

REPLIES

Paul Faulkner

In the following sections I try to offer some reply to the contributions made to this *Abstracta* edition discussing *Knowledge on Trust*. For this opportunity of conversation I am extremely grateful and would like to thank both the contributors and the editors of *Abstracta*.

Reply to Longworth

I'd like to concentrate my reply to Longworth along three axes: my account of how we satisfy the demands of doxastic responsibility since this is the only place where I think Longworth does not get me exactly right; the epistemic externalism of my account since Longworth thinks that this is problematic; and the presumption I make about the nature of practical reason since this is the nub on which everything turns, as Longworth is correct to observe.

First, doxastic responsibility requires that the uptake of a piece of testimony be based on reasons. This condition, I think, is established by the problem of cooperation; it is established by the fact that there are a multitude of potential explanations of any given bit of testimony. As Longworth observes, I think that this tells against the non-reductive theories of McDowell and Burge. The reasons required for responsible uptake must be available *in advance* of that uptake, contrary to McDowell. And the origin of this demand for reasons implies that there can be no *pro tanto* entitlement, contrary to Burge. Given that I reject these options, Longworth observes that an alternative account of how uptake satisfies the demands of doxastic responsibility is required. And then notes that the account I offer suffers the same problem as McDowell's account; that is, even if trust can serve as evidence of a speaker's trustworthiness "by conditioning them to be trustworthy", it is "not clear how that evidence could ground the *initiation* of that trust." (n.10, p.24).

Questions of doxastic responsibility, I think, Longworth is right to observe, are to be settled by appeal only to internalist reasons; or reasons that are reflectively accessible. Doxastic responsibility tracks how a subject thinks about a situation, which in this case is the potential uptake of a piece of testimony. As such, it is not an audience's attitude of trust *qua* piece of evidence, if indeed it is that, which satisfies the demands of doxastic responsibility. Rather, it is the audience's attitude of trust itself, and specifically the presumptions that this attitude carries with it – how the audience thinks about the trust situation – that satisfies these demands. It does so because it identifies one of the multitude of potential explanations of a bit of testimony as the actual explanation; and because there is a potential explanatory connection between this act of selection and the speaker's trustworthiness. In this respect, the attitude of trust shares a property with something that is a bit of evidence for a speaker's trustworthiness. And it is its sharing this property that makes it fit for serving an epistemic role, which is rendering uptake doxastically responsible. So while it is true that an actual explanatory connection between trust and trustworthiness could not ground the *initiation* of trust, what grounds this initiation is not this connection but how the subject thinks about the trust situation in adopting the attitude of trust (with the presumptions it expresses).

It is worth elaborating this point somewhat. A full account of our reasons for uptake should, I think, distinguish between our situation, and the genealogically basic case characterized in the problem of cooperation. In our case, the demands of doxastic responsibility can be satisfied by all the empirical reasons that we have for thinking that a bit of testimony is true. These reasons are numerous. However we possess these reasons only insofar as uniformities exist within the testimonial domain, and this requires uniformities in intention. Our situation is one where there are *norms of trust* that determine these uniformities. These norms also determine a potential explanatory connection, by way of common cause, between trust and trustworthiness. (This explanation ordinarily follows the logic of the norm but it can also be mediated by the desire to avoid the punitive attitudes that are associated with norm transgression.) And these norms determine a speaker's expectation that an audience's reason for uptake be trust – or faith as Longworth calls it – and not empirical judgement. So our situation is one where we live in two 'worlds': we

have many empirical reasons to believe but are often called to believe on the basis of trust or faith; and where the tension between these reasons is also felt when a speaker asks to be given the benefit of the doubt, or when trust can otherwise take one out on a limb. The situation in the genealogically basic case is much bleaker. There is not yet the normative structure to ground either the presumptions of trust or the uniformities needed for the possession of empirical reasons. The problem of cooperation thereby confronted is then only resolved through an intrinsic valuation of trust and trustworthiness, where this evaluation is achieved in our case through our thinking about trust in the manner encoded in the norms of trust. This is the normative structure that then provides for both trust and empirical reasons. So the focus on trust is due to its role in the provision of reason, in setting up our world where empirical reasons are then bountiful and trust often, but not always, falls into the background.

Second, the idea that knowledge can be got on trust forces one to adopt an epistemic externalism about knowledge. Longworth argues for this consequence as follows. Call the case where a speaker knows what she says ‘the good case’, and the case which differs only in that the speaker does not so know, ‘the bad case’. The internalist reasons that support uptake in the good case – be they empirical reasons or trust – cannot outstrip the reasons available in the matching bad case. So if one can get to know something on the basis of testimony in the good case, that knowledge must be accounted for in externalist terms. This argument is good, I think; but the case for an externalist account of testimonial knowledge can be put more directly. Suppose, on the basis of testimony, I get to believe that spider silk is five times stronger than steel.¹ For this to be an item of knowledge, someone has to have done the science needed to establish it. However, the reason I possess for believing that spider silk is five times stronger than steel might be no more than that I trust some source. But if this is so, my belief is not based on my having done the science, so someone other than me must have done this science if I am to know this claim to be true. So if I do possess this knowledge, at least part of the explanation of my knowing this must refer to this science. This reference makes the account of my (testimonial) knowledge externalist.

¹ <http://www.guardian.co.uk/science/video/2012/jul/06/unravel-secret-spider-silk-video> . Accessed 9/7/12.

What is problematic here, according to Longworth, is that this externalism is at odds with the claim that questions of doxastic responsibility must be settled only by reference to internalist reasons.

It would be natural to expect facts about what one knows or doesn't know to play a role in determining whether or not one meets the demands of doxastic responsibility. ... However, with respect to the cases of knowledge that are of greatest interest to Faulkner—the normal cases in which one acquires knowledge from another—the expectation is not met. Some of the reasons that play an essential role in determining whether or not one knows are externalist reasons. And those externalist reasons are precluded from figuring in the determination of whether or not one's beliefs meet required standards of doxastic responsibility. On Faulkner's account, wherever one knows just on the basis of (responsibly) accepting what one is told, it is one's belief only, and not one's knowledge, that figures in determining whether one is, in further cases, being responsible in believing as one does. (p.25)

The solution, Longworth thinks, is to assert that the reasons made available by the good case are different to those made available by the bad case. However, this solution is McDowell's and adopting it means that the phenomenological objection I press on McDowell – against his account of our reasons for testimonial uptake – no longer stands.

The solution Longworth proposes, following McDowell, shows that epistemological externalism is not forced upon us by the suggested argument; that is, by the fact that we can imagine cases that are good and bad in the described way. However, epistemological externalism is independently plausible because of cases like the one stated, whereas the strategy Longworth proposes for avoiding externalism becomes less plausible when applied to such cases. Take the good case to be the actual case – wherein I get to know that spider's silk is five times stronger than steel – and the bad case to be one where the background scientific calculations and experiments were done badly (the result is still announced to be the same but the science does not establish it). The proposal is that the reasons that are reflectively accessible in the good case are different to those that are reflectively accessible in the bad case even when the only difference between the cases is that the scientific calculations were done badly in the latter case. Since this difference is not one that the

audience is cognisant of at any level – if he were, the bad case would not be matching – any sense of ‘reflectively accessible’ in play is implausibly etiolated. Moreover, and importantly for Longworth’s objection, it is impossible to see how this difference in what could be reflectively accessed could make a difference with respect to questions of doxastic responsibility such that the audience is rendered responsible when further uptake is based on the good case belief but not the bad case one. But if this is so, it cannot be a unique criticism of externalist accounts – and in particular my externalist account – that matters of external difference might equally fail to make such a difference. The problem is general to any account that allows testimonial knowledge of scientific claims (or claims that require a specific kind of warrant, which may not be possessed by an audience).

Third, the argument I give against Burge requires a specific conception of the nature of practical reason. As Longworth observes, it is the connection Burge presumes between belief and speech that I take issue with, and specifically the presumption of sincerity. One might trust a speaker to be sincere, and have reasons for thinking that the speaker is so, but there is no entitlement to presume sincerity established merely by the rationality of speech. This is because speaking is a practical activity and so responsive to the speaker’s interests, which need not be the same as the audience’s interests. Moreover, with respect to conversations as to the facts, the primary concern of any epistemology of testimony, if there is any presumption, it is that there will be such a divergence of interest. This is the root of the problem of cooperation, and it is this problem that Burge’s argument does not adequately address. Longworth’s acute observation is then that this problem presupposes a conception of practical reason – servicing an individual’s ends – to which one might object. In particular, there is an alternative conception according to which the primary task of practical reason “is the selection and attainment of ends that are *good*” (p.28). So its first task is working out what is moral and its second task is then servicing this end. If this account is adopted *and* it is proposed that there is “a *pro tanto* moral requirement to the effect that one should do one’s best to service an audience’s epistemic ends” (p.29), then Burge’s argument goes through because “[w]e would be entitled to faith, not only in others’ epistemic capacities, but also in the goodness of their wills.” (p.30)

This much, I think, is true: in trusting others, we do rely on the goodness of their wills; and in thinking about the trust situation as the trusting and trustworthy person do we take it that there is *pro tanto* moral requirement to the effect that one should do one's best to service an audience's epistemic ends. This requirement is effectively a particular expression of what I have termed the *norm of trustworthiness*. In thinking in terms of this norm, we think in a way that suggests the truth of the stronger conception of practical reason that Longworth, following Aristotle and Kant, proposes. However, the issue is whether this conception of practical reason is true, or merely a true description of how we think. The problem, I think, is that when we imagine the genealogically basic case what we imagine is a state of affairs where this conception of practical reason does not hold. The result is that we confront the problem of cooperation and what then matters is securing the motivations needed to solve this problem and so make society a possibility. Doing this involves finding some way of intrinsically valuing trusting and trustworthy behaviour. The way that we have secured these motivations is via a certain way of thinking about trust and trustworthiness, where this way of thinking is encoded in our norms of trust and articulated by the stated moral requirement. So *given* our set of thick ethical concepts, there can be the presumption of others' goodwill that suggests that the stronger conception of practical reason is true of us. But here my sympathies, as elsewhere, lie with Bernard Williams: we could have had a different set of thick ethical concepts. What matters is securing the motivations necessary for resolving the problem of cooperation, not how these motivations are secured. It then seems plausible that not all ways of securing these motivations would result in anything like our norm of trustworthiness or the stated moral requirement. But if this is the case, it would seem that the stronger conception of practical reason is no more than a truth about us, and the minimal more egoistic conception presumed in setting up the problem of cooperation is the basic one. We cannot, I have been at pains to argue, make sense of our lives in terms of this more basic conception. But this is to say something particular about us and not something general.

Reply to Keren

Keren though sympathetic to the shape of the theory proposed in *Knowledge on Trust* criticizes it in three ways. First, the argument from cooperation has a falsehood as a major premise. Second, the account I give of the psychology of trust cannot explain how belief can be based on trust. Third, I fail to explain how trust provides an *epistemic* reason for belief.

The basic testimonial situation, I argue in *Knowledge on Trust*, presents a problem of cooperation, which can be illustrated by reference to the Trust Game. The problem of cooperation for this game is: why should the investor cooperate and make an initial transfer? There are many answers that can be given to this question and Keren offers several, where that the investor believes the trustee to be kindly is one. What raises the question – why there is a ‘problem’ to be addressed – is that without the further information that explains cooperation, it seems that the game will default to the uncooperative outcome. This is because it would seem to be in the trustee’s interest to keep whatever monies he gains, and so in the investor’s interest not to trust him with any in the first place. The problem of cooperation for the Testimony Game is: why should an audience cooperate and believe what a speaker says? Again there are lots of answers that can be given, where again that the audience believes the speaker to be kindly is one. But the point is that there is a problem at this juncture that needs to be addressed, and this I argue by way of what Keren calls the *parity claim*: “the Testimony Game has the same payoff structure as the Trust Game”. (p.35) This premise Keren believes to be false.

The parity claim is false, Keren contends, because information, and indeed knowledge, “is a good whose consumption is non-rivalrous” (p.36); that is, unlike money, it can be shared without loss. However, it is not the similarity of goods that underlies the similarity in payoff structure or the parity claim. And though it is true that the conflict of interests in the Trust Game derives from the fact that both parties want the same good whose consumption is rivalrous, the conflict of interest in the Testimony Game rather derives from a *difference in interests*. An audience’s basic interest is acquiring a bit of information, and this because of some further need. A speaker’s basic interest is being believed because this is a way of exerting influence, which can similarly further some

particular need. A commitment to sincerity constrains this basic interest, so it is better for a speaker not to be so constrained, and merely to tell the truth when this suits. But this is to be non-cooperative. So while any given conversation will be shaped in determinate ways that can make it cooperative, the basic shape of our conversational engagement is such that without further information that explains cooperation, the game will seem to have the non-cooperative outcome as its default.

Keren's second problem concerns the central claim that testimonial uptake can be based upon affective trust; that trust in this sense can be a reason for belief. Whether trust provides an *epistemic* reason is Keren's third problem, which I will come to shortly, but for now the issue is whether affective trust could be a reason in a "causal-motivational" sense. (p.40) Were it the case that affective trust implied the *belief* that the trusted is trustworthy, there would be no issue, but as I characterize it affective trust merely implies the *presumption* that this is so and, Keren asks, "why should we expect this presumption to result in their believing that the speaker's testimony is true, and not merely in their *presuming* that the testimony is true?" (p.41) A presumption could underwrite belief if it were a presumptive right to believe; that is, if the presumption of trustworthiness were the presumptive right to believe this. However, Keren rightly observes that this is not what is intended, so then he challenges that if the presumption is little more than an assumption made for the sake of argument – as when a judge presumes the defendant innocent – the most it could suffice for would be a state of acceptance. But what speakers expect when they expect to be trusted is to be believed, and trust must suffice for belief if it is to credit this fact and find its place in an epistemology of testimony.

On the other side it is, I think, an important feature of (affective) trust that it does not require a belief in trustworthiness. This is needed if we are able, as I think we are, to choose to trust and to extend trust in the face of doubt. The challenge is then: how can anything short of belief ground testimonial uptake? The idea that trust entails a presumption of trustworthiness was meant to address this challenge, Keren's contention is that it does not. In reply the best I can do is say why I think it does. This presumption is not comparable to that made by the judge, nor is it comparable to an assumption made for the sake of argument because such assumptions are singular and isolated, whereas in the case

of trust the presumption of trustworthiness follows from a set of further acceptances, where these constitute the background set of attitudes necessary for adopting the attitude of trust in the first place. Affective trust is a ‘thick’ ethical concept so that using it is part of living in a particular social world, which is a world where the truth of the set of things accepted in trusting someone is a commonplace. So the presumption is not isolated but a consequence of other things accepted, and these other things, even though they do not need to be believed, will be recognized as things whose truth is commonplace. As such the presumption made in trust might be characterized, if this is any clearer, as a presumptive belief: it is a taking that things are a certain way, but it is a taking that is based on an optimistic view of a world (in that it is both optimistic and expresses a view of the world) rather than a taking based on evidence. So it would be wrong to describe it as a belief, but it has the substance of belief rather than acceptance, which need not be similarly coherent, and it is this that enables it to be the “causal-motivational” basis for belief. This might then be put the other way around: what reflection on cases of trust shows is that we need to refer to an attitude that is belief-like but that falls short of belief and is not constrained in the way that belief is constrained. This is the attitude that I have called a *presumption*. Similarly, that we need such a type of attitude can also be argued by reflection on other cases; for instance Holton argues that we can only make sense of certain cases of intending if we have a notion of *partial belief*.²

The third problem Keren finds is with my account of how trust can be an epistemic reason for belief. It is part of my account of how trust can be such a reason that it can be potential evidence. That is, the fact that A trusts S for the truth as to whether p, when S tells A p, can be potential evidence for p; it is so when S is trustworthy. Keren then observes, “even if A believes that p because she trusts S, and her trusting S is potential evidence for p, it is not at all clear that it would be correct to say of her belief that it is based on this evidence” (p.44). However, potential evidence for p can justify A believing that p if and only if A believes that p on the basis of this evidence. Since this condition is not satisfied

² See Holton (2009), pp. 29-34. Though the presumption of trustworthiness is not a partial belief as a key feature of this presumption is that the trusted party’s being untrustworthy is not a “live possibility”.

when testimonial uptake is based on trust, this account cannot explain how it is that trust provides an epistemic reason for belief in these cases.

If the warrant that trust provides comes by way of trust being a piece of evidence, then this is correct: the audience's belief needs to be based on trust *qua* evidence, and it is not so. But the warrant that trust provides does not come this way, rather trust provides a way in which an audience's belief can be based on the extended body of warrant possessed by the speaker (and testimonial chain) and it is this that warrants the audience's belief. However, in order to play this role of basing the audience's belief on this extended body of warrant, trust needs to make testimonial uptake epistemically reasonable. The argument that Keren considers is then addressed to this question: not how does trust warrant but can trust be considered to be an epistemic reason? The argument is this: trust provides a "causal-motivational" reason in that when A confronts S's telling him p, A's trusting S for the truth makes it subjectively probable for A that p; and this reason is an epistemic reason because A's trust could be evidence for p. Thus, it is not that trust provides the reason of evidence but that the reason that trust provides is an epistemic reason because trust could figure in an argument for the truth of that which it is a reason for, because it could be a piece of potential evidence.

Reply to Hinchman

There are, I think, essentially three criticisms that Hinchman directs against the theory of testimony put forward in *Knowledge on Trust*. The account I offer of trust-based reasons is implausible. The solution I offer to the problem of cooperation is incoherent. And the analysis I give of affective trust fails of sufficiency.

First, affective trust, as I conceive it, provides a reason for belief because it is essential to trusting that one presume the trusted will behave as one expects; that is, and in short, because one presumes the trusted to be trustworthy. Take the case of an audience A trusting a speaker S to tell the truth. As Hinchman observes, it is consistent with A trusting S in this respect that S be quite untrustworthy. From this observation Hinchman then infers that "trust cannot itself provide a reason to believe the trusted, independently of the speaker's status as reliable". (p.56) And since trustworthiness implies reliability – merely

saying what one believes is not enough – the inference is: the trusted has a reason to believe *only if* the presumptions made in trust are true. This condition is necessary because “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.” (p.55). And this, Hinchman thinks, is implausible.

Now the truth of what is presumed in trust matters. It matters to A’s epistemic standing: whether trust results in knowledge or warranted belief or not, will hinge on whether S is trustworthy. And it matters to the status of A’s reason for belief: were A’s presumptions true, that A’s trusts would be a *good* reason for belief in the sense that it could be used by a third-party in a justification for p, when this is what S tells A; it would be a bit of potential evidence for p. But the truth of what is presumed in trust is *not necessary* for trust being a reason for belief for A. Rather, trust is a reason because in trusting S for the truth, A’s attitude makes it subjectively probable for A that p, when this is what S tells A, and our social world is such that A’s trust *can be* a good reason in the sense noted. In allowing that trust can provide a reason for belief *and* allowing that trust can be chosen, it follows that A can bootstrap his way into the possession of a reason. But this conclusion is not implausible given its presuppositions, namely that this possibility requires as a background the social institutions of trusting and telling.

Second, the solution I offer to the problem of cooperation, Hinchman argues, is incoherent in that it requires both an internalist and an externalist conception of practical reasons. What is presumed in trust is that the trusted will see things in a certain way, and as a result have certain motivations. In trusting S for the truth as to whether p, A presumes that S will be motivated to tell A the truth by A’s need for it. Here reasons are conceived in internalist terms. However, and here Hinchman asserts the necessary condition just disputed: A possesses a trust-based reason *only if* S does in fact have this reason to tell A the truth. “Faulkner’s strategy”, Hinchman continues, “is to turn that necessary condition on trust-based testimonial reasons into a sufficient condition.” (p.57). Genealogy accomplishes this transformation: if we have escaped the State of Nature, social conditions must be in place such that *anyone* in S’s position would have a reason to tell A the truth. Thus trust is sufficient for reasonable belief and the problem of cooperation is resolved.

However, this resolution presupposes an externalist conception of reasons – it presupposes that S has a reason merely because of the position that S is in. Given that both internalist and externalist conceptions of reason are presumed, the result is incoherence.

It is true that it is the availability of trust-based reasons that resolves the problem of cooperation, and that what makes these reasons available is that certain social conditions are in place, but the genealogy does not thereby establish their availability through showing that *everyone* has the reason to be truthful that it is presumed the trusted has in trust. As Hinchman observes, it is not true that everyone would care about the trusting party's needs, and there are those, who I follow Hinchman in calling 'psychopaths', who could not be brought to care. However, the argument is not that even psychopaths have a reason – an externalist reason – because what makes trust-based reasons available is not that the putative condition on their possession is universally satisfied; rather, it is simply that the social conditions are such that trust generally, though not always, is an option. Its being an option is what the genealogy establishes, and it is not always an option because one might, for instance, believe the other party to be a psychopath.

If there are norms of trust, then it is possible to go wrong in one's thinking about the trust situation. And in terms of these norms the psychopath does go wrong. However, to suppose that any failure to think about the trust situation in the prescribed way is a failure to be moved by *reason* requires supposing that rational deliberation in conformity with the norms of trust is just a matter of reasoning as the fully rational person would. Here I agree with Williams: accusations of irrationality are mere "bluff".³ It seems reasonable, for instance, to think about the trust situation in terms of the logic of tit-for-tat even if to operate in accordance with this strategy is to demonstrate little trust (in the affective sense). We think about things in terms of the norm, this is how we have resolved the problem of cooperation confronted in the State of Nature, but the psychopath has slipped through the net. Thus a short response to this criticism is simply that the only conception of practical reason in play is internalist (and Hinchman only thinks otherwise because he takes the disputed necessary condition to hold).

³ Williams (1980), p.111.

Third, the analysis I give of affective trust fails of sufficiency because it aims to capture the distinctively second-personal character of trust as it can be found in the testimonial situation and yet the conditions proposed can be satisfied by trust as it is found in *institutional settings*, which are not second personal in any way. Hinchman gives the example of A (an account holder) trusting S (a bank teller) to transfer some monies. In this case, the dependence condition is satisfied: A depends on S transferring these monies. And the expectation condition is satisfied: A expects S to transfer the monies precisely because this is what he, A, depends on S doing. Moreover, this expectation is normative: in the given situation, A thinks that this is what S *ought* to do and this is what A *expects of* S. However, the normativity here is not second-personal but is mediated by the institutional setting. It rests on no more than the thought that S should do her job. Since the conditions on affective trust are met but the trust is not second-personal, these conditions do not suffice to give an account of this dimension of trust.

This is an interesting case but, on balance, I think the best response is that it is not a case of affective trust, and so not a counterexample to the use of this notion of trust in giving a theory of testimony. It is *not* that A (the account holder) expects S (the bank teller) to respond to *his*, A's, dependence but that A expects S to respond to the demands of account holders generally. This is what is involved in S doing her job: *not* responding to the needs of particular individuals but responding to the demands of individuals *qua* account holders. So what the case combines is predictive trust with normativity. It combines a belief that someone will act in certain ways with the belief that acting in these ways is prescribed. Here it is worth stressing that our reason for testimonial uptake can often be such a straightforward assessment of the likelihood truth and that such an assessment can be, as in this case, straightforwardly based on assessments of motivation, which in turn can revolve around judgements about social and institutional norms. Testimonial trust does not always have the second-personal character that I hope to capture by reference to affective trust; it can be simply predictive. The philosophical point is then that the success of this kind of empirical judgement presupposes large-scale consistencies in motivations, which are secured by the norms of trust (even if these norms do not shape motivations universally and there remain psychopaths). So, in the genealogical story, assurance comes before evidence.

Reply to Hawley

Like Keren, Hawley thinks that the argument from cooperation is based upon a falsehood. The falsehood is that the Testimony Game is analogous to the Trust Game. Hawley argues that it is not, and what follows from this is that I have failed to establish that the uptake of testimony requires a reason if it is to be reasonable. That is, what follows is that the argument from cooperation fails to conclude. In reply I would argue that the analogy between the Testimony Game and the Trust Game amounts to a specific claim: these games have the same pay-off structure. This I argue on the basis of assumptions about the basic interest each player has in the game, where these assumptions are defeasible and their defeat would then constitute a specific reason for expecting a cooperative outcome. In arguing that the games are disanalogous Hawley does not question these assumptions or that the games thereby have the same pay-off structure. Rather, she argues that the games are disanalogous on the basis of a number of differences between them. Of course, that the games are analogous in the sense specified is consistent with their being different in other ways. So the question is: *are the differences Hawley highlights such that the similarity between the games with respect to pay-off structure cannot be used as a premise in the argument from cooperation?* In what follows I argue that the differences Hawley highlights are not as significant as she suggests, and that these differences do not undermine the argument from cooperation.

There are, Hawley argues, four differences between the Trust Game and the Testimony Game: there is a temporal difference in the order of play; there is a difference in the strategies that can be adopted; there is a difference in the currency of rewards; and the roles played in each game differ in their stability. However, Hawley only uses the first two of these differences in arguing that the games are disanalogous in a way that undermines the argument from cooperation, so I focus only on these two.

First, there is a temporal difference in the order of play. In the Trust Game the player who must decide whether to trust (the investor) makes the first move in making a transfer of monies or not, whereas in the Testimony Game the player who must decide whether to trust (the audience) makes the second move in that this decision is a response to a speaker's assertion of something. This undermines the argument from cooperation,

Hawley claims, because the very fact that a speaker has made an assertion can be taken as a reason for thinking that the speaker is cooperating.

That there is a difference here, I think, must be acknowledged, but it is less than it seems and does not undermine the argument from cooperation. In the Testimony Game, a speaker's assertion purports to address an audience's need for information, say as to whether *p*, and can, as such, be regarded as a response to this recognized need. That is, the audience's need is regarded by the speaker as posing the question, *p*? Were this question vocalized by the audience, the difference Hawley highlights would vanish. But then does this vocalization really make such a difference? That a speaker's assertion comes first undermines the argument from cooperation only if this assertion in itself constitutes a reason or thinking that the speaker is being cooperative. Later in *Knowledge on Trust* I offer one account of how this might be so in outlining the assurance view, and endorse something like this view. But the problem for this view, and indeed for any suggestion that a speaker's assertion itself offers a reason, is that multiple explanations can be given of utterance, including, in particular, explanations that do not start from a speaker having any informative intention. This account denies the assumption about interest made in setting up the Testimony Game, but if this assumption is accepted, and Hawley does not question it, then the temporal order of play makes no difference.

Second, there is a difference in the strategies that can be adopted by the players in the two games. In the Trust Game the player who must decide to trust (the investor) can adopt the strategy of "never take a risk" – or "never transfer cash" – and so accumulate capital, whereas in the Testimony Game if the player who must decide to trust (the audience) were to adopt the strategy of "never take a risk" – which in this case amounts to "never give credence" – the result would be unremitting ignorance. This undermines the argument from cooperation, Hawley suggests, because it implies that uptake must be default entitled.

Again I think that there is a difference here but it is less than it seems and it does not undermine the argument from cooperation. It is true that the non-cooperative strategy is not a bad strategy for the investor in the Trust Game in that by following it he might accumulate cash. But it is still a worse strategy than cooperation when this is successful; the

investor might get rich through taking no risks but could get much richer by taking risks. And it is true that the non-cooperative strategy is not a good strategy for the audience in the Testimony Game in that by following it he is simply left in a position of ignorance. But it is still a better strategy than cooperation when this is unsuccessful; the audience might remain ignorant through taking no risks but could end up in error by taking risks. It is better to give no credence than give credence wrongly just as it is better to lose no money than lose all one has. So the difference is not so great. And rather than undermine the argument from cooperation all that is implied by the fact that we are not ignorant is that the problem of cooperation is empirically resolvable. Our risk taking is often, even ordinarily, well judged in that we tend, by and large, to give credence when credence is due. What I then argue in *Knowledge on Trust* is that this (reductive) solution to the problem of cooperation is limited in various ways; in particular, it is limited in that it accepts that there is a problem confronted. A more radical solution involves denying the background assumption, which leads to the problem, that competing explanations of utterances are equal. But there is no short step from noting the sceptical nature of the problem of cooperation to such a denial. In particular, our possession of an entitlement to believe testimony would require some basic connection between testimony and truth but if the assumptions about interest that generate the problem are correct any such connection is contingent.

Reply to Graham

I would like to make three points in reply to Graham. First, I would like to try and clarify the psychology of trust and the role that trust plays in the epistemology I propose because Graham does not get these quite right and his not doing so feeds into his first, and principle, objection. And then, second and third, I will try and respond to Graham's two objections: that there can be no epistemological role for trust once social norms of trust are recognized; and that my theory of testimony is refuted by the fact that children can acquire testimonial knowledge.

As Graham reports my account, the psychology of trust involves the trusting party having a number of beliefs and making a number of presumptions. For the case where A trusts S to ϕ , trust involves A believing that S can recognise his, A's, dependence and

believing that S can recognise that his, A's, attitude towards this dependence is one of trust. And trust involves A presuming (i) "that if S recognizes A's expectation that S should prove informative, then other things being equal, S will prove informative for this reason"; and so presuming (ii) "that the trusted will prove trustworthy". (p.99) This description gets the belief side right, and trust does involve presumption (ii), but this presumption does not rest on (i), which is no part of trust. What is presumed is that S will be moved by the *fact of* A's dependence, which A believes S recognises, not by any normative expectation that A might have. To see this suppose that S is trust-responsive in Pettit's sense and is particularly sensitive to A's opinion and desirous to avoid A's resentment. If this were the sole reason that S ϕ -ed, then presumption (i) would be true and yet presumption (ii) false because S would not have the motivations characteristic of the trustworthy person. So (i) could not be the basis for concluding (ii) and, indeed, it is not so.

This is then relevant to the main criticism Graham articulates because he finds two roles for trust in my theory – a rationalizing role and a motivating role – and his criticism is that once social norms are recognised both roles are epistemologically otiose. The motivating role is meant to be that just described. The trusting party A is meant to presume (i) and this presumption is meant to come out true such that S is motivated to ϕ by A's trusting S to do this. This can be the case: S can be trust-responsive, and so A's trust can give S an instrumental reason to behave as A expects. However, this reason is not the reason that would move the trustworthy person, who would be moved simply by A's need, rather than by A's attitudes. So trust does not play the motivating role Graham describes. So it can be no criticism that this role is epistemologically superfluous. However, trust is meant to play a rationalizing role, so Graham's main criticism still needs to be addressed. I turn to this now.

In *Knowledge on Trust* I argue that the problem of cooperation establishes the requirement that the uptake of any piece of testimony needs to be rationally supported. And I argue that the attitude of (affective) trust can play this role: it can be reasonable to base belief on trust. Graham's criticism is that the existence of a social norm of truth-telling establishes an entitlement to believe testimony. So if there is such a norm, there can be no requirement that testimonial uptake be supported in every case. So there is no essential

rationalizing role for trust to play. Of course, our entitlement to testimony might be defeated in any particular case, so the uptake of a piece of testimony might require supporting reasons and trust could then play a rationalizing role. But this role is non-essential and what matters is the fact that speakers are by and large reliable, where their being so is established by the social norm of truth-telling; “[s]peaker reliability, grounded in pro-social preferences, is enough.” (p.115)

Suppose it were established that there is a universal norm of truth-telling. (For instance by appeal, following Longworth’s suggestion, to a stronger conception of practical reason than that which underpins the problem of cooperation.) If this were established, then there would be a basis for arguing that we have an entitlement to believe testimony. However, the claim that there is a social norm of (informative) truth-telling is not this universal claim. It is the claim that there is a *social* norm, a norm that operates in this vicinity or for people like us. By contrast, the requirement established by the problem of cooperation is meant to be universal. This is because it is an essential feature of communication that speaker and audience have different basic interests in the communicative exchange. So any instance of testimonial uptake must be rationally supported. What the institution of social norms of trust then establishes is one way in which this requirement can be met: in this locality, uptake can be based on trust. It might then be claimed that what is thereby established is that *we* have an entitlement to believe testimony, where the referent of ‘we’ ranges over the domain of the social norm. However, this would be an entitlement in name only because it would be grounded on a particular empirical claim, and so, in fact, would be a reductive justification with a major premise of the form ‘speakers of type X tend to be reliable’. This gives a straight solution to the problem of cooperation; but this is not the solution proposed in *Knowledge on Trust*, which is rather that where there are social norms of trust interlocutors can operate with a series of presumptions that make trust non-problematic. This dissolves the problem not, as Graham suggests, through adding an additional set of pro-social preferences, but through denying the conception of rationality presupposed in setting up the problem of cooperation by widening the class of things that can act as a reason for action. Thus reasons are not limited to complexes of belief and desire but can include the trustworthy speaker’s perception of

the audience's informational need. However, this dissolution travels only as widely as the norms of trust are internalized, and these are merely local distances.

The second objection put by Graham is that facts about children refute the Trust Theory of testimony, and my argument otherwise fails. Specifically the following facts form an inconsistent set.

1. An audience A is warranted in the uptake of testimony to p if and only if it is reasonable in the light of A's other attitudes to believe that p.
2. The attitudes that can render the uptake of trust reasonable are either A's beliefs, or A's trust.
3. Children do not possess enough background of belief to render uptake reasonable.
4. Children do not have the sophistication needed for adopting an attitude of (affective) trust.
5. "Nonetheless, children learn (come to know) a great deal by believing what their parents, relatives, teachers, caregivers and even strangers tell them."
(p.112)

In *Knowledge on Trust* I considered this argument and proposed that neither 3 nor 5 were clear cut. At some point, children gain sufficient sophistication and belief for uptake to be reasonable in the light of it, and at some point children come to know things on the basis of testimony. But at no point is it clear, as the argument requires it to be, that children lack sufficient sophistication and belief for uptake to be reasonable but nevertheless come to know things on the basis of testimony. Graham rejects this response to the argument. First, what I cited as evidence of the sophistication of children, namely their sensitivity to a speaker's communicative intentions, is not sufficient for the possession of a reason. Second, denying that children acquire testimonial knowledge amounts to just "denying the data".
(p.113)

In response, what I cited as evidence that children are more sophisticated than proposers of the argument would assert in asserting 3 was in fact an article by Paul Harris. I now quote from this.

I resist the claim that credulity is strongest in early childhood. Young children are endowed with two protective devices. First, they are alert to a speaker's intention – they do not systematically confuse fictional and factual claims. Second, for comprehension to proceed smoothly, any new piece of testimony needs to be consistent with, and integrated into, what is already known about the topic in question. Hence we may assume that young children will find it difficult to accept and integrate what they regard as anomalous statements. Indeed, to the extent that they know less than adults, they may ultimately be less credulous. Their impoverished knowledge base will make the integrative process slower or more taxing. Stated simply, children's ignorance may often safeguard them from misplaced trust.⁴

If Harris is correct, then testimonial uptake can be reasonable even for young children: testimony which is believed is believed because it fits with the child's background of belief. Now I do not want to assert this consequent, but I think that this possibility is enough to establish that 3 is not clear cut, as claimed. What, then, of the 'data'? It cannot be that children learn things from testimony. This is certainly true but it is so because it can be interpreted as no more than the claim that children get to form true beliefs on the basis of testimony. What is needed is the claim that children acquire *knowledge* in some stronger sense. And there is just no clear cut data here as to when this happens in part because of disagreements as what knowledge in this stronger sense amounts to. However, if one takes knowledge to be a certain 'standing in the space of reasons', as I do but Graham doesn't, then it would be natural to regard the shift from 3 being true to being false that occurs with a child gaining belief and sophistication as tracking the shift from 5 being false to being true. But if this is the case, there is no point at which 3 and 5 are both true, as the argument asserts.

⁴ Harris (2002), p.331.

Paul Faulkner

University of Sheffield

paul.faulkner@sheffield.ac.uk

References

- Harris, P.L. (2002) 'Checking our sources: the origins of trust in testimony'. *Studies in the History and Philosophy of Science*, 33:315-333.
- Holton, Richard (2009) *Willing, Wanting, Waiting*. Oxford: Clarendon Press.
- Williams, B. (1980) 'Internal and External Reasons'. In *Moral Luck*, edited by B. Williams. Cambridge: Cambridge University Press.