance view, and it is plausible that Williams's emphasis on assessment is inconsistent, as Faulkner argues, with his conception of himself as in the business of vindicating norms of trust. Assessing the speaker for truthfulness is incompatible with simply trusting her,¹⁸ and an emphasis on assessment erases the key distinction between believing a speaker and believing merely what she asserts. Since Faulkner wants to dispense with that need for assessment, he thinks he needs to show that every speaker has a reason to be truthful. And, as we've noted, it is hard to see how a genealogical argument could show that. This is the most fundamental problem confronting Faulkner's attempt to vindicate trust-based reasons. Is it fatal to any assurance-theoretic approach to testimonial reasons? How might an assurance view of testimonial reasons make better use of a genealogical argument?

IIIb. How an assurance view can ground reasons in reliability

A genealogical argument for an assurance view need not show that every speaker has a reason to be truthful, because an assurance view need not embrace the idea that a

¹⁸ In "Telling as Inviting to Trust" (*Philosophy and Phenomenological Research*, 70:3 (2005)) and elsewhere, I have argued that trust can be rational because it crucially rests on a counterfactual sensitivity to evidence of untrustworthiness: if there had been available evidence that the trusted is not worthy of your trust, you would not have trusted. Trust can thus be rational without the assessment for positive trustworthiness that Faulkner rightly regards as incompatible with trust.

testimonial reason can be grounded in trust. A testimonial reason cannot be grounded in trust, I've argued, because that would amount to illicit bootstrapping. The only way an assurance view can avoid illicit bootstrapping is to appeal to the speaker's status as relevantly reliable. But an emphasis on reliability need not be incompatible with an assurance view. Of course, if the testimonial reason rested entirely on the speaker's status as truth-conducively reliable, that would bypass the second-personal relation between speaker and addressee on which an assurance view focuses. But there is another dimension of reliability in play. To give a testimonial reason, a speaker must indeed be truth-conducively reliable. But she must also be appropriately responsive to how the quantity and quality of her evidence bears on her addressee's specific doxastic circumstances – that is, the context in which he would, if he trusted her, come to believe what she tells him. An emphasis on reliability can serve an assurance view if it targets the speaker's relation not only to the truth but also to her addressee's epistemic needs.

Like other proponents of the assurance view, Faulkner writes as if an addressee's epistemic needs are exhausted by his need to believe the truth. But that's too simple a conception of epistemic needs in general. When you wonder whether *p*, you're wondering not only whether it is true that *p* but whether you have sufficient epistemic warrant, in your actual circumstances, to believe that *p*. When you wonder whether to trust *S*'s testimony that *p*, you're wondering not only whether her assertion is true but whether it gives you, in your actual circumstances, sufficient warrant to believe that *p*. There are thus two burdens on the shoulders of any speaker who tells you that *p*: to tell you that *p* only when it is true that *p*, and to give you what would count, in your actual doxastic circumstances, as sufficient warrant for believing that *p*. As we'll see, we can easily imagine cases in which a speaker falls short of being a reliable testifier by doing the former but not the latter. That is, we can easily imagine how a speaker might tell you the truth without being appropriately responsive to your epistemic needs.

One key difference between my approach and Faulkner's thus derives from a difference in how we conceive the addressee's epistemic needs. Faulkner endorses the point that Williams emphasizes: that beyond telling what you believe to be the truth, sincerity requires doing so in a way that does justice to how much of the truth your

addressee needs and deserves from you.¹⁹ We thus get Williams's case in which S is opening A's mail but when confronted tells A only that *someone* is opening his mail, not that she herself is. As Faulkner observes, S is here violating a crucial norm of trust – in the aretaic sense, S is being untruthful – despite telling A nothing but the truth.²⁰ Testimonial trustworthiness and the norm of sincerity require being appropriately informative in this respect: you must tell your addressee as much of the truth as he needs and deserves from you. That 'needs and deserves' marks a complexity that Williams discusses at length: A may make an unjust request for information that he does not deserve.²¹ Whatever we say about that issue – the complexity that the norm of sincerity seems to rest on a norm of justice – it is undeniable that testimonial trustworthiness requires not merely that you tell your addressee the truth but that you more broadly do justice to his epistemic needs, giving him as much information as he deserves from you.

I accept that point, but my emphasis on the addressee's epistemic needs is different. When I note that a speaker might fail to be appropriately responsive to her addressee's epistemic needs, I mean that she might either tell him something that he doesn't, in his context, have sufficient epistemic reason to believe or fail to tell him something that he does in context have sufficient reason to believe (and that is supported by evidence available to S). It isn't at all difficult to come up with cases in which S has sufficient reason to believe that p in her context and does on that basis believe that p, but without being entitled to tell A that p – simply because it takes more evidence to count as sufficiently warranted to believe that p in his doxastic context than it does in hers. Allergy cases show this vividly: my pretty good evidence may suffice for me to believe that this bowl of snacks is nut-free but not for you to believe it, given your nut-allergy. I looked to see if the snacks contained nuts before I began to eat because I dislike the taste of nuts. "No nuts," I concluded, so I scooped up a handful. Now you arrive and ask me, "Does the bowl contain nuts?" I'm about to tell you that it does not contain nuts, since that's what I believe, but then I remember your allergy. "I can't say," I reply. Of course I *could* say, and with no impropriety – if I thought you merely shared my distaste for nuts. But your allergy

¹⁹ Cf. note 17 above.

²⁰ Williams's mail-opening case is Faulkner's case 32.

²¹ *Truth and Truthfulness*, 110-122.

imposes a higher standard on my telling. It would be a violation of illocutionary norms – of what Faulkner calls ‘norms of trust’; Austin called it an illocutionary ‘abuse’²² – to treat you as entitled to believe by the evidence that suffices to entitle me to believe.

Though one may balk at calling a violation of illocutionary norms an ‘insincerity’ when the speaker does believe what she asserts, we may naturally call it a failure to be properly informative.²³ To get a case of this sort, let’s flip the previous case around and imagine that you’re the speaker, with a severe nut allergy, and I’m your addressee, known merely to dislike nuts. You’ve been checking out the snack bowl and are confident enough for my needs but not for your own that the bowl does not contain any nuts. I ask you if the bowl contains nuts, and you tell me that it does not – despite not yourself believing what you assert. Are you insincere? Are you in any respect attempting to deceive me? Well, we may imagine that your refusal to explain why you are not yourself eating from the bowl manifests an attempt to deceive me about your allergy. But that’s a different matter, and its relevance to the present issue is merely that it helps distinguish your illocutionary obligations from other aspects of your relationship with a given interlocutor. (Why, anyway, should I have a right to know about your medical status?) The topic of our actual conversation is this bowl of snack food, and you aren’t attempting to deceive me about *that*. If without explanation of your allergy (and, again, you may well not owe me any explanation) you refrained from telling me what you believe me entitled to believe in answer to a question I’ve just asked you, that *would* amount to an illocutionary violation – to an Austinian abuse – since it would count as withholding information that I manifestly need and deserve and that I am now explicitly requesting from you. But that is not the case we’re imagining.

Illocutionary norms require telling your addressee what he needs and deserves to know given *his* epistemic standard, the standard against which he would form a belief. (Of course, you use *your evidence*, not his. Obviously, the idea is not that you should tell him only what he’s in position to figure out on his own!) In cases where your doxastic circumstances differ, the epistemic standard that governs your testimony may well differ

²² See note 14 above.

²³ In “Assertion, Sincerity, and Knowledge” (*Noûs*, forthcoming), I argue that *either* this is insincerity *or* sincerity does not mark an illocutionary norm.

from the standard against which you would form and retain a belief of your own. The norm governing testimony and the norm governing the speaker's belief are different norms that may give the speaker different directives: either 'tell that p but don't believe that p,' or 'believe that p but don't tell that p.' The norms can come apart in these ways because tellings are not announcements of what you believe but acts of assurance, wherein you put yourself under a norm defined from your addressee's perspective.

You must meet this norm – as we might put it, you must be *properly informative* – by being such that your addressee can rely on you as an informant. Flipping the testimonial relation around, you can get a reason to believe a speaker only if the speaker is a reliable informant – that is, reliable not merely as a speaker of truth but as a giver of such a context-sensitive assurance. The core of an assurance view of testimonial reasons lies in how testimonial trust – the species of trust that you manifest when you believe the speaker, not merely what she asserts – presumes that the speaker is reliable in both of these dimensions. Testimonial trust presumes that the speaker is not only truth-conducively but also, as we might put it, closure-conducively reliable – 'closure-conducively' because a speaker acts as a genuine informant only when her assertion that p permits her addressee to close a deliberation whether p, or to treat this matter of possible deliberation as closed, simply by believing her – that is, by accepting her assertion on trust. To accept an assertion on trust, you cannot be in the business of assessing the speaker for reliability. As Faulkner rightly emphasizes, when you assess for reliability you violate the terms of the trust relation, effectively stepping outside any relation of trust and believing (or not) on the basis of your independent assessment of the evidence. When you assess for reliability, you may wind up believing what the speaker asserts, but you don't believe the speaker: you don't accept what she says on her say-so.²⁴ So an assurance view needs to be anti-reductionist and emphasize that the speaker's reliability figures not in the addressee's assessment but as a defeating condition on whether an addressee who trusts without thus assessing counts as acquiring a testimonial reason. Anti-reductionists who are not assurance-theorists emphasize such a defeating condition on truth-conducive reliability, without realizing that there is a second

²⁴ I emphasized this point at length in "Telling as Inviting to Trust," where I also emphasized the need for reliability.

dimension of reliability in play. When we place equal emphasis on closure-conducive reliability, we open the door to an assurance view of testimonial reasons that avoids the problem of illicit bootstrapping.

Is closure-conducive reliability second-personal? From the perspective of an assurance view, the problem with appeals to truth-conducive reliability is that a speaker's reliability as an asserter of truth is not second personal: when she tells A that p, S's reliability in asserting the truth on the question whether p has nothing directly to do with her relation to A. But now what of S's reliability in giving A a reason that could serve to close A's deliberation whether p? Perhaps there's a guru with a website who can determine what would be a deliberatively sufficient reason for this or that abstractly characterized doxastic context, including a context like A's. Say S consults this website before telling A that p, and that the consultation is what makes S closure-conducively reliable in telling A that p. Does that show that closure-conducive reliability is not second-personal? It does not. It may be that the guru's website explains how S *became* closure-conducively reliable, but what her closure-conducive reliability *is* in this case directly involves her relation to A. The observation applies a point that I emphasized in a different application in section I: the nature of a reason is one thing, how one became able to give that reason another. In section I, I argued that a second-personal relation, mediated by affective trust, might explain how S comes to be truth-conducively reliable in addressing A, where truth-conducive reliability is not itself second-personal. Here I'm arguing that something that is not second-personal, consulting a third-party's website, might explain how S comes to be reliable in a way that is nonetheless second-personal. However it is produced, closure-conducive reliability involves a second-personal responsiveness, not mere possession of information or just any ability to 'get it right.' The guru's website may give S information that in turn enables A to count as closure-conducively reliable in addressing A. But what it is for S to be thus reliable makes essential reference to S's relation to A. One might object that what matters is S's relation to the doxastic context that A happens to be in, a relation that doesn't itself essentially involve A. But that context is defined by A's actual epistemic needs. Those needs can be characterized abstractly, but S's reliability consists in her responsiveness to those needs conceived *as A's* – that is, as the needs of this particular person to whom she

addresses her speech act.

One might still object that even if S must be responsive to A's epistemic needs, the species of reliability in question is reliability in addressing the needs of anyone in a doxastic context relevantly similar to A's. One might think that closure-conducive reliability must, like truth-conducive reliability, be implicitly general – a matter of getting a kind of thing right. A truth-conducively reliable speaker is reliable in getting the proposition right – that is, in asserting that p only when p really is true. And a closure-conducively reliable speaker is reliable in getting her interlocutor's doxastic context right – that is, in assuring A that p only when A really is entitled to close doxastic deliberation with the belief that p (or to treat that deliberative matter as closed). Viewed from this angle, S's reliability in getting her interlocutor's context right is no more second-personal than her reliability in asserting the truth.

A full reply to this objection would take us quickly into deep issues in epistemology and the philosophy of mind, since I would argue by drawing a comparison between testimonial reliance and a single subject's self-reliance when the subject forms a judgment that p . In a fuller treatment, I would argue that your self-relations when you form a judgment crucially include a relation isomorphic to the relation in which you stand to a speaker when you depend on her status not merely as truth-conducively but also as closure-conducively reliable.²⁵ When that relation is realized intrapersonally it becomes clear that it is not a relation between the subject and a source of information, conceived as a guide to 'getting it right.' Of course, you do rely on yourself to 'get it right' when you form a judgment, but that is a question of truth-conducive reliability: you rely on your epistemic faculties to give you the truth. The question of closure-conducive reliability here is question of self-concern: is your disposition to treat the doxastic question whether p as settled by your evidence a disposition that does justice to your epistemic needs? The question is not, of course, how you 'feel' about yourself but whether you are actually meeting those needs. You typically do trust yourself in this way, relying on your status as closure-conducively reliable every time you form a belief. Such self-trust is required even

²⁵ For part of this argument, see my "Reflection, Disagreement, and Context," *American Philosophical Quarterly* 49:2 (2012). For another part, see my "Judging as Inviting Self-Trust," in preparation.

in simple cases of perception. When you form a belief you treat the disposition that gives content to the self-trust – the disposition to treat your present evidence as sufficing to settle some matter that you might have deliberated, or deliberated further – as manifesting appropriate self-concern, by which I simply mean: as adequate to your actual epistemic needs. In any given case, you might not have done so; you might have treated the disposition as manifesting a mistake or confusion about your epistemic needs – or, in the most interesting cases, as manifesting a kind self-sabotage. This important dimension of your self-relations reveals a striking parallel between interpersonal and intrapersonal trust. In each dimension of trust, you rely on the trusted both truth-conducively and closure-conducively. The latter species of reliance is second-personal, or directed in the way of a second-personal relation, even in the intrapersonal dimension. Whenever you rely on a source or ‘mechanism’ of belief formation closure-conducively, you treat it as adequately concerned or ‘caring’ for – that is, as taking care of – your actual context-sensitive epistemic needs.²⁶

These remarks raise large issues. But setting aside that larger inquiry, we can note that the presumption of closure-conducive reliability figures at the core of testimonial trust. Whatever we say about the deeper role of trust in judgment, your trust in a speaker who tells you that *p* rests most fundamentally on your dependence on her as closure-conducively reliable. We can see this clearly when we ask what drives the distinction that assurance theorists emphasize, between believing the speaker and believing merely what she asserts. In each case, you depend on the speaker as truth-conducively reliable. The distinction emerges insofar as in the former case, but not in the latter, you depend on the speaker as closure-conducively reliable. To depend on a speaker as closure-conducively reliable is to grant her executive authority over your beliefs – to treat her assertion not as mere input to your independent deliberation whether *p* but as on its own settling whether *p*. When you merely believe what she asserts, you do not treat her speech act as settling whether *p*; but when you believe the speaker, you do. To treat her speech act as settling whether *p* is to treat the speaker as not only truth-conducively reliable but also as closure-conducively

²⁶ One might try to distinguish trustworthiness in testimony from trustworthiness in judgment as merely an application of the distinction between the other-regarding and the self-regarding. In “Assertion, Sincerity, and Knowledge” I argue that in this application that distinction is not so simple.

reliable on the question whether *p*. Unlike truth-conducive reliability, closure-conducive reliability goes right to the core of the testimonial truth relation.

IV. Comparing the two versions of the assurance view

Though I reject Faulkner's aim of vindicating trust-based reasons, I agree with his observations about the curiously self-reflexive nature of trust. When you trust a speaker, you expect her to acknowledge your dependence on her. When you invite such trust, by telling someone that *p*, you expect that he will acknowledge that you aren't merely asserting that *p* within his earshot but assuring him, whereby you expect him not merely to listen to you but to trust you. As Faulkner emphasizes, these are normative expectations: not merely something you expect him to do, but something you expect *of* him. I would explain the key normative expectation simply as the expectation that the speaker is closure-conducively reliable. The self-reflexive nature of testimonial trust derives from the fact that when you trust a speaker you are relying on her not merely to speak the truth but to do justice to your epistemic needs – that is, to your need to be warranted in closing deliberation, or in treating it as closed, in your particular circumstances. Faulkner is right that testimonial trust is self-reflexive but wrong about how.

I agree with Faulkner that testimonial tellings are structured by norms of trust, but I regard those norms as themselves structured by a mutual recognition between speaker and addressee. What makes the structure of recognition normatively engaging is the acknowledgment that the perspective on the other side is a perspective on, among other things, *one's own perspective*. The speaker imputes to the addressee a perspective from which he looks back on her perspective and trusts her. The addressee imputes to the speaker a perspective from which she looks back and either provides or fails to provide an adequate basis for that trust – a basis that would lie, as we've seen, in both truth-conducive and closure-conducive reliability. If his sensitivity to evidence of untrustworthiness is triggered, thereby alerting him to the absence of that basis, his refusal to trust constitutes a claim of right: 'Hands off, *you* don't get to define *my* epistemic needs.' If he does trust, he accepts the provision of care. Of course, he could instead believe what the speaker tells him without trusting her, by refusing the illocution but accepting the evidence it gives him.

But that would also amount to a refusal of the care. Putting the addressee's recognition as a question of the speaker's care and the speaker's recognition as a responsiveness to the addressee's needs captures the second-personal dynamic at stake in the speaker's status as closure-conducively reliable.

Here, once again, is where Faulkner and I part ways. From my alternative assurance view, I'd explain the divergence as follows: Faulkner overlooks how closure-conducive reliability informs the recognitional structure that in turn informs testimonial trust. If you fail to see how reliability can be closure-conducive as well as truth-conducive, and how both forms of reliability inform warranted belief, you'll have a hard time resisting the idea that any role for trust in testimonial reasons must posit that trust as the ground of the reasons. I've argued that the idea is an illusion. There is a key role for trust in a broadly reliabilist account of testimonial reasons: trust necessarily structures the recognitions whereby the speaker brings her reliability to bear on her addressee's epistemic needs. In the natural and compelling metaphor, she invites his trust. She invites his trust by presenting herself as reliable in a way that would give him a reason to believe what she asserts. One puzzle that informs Faulkner's book is how presenting oneself as reliable could amount to inviting trust. My solution is that presenting oneself as closure-conducively reliable amounts to providing, or at least to undertaking an illocutionary commitment to provide, a distinctively epistemic species of care.

Again we might make the point vivid with a creature construction. Consider three cases: (a) you encounter a sign warning you of danger further down the path you're traveling; (b) you encounter a man making assertions, though none addressed to you, about the danger down the path; (c) the man turns to you, looks you in the eye, and tells you of the danger. On my view, the man in (c) is doing something of an epistemic nature that the man in (b) simply is not doing: presenting himself as closure-conducively reliable about your epistemic needs. On Faulkner's view, the man in (c) does nothing of an epistemic nature that the man in (b) is not doing. Indeed, the only thing of an epistemic nature that either man does is already done by the sign in (a): each presents itself as truth-conducively reliable. On Faulkner's view, the extra epistemic work in (c) is done by your presumption, as addressee, that your trusting dependence gives this man a reason to be truth-conducively

reliable that the man in (b) does not possess. My argument against that move has been two-fold: (i) even if your trusting dependence does give the speaker a reason – and Faulkner hasn't argued that it would give someone with psychopathic tendencies a reason – that would merely be an observation about how the speaker came to be reliable; (ii) Faulkner has not explained how this mechanism differs from a merely institutional mechanism whereby the man in (b) might likewise have a reason to be truthful, since he may be aware of your dependence on him despite failing to address you. But we can now see a broader rationale for suspicion of Faulkner's entire approach. An assurance view of testimonial reasons ought to give an epistemic basis for distinguishing the speech act in (c) from the speech act in (b), since the former contains an assurance that the latter lacks. An assurance-theoretic approach to the epistemology of testimony ought to say how it makes an epistemic difference that the man in (c) looks you in the eye and invites your trust. I think this difference lies in how his invitation to trust purports to manifest closure-conducively reliability, a status that is itself directly epistemic. An assurance view of testimony needn't eschew reliability because a species of reliability serves to distinguish testimonial assurances.

Can we in general count on speakers to manifest the second-personal concern at the heart of closure-conducive reliability? We thus ask whether the norm informing the practice of testimonial assurance can be rationally sustained. It is not obvious that it can be rationally sustained. Perhaps Williams is right, and all we can rationally sustain is a practice that institutionalizes a need to assess speakers for reliability. If he's right, then the assurance view of testimony is fundamentally wrong. It would take a genealogical argument, or something like it, to decide this matter. It is, I think, undeniable that we treat others as if they could give us reasons grounded not merely in truth-conducive but also in closure-conducive reliability. It is undeniable that we often, if not always, treat each others' assertions as inviting trust. Is this a rationally sustainable practice? I have suggested that our interpersonal presumptions mirror intrapersonal presumptions that lie at the heart of individual judgment. This amounts to the suggestion that a debunking genealogy would overturn more than merely our testimonial presumptions. If we aren't testimonially trustworthy in this dimension, it's hard to see how we could be worthy of our

own trust even as we form our solitary judgments.²⁷ Again, the question isn't whether trust is risky. If there weren't a risk that the trusted would prove unworthy of your trust, your reliance wouldn't count as trust. The question is whether the practice and presumption are rationally sustainable – whether it makes sense in general to let our social transaction, and perhaps our self-transactions, rest with a counterfactual sensitivity to evidence of untrustworthiness, rather than working from a positive assessment of trustworthiness. No genealogist of trust has attempted such a vindication, and I'm not sure what it would take to pull one off. We need a better understanding of the risks of relying on a closure-conducive provision of care, and of how we might manage if we had to do without it.

Edward S. Hinchman

University of Wisconsin-Milwaukee

hinchman@uwm.edu

References

- Austin, J.L. (1975) *How To Do Things With Words*. Cambridge, MA: Harvard University Press.
- Faulkner, P. (2007) 'What is Wrong with Lying?' *Philosophy and Phenomenological Research* 75 (3): 535-557.
- Grice, P. (1991) 'Method in Philosophical Psychology', in *The Conception of Value*, Oxford: Oxford University Press.
- Hinchman, E.S. (2003) 'Trust and Diachronic Agency', *Noûs* 37 (1):25–51.
- _____. (2005) 'Telling as Inviting to Trust', *Philosophy and Phenomenological Research* 70 (3): 562–587.

²⁷ Compare this with my conclusion in "Trust and Diachronic Agency" (*Noûs* 37:1 (2003)): "If no one is trustworthy, no one has a reason to follow through either on his or her intentions or on others' advice. If no one is trustworthy, diachronic agency – the life of practical reason – is simply impossible" (45). We need a vindicating genealogy – or some kind of vindicating account – in the practical case as much as in the epistemic.

- Hinchman, E.S. (2012) 'Reflection, Disagreement, and Context', *American Philosophical Quarterly* 49 (2):95-111.
- _____. (forthcoming) 'Assurance and Warrant', *Philosophers' Imprint*.
- _____. (forthcoming) 'Assertion, Sincerity, and Knowledge', *Noûs*.
- _____. (in preparation) 'You May Rest Assured: A Theory of Normative Powers'.
- _____. (in preparation) 'Trust and Reasons'.
- _____. (in preparation) 'Judging as Inviting Self-Trust'.
- Williams, B. (1981). "Internal and External Reasons", in *Moral Luck*, Cambridge: Cambridge University Press, pp.101-113.
- Williams, B. (1995a) 'Internal Reasons and the Obscurity of Blame', in *Making Sense of Humanity and Other Philosophical Papers*. Cambridge: Cambridge University Press, pp. 35–45.
- Williams, B. (1995b) 'Nietzsche's Minimalist Moral Psychology', in *Making Sense of Humanity and Other Philosophical Papers*. Cambridge: Cambridge University Press, pp. 65–76.
- Williams, B. (2002) *Truth and Truthfulness: An Essay in Genealogy*. Princeton: Princeton University Press.