

KNOWLEDGE ON AFFECTIVE TRUST

Arnon Keren

Epistemologists of testimony widely agree on the fact that our reliance on other people's testimony is extensive. However, they have not always paid enough attention to the fact that our reliance on testimony is not only extensive, but also varied. That is, there is more than one way in which we can rely on a speaker's testimony to form a belief: Sometimes we treat a person's testimony that *p* as we treat any other piece of evidence. At other times, we just trust the speaker for the truth of what she says. Arguably, the latter is the normal way of forming a belief on the basis of other people's testimony. And while Paul Faulkner's important book presents an account of how we can obtain knowledge in both these ways, its focus, as its title suggests, is on knowledge obtained through trust.

Faulkner is one of a number of epistemologists who have attempted in recent years to explain how we can obtain knowledge by trusting others (Faulkner 2007; Fricker 2006; Hinchman 2005; Keren 2007; McMyler 2011; Moran 2005). It is hard to deny that knowledge can be obtained in this way, and Faulkner's attempt to explain how this can be possible is both original and sophisticated. It thus constitutes a significant contribution both to the epistemology of testimony, and to the philosophical study of trust.

Faulkner's book is divided into three major parts. The first (chapter 1) presents a problem—the problem of cooperation—and draws a conclusion from the analysis of the problem that underlies much of the discussion throughout the book: This is the claim that it would be reasonable for an audience to trust a speaker only if she has some positive reason to think that the speaker is trustworthy or cooperative.

This claim is at the basis of some of Faulkner's central arguments in the second part of the book (chapters 2-5) in which he discusses and rejects reductive and non-reductive approaches to testimony, and defends instead a distinct theoretical view in the epistemology of testimony: one that combines some of the more plausible elements of both, while rejecting the less plausible ones. Thus, with the non-reductionist he claims that testimonial

knowledge is transmitted knowledge. Accordingly, a necessary condition for obtaining *testimonial* knowledge by trusting a speaker is that the speaker herself knows. And it is in virtue of its being transmitted knowledge that testimonial knowledge is a distinctive kind of knowledge. At the same time he agrees with the reductionist that one can obtain knowledge by trusting a speaker only if one has reasons, epistemic reasons, for trusting her. Thus he rejects the non-reductionist claim that we have a default entitlement to accept the testimony of a speaker.

The third part of the book (chapters 6 and 7) proposes a positive solution to the problem of cooperation, explaining how we can have reasons to trust a speaker. Faulkner adopts a non-doxastic account of a form of trust—*affective trust*—according to which, to trust a person, one need not believe that the person is trustworthy. Moreover, according to Faulkner's account, to be rational in trusting a person, one need not have evidence that the trusted person is trustworthy. While Faulkner's account of trust shares important features with other accounts of trust found in the literature, e.g., in Holton (1994), Faulkner's account is much more developed, and his explanation of how trust can be rational, and how beliefs formed through trust can be epistemically warranted, is highly original. Ultimately, he claims, that affective trust is rationally self supporting, and that it is this trust itself which provides an audience with an epistemic reason for her belief in the testimony of the speaker.

There is a lot that is going on in this rich book. In my discussion, I will therefore only address some key aspects of his study of knowledge formed through trust, focusing on those aspects with which I disagree. I will first raise some worries about Faulkner's analysis of the problem of cooperation. Ignoring much of Faulkner's discussion of reductive and non-reductive position, with which I largely agree, I will then turn to a discussion of Faulkner's suggested solution of the problem. As I will argue, this latter part of the book raises a number of fundamental worries. It is therefore doubtful whether Faulkner has presented us with an adequate explanation of how knowledge is obtained through trust.

The Problem of Cooperation

Faulkner opens his book by presenting what he calls "the problem of cooperation". From his analysis of the problem, he concludes that it would be reasonable for an audience *A* to trust a speaker *S*, only if *A* has some positive reason to think that *S* is trustworthy or cooperative. This claim is at the basis of Faulkner's arguments against both reductive and non-reductive approaches to testimony. And it lies at the basis of the problem that Faulkner's trust theory of testimony is designed to solve: that of explaining the nature of the positive reasons that make it reasonable for *A* to accept *S*'s testimony, and in particular, to trust *S* for the truth of what she says.

Faulkner presents the problem of cooperation by drawing a parallel between two choice situations in contexts where one subject can decide whether to trust another. The first is a purely practical decision, of the kind subjects need to make in the 'Trust Game'. In this game, an 'investor' is given an initial sum of money (say £10), and is asked to choose whether to cooperate, and to transfer some or all of it to the 'trustee', or whether to defect and keep it all to himself. Whatever sum is transferred is then multiplied, say by a factor of 4, and the trustee is asked to decide whether to cooperate, and give back all or part of the multiplied sum to the investor, or to defect and keep it all to herself. Given very minimal assumptions about the two parties' preferences, it is plausible to presume, first, that the trustee would prefer to keep whatever money is transferred to her; and, second, that the best outcome for the trustee—where the investor cooperates and the trustee defects—is also the worst outcome for the investor. Because of this, Faulkner concludes, it would be unreasonable for the investor to cooperate, and to make the initial transfer, unless she has reason to think that the trustee would cooperate.

Faulkner then draws a parallel between the trust game and the 'Testimony Game', to argue that a similar conclusion applies to the case of testimonial encounters: audience *A* is unreasonable in accepting speaker *S*'s testimony unless *A* has positive reasons for thinking that *S* is cooperative or trustworthy. Faulkner bases this conclusion on a *parity claim*, according to which, given some very minimal assumption about the preferences of the parties, the testimony game has the same payoff structure as the trust game: the best option for the potentially trusted party, the speaker, is the worst outcome for the potentially

trusting party, the audience (2011, 6). He then goes on to claim that this parallel implies that, as in the trust game, "one can equally draw the conclusion that [in the testimony game] it is not reasonable to trust without a supporting reason that rationalizes trust" (2011, 6).

While I ultimately agree with Faulkner's conclusion, I have worries about his parity claim, and therefore with his argument for this conclusion. Why should we think that in the testimony game, as in the trust game, the best outcome for the speaker is the worst for the audience? The good that the audience hopes to obtain in the testimony game—information—is very different from that which the investor hopes to obtain in the trust game. Unlike money, information is a good whose consumption is non-rivalrous, so that providing the audience with that good does not generally reduce the benefit available to the speaker from its use. So there is no reason to presume that the speaker should generally prefer to provide the audience with misinformation rather than with information. Of course, there are testimonial encounters, such as that between a car dealer and a potential buyer, in which the speaker might have an interest in misleading the audience. However, Faulkner's claim is that such a conflict between the interests of speakers and audiences is not peculiar to some particular testimonial interactions. Instead it is a general feature of testimonial encounters, as exhibited by the testimony game.

While Faulkner's parity claim does not seem to be true if the audience is represented as merely seeking to form accurate, true beliefs, it might be thought that the claim holds if we conceive of her instead as seeking knowledge. Indeed this appears to be how Faulkner attempts to base his parity thesis: The ranking of outcomes attributed to the audience is supposed to be true "to the extent that an audience's interest is epistemic". Our interests as audience is in "learning the truth," which, I take it, involves more than just obtaining a true belief (2011, 5); it involves obtaining knowledge. Arguably, to obtain knowledge requires a speaker who does not only tell the truth when it suits him. It requires a speaker who is trustworthy. Accordingly, Faulkner claims that the best outcome for the audience is obtained when she trusts a trustworthy speaker. But, he claims, the best option for the speaker is to be believed without being trustworthy, and this is the worst outcome for the audience.

However, the assumption that the goal of the audience is the attainment of knowledge only makes matters worse for Faulkner's parity claim. For it points at a fundamental reason for thinking that the trust game and the testimony game cannot have a similar payoff structure. For if what the audience is after in the testimony game is knowledge, then her payoff depends not only on the speaker's decision and on her own; it also depends on whether her trust of the speaker is reasonable. Arguably, even if the speaker happens to be trustworthy, and the audience trusts her, the audience will not end up with knowledge if it was not rational for her to trust the speaker. In contrast, in the trust game, the outcomes for the parties depend only on their decisions, and not on their reasons for deciding as they did.

Whether or not Faulkner is correct in claiming that it is unreasonable to trust a speaker without positive reasons supporting such trust, the claim is not implied by his conclusion about the trust game, because of the disparity between the games: In the trust game, payoffs are not a function of the parties' reasons for their choices, but in the testimony game they are. As a result, the argument from cooperation fails, and some of Faulkner's key arguments against reductive and non-reductive approaches to testimony are lacking. This is so in as much as these arguments attempt to show that reductive and non-reductive approaches to testimony fail to adequately address the problem of cooperation. Thus, a central objection to the reductive theory is that the reductive solution to the problem of cooperation is too restrictive, in that it misses the central reason for trusting the testimony of the speaker: that we trust the speaker (2011, 53-55). Similarly, the central objection made against non-reductive approaches to testimony, such as those offered by Burge and McDowell, is that the acceptance principles offered by them are inconsistent with the principle of reasonable uptake (R), which is supposedly supported by the argument from cooperation: the principle according to which "an audience *A* is warranted in believing [a speaker's testimony] that *p* if and only if *A*'s other attitudes make it reasonable for *A* to believe that *p*" (Faulkner 2011, 119). Since the argument from cooperation fails to establish (R), proponents of the non-reductive view will probably not be convinced by this argument.

The Trust Theory: Solving the Problem of Cooperation

While I have some reservations about the way Faulkner relies on the argument from cooperation to reject reductive and non-reductive approaches to testimony, I think that he is ultimately correct both in his rejection of these approaches, and in his reasons for rejecting them. We can obtain knowledge by trusting speakers, and both approaches fail in explaining how we do so: the reductive theory is unable to address what is distinctive about knowledge based upon trust; and the non-reductive theory is mistaken in suggesting that we do not need positive reasons to trust. So Faulkner is correct in arguing that a main task of the epistemology of testimony, perhaps *the* main task, is that of explaining how such trust can be reasonable. Whether or not his attempt to meet this task is successful, Faulkner's book makes an important contribution to the epistemology of testimony by first focusing our attention on this challenge; and second, by presenting an original approach in the epistemology of testimony, at the heart of which is a sophisticated theory which attempts to meet this challenge. Whether or not Faulkner's trust theory succeeds in doing so, his discussion of the theory clearly advances our understanding of the theoretical options available to us, and of the challenges that we must face on our way towards a solution of the problem.

There are a number of elements in Faulkner's solution to the problem. At the heart of the solution is an account of a thick form of trust, *affective trust*, which Faulkner distinguishes from a thinner notion of *predictive trust*. The main difference between the two forms of trust is that they involve two very different notions of expectation. In trusting *B* to Φ in the predictive sense, we expect that *B* will Φ . In trusting *B* to Φ in the affective sense, our expectation of *B* that she Φ 's is a normative expectation. We expect this of her. This kind of normative expectation is to be cashed out in terms of the reactive attitudes we shall be prepared to feel if *B* fails to Φ . The main task Faulkner sets to himself is that of explaining how we can obtain knowledge by affectively trusting speakers.

The first important element of the suggested solution to the problem of cooperation is a claim about the internalization of norms of trust. Here Faulkner's suggested solution builds on one proposed by Bernard Williams (2002), but also departs from it in a significant way. Williams suggests that as speakers we have internalized the value of sincerity, and

that it is our intrinsic valuation of sincerity that explains the reliability of testimony. His suggestion is that audiences are then rational in trusting speakers to the extent that they have grounds for believing that speakers would be trustworthy. Faulkner rejects the latter element of this solution, arguing that Williams' account of what makes it reasonable for audiences to believe speakers' testimony is at odds with the claim that we are trustworthy as speakers because we intrinsically value sincerity (Faulkner 2011, 177). Instead, he suggests, as audiences we trust speakers because we have internalized norms of trust requiring of us that we trust speakers, in the affective sense, just as we are trustworthy as speakers because we have internalized norms of cooperativeness in conversation.

It is not entirely clear however what kind of role the internalization of norms of trust is supposed to have within the solution of the problem of cooperation. As said above, Faulkner objects to Williams's reductive account of what makes it reasonable for audiences to believe speakers' testimony. He therefore replaces a key element in Williams account—the claim that audiences are rational in trusting speakers when they have grounds for believing that speakers are trustworthy—with the claim that audiences internalize norms of trust. But it is hard to see how the latter claim can play the kind of role played by the rejected claim: that of explaining the rationality of beliefs formed by trusting speakers. The problem emerges from a second fundamental difference between the trust game and the testimony game: This difference lies in the relations between the rationality of choices and the valuation of outcomes. When it comes to the decision how to act, the practical rationality of one's action can be a function of one's own valuation of possible outcomes. In contrast, the epistemic rationality of one's belief, and hence its epistemic status, is arguably not a function of one's own valuation of knowledge, or of various ways of forming beliefs. Thus, in the trust game, if the trustee cares more about the investor's well-being than about her own, or if she intrinsically values performing acts of kindness, then it might be perfectly rational for her to give back all the money transferred to her by the investor. If, on the other hand, she does not place any extraordinary value on the performance of acts of kindness, then giving all the money back to the trustee would arguably not be reasonable for her. In contrast, the audience's valuation of knowledge and trust does not seem to matter for the epistemic rationality of her trusting the speaker. That a thinker does not value

knowledge and accuracy as we do does not mean that she can be *epistemically* rational in forming beliefs contrary to the evidence. Similarly, the fact that a thinker intrinsically values the formation of beliefs in certain ways—by consulting the stars, or by trusting a speaker—does not seem to matter at all for the epistemic rationality of beliefs formed in these ways.

It may therefore seem unclear what role audiences' internalization of norms of trust is supposed to play within an explanation of the rationality of trust and of beliefs thus formed. Speakers' internalization of norms of sincerity can make it rational for them to speak truthfully, even in situations where speaking truthfully would not have been rational otherwise. However, unlike audiences' having evidence about speakers' internalization of norms of trustworthiness, it is difficult to see how audiences' internalization of norms of trust can make a difference to whether it is *epistemically* rational for them to trust speakers. Having rejected the idea that it is evidence that makes our trust rational, Faulkner's alternative suggestion may therefore appear to be irrelevant to the rationality of audiences' trust.

Perhaps Faulkner's idea is this: an explanation of knowledge obtained through (affective) trust should do two things. It must first explain why, faced with the problem of cooperation, audiences do indeed form beliefs by trusting speakers, in the affective sense. Call this the causal-motivational part of the explanation. The second part of the explanation is not causal, but normative, and addresses the epistemic status of beliefs thus formed: At this stage what needs to be explained is how beliefs thus formed can have the epistemic status of knowledge. It might be suggested that the discussion of the internalization of norms of trust plays a role only in the causal-motivational part of the explanation. The normative part of the explanation is done elsewhere (in chapter 6), where Faulkner explicitly sets out to explain how affective trust can be rationally-self supporting, and how affective trust can be epistemically rational.

This appears to be a plausible interpretation of Faulkner's suggested solution. There are, however, reasons for doubt about both parts of the explanation. The first problem involves the causal-motivational part. Our internalization of norms of trust may explain why we trust speakers, in the affective sense. But this is only part of an explanation of why

audiences form beliefs by trusting speakers. What must also be explained is why audiences who affectively trust speakers also believe what speakers tell them. But this part of the explanation is missing from Faulkner's account. And it is not clear whether this is something that Faulkner's account can adequately explain. This is because affective trust, unlike predictive trust, does not seem to involve belief in the trustworthiness of the trustee. Predictive trust involves the belief that the trusted party would do what she is trusted to do, and so it is clear why an audience who predictively trusts a speaker to speak knowledgeably and honestly would believe the speaker's testimony (provided that she understands what the speaker tells her). However, affective trust, on Faulkner's account, does not seem to involve the *belief* that the trusted party would do what she is trusted to do. Instead it involves the *presumption* that the trusted party would do so. But when audiences merely presume that speakers are trustworthy without believing this, 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? Why would trusting the speaker invariably involve believing what she says (provided adequate understanding)? After all, we often presume that a defendant is innocent without forming the belief that she is. Faulkner needs to explain not only why trusting a speaker results in belief, but why it invariably has this result. And he must do so while maintaining the distinction between predictive and affective trust. It is not clear how this can be done.

In previous writings (such as Faulkner 2007), Faulkner in fact denied that affectively trusting a speaker always involves believing her testimony. Instead, he suggested there, it involves accepting the speaker's testimony as true, where acceptance, as suggested by Cohen (1992), is to be distinguished from belief, *inter alia*, in being under our direct voluntary control. This allowed Faulkner to claim that we have a kind of voluntary control over our trust, that we don't have over what we believe (2007, 894). But in *Knowledge on Trust* Faulkner withdraws from the claim that trusting a speaker is consistent with not believing what she says—and rightly so. For in as much as our reactive attitudes are a relevant indication, they seem to suggest that what speaker *S* expects of audience *A*, when she invites *A* to trust her, is not merely that *A* accept her testimony as true. "In telling *A* that *p*," he writes, "*S* will expect to be believed." (2011, 182). Accordingly, "we are liable to

resent...audiences who do not believe us...and [as audiences, to] feel pressure to believe what speakers tell us" (180). Faulkner is therefore right to claim that "When acceptance is motivated by an attitude of trust—when it is a case of trusting—it issues in belief. The act of trusting testimony is the uptake of testimony" (23). What is left unclear is how his account of affective trust can explain this invariable relation between trusting a speaker, and believing her testimony.

Perhaps Faulkner can meet this challenge by suggesting that affective trust involves having *both* normative expectations of the trusted person, and, in addition, the belief that the trusted person will be trustworthy. Indeed some passages in Faulkner (2011) seem to be consistent with such an interpretation. Thus Faulkner writes of the expectation involved in affective trust that "this expectation is *more than* a statement of our subjective probabilities" (146; emphasis mine). And if affective trust involves believing that the trusted party is trustworthy, then it is no mystery why affectively trusting a speaker should result in believing that her testimony is true. However, there are ample indications that Faulkner does not conceive of affective trust as involving belief in the trustworthiness of the trusted party. Faulkner does not conceive of the presumption involved in trust as a kind of belief: It is a different mental state, which like acceptance, is under our voluntary control in ways that belief is not (149-150). It is therefore unclear how Faulkner's account of affective trust can adequately explain the systematic relation between trusting a speaker and believing what she says.

The Epistemic Rationality of Trust

Let us suppose that in spite of what was said in the previous section, Faulkner's account of affective trust can be part of an adequate explanation of the fact that we form beliefs on the basis of trust. We must still ask whether the account of affective trust can serve in the way suggested by Faulkner within an explanation of the epistemic status of belief formed in this way. Can we explain why beliefs formed by affectively trusting a speaker often count as knowledge?

Faulkner suggests a highly original explanation of the epistemic status of trust-based belief. It is original, first, in that it departs from most reductive and non-reductive

accounts, combining plausible elements of both, while rejecting less plausible ones: claiming with the reductionist that A can obtain knowledge by trusting S only if A has epistemic reasons for trusting S , and with the non-reductionist that testimony serves to transmit knowledge. It seems to me that Faulkner is correct in both of these claims.

However, it is Faulkner's account of the epistemic rationality of trust which is, in mind, the most original part of the book. There are two parts to Faulkner's account. First, he argues that it is A 's trust of S that provides A with an epistemic reason for believing S 's testimony that p , in the minimal sense that it makes p subjectively probable for A (153-4). Second, he argues, when the speaker is in fact trustworthy, the fact that the audience trusts the speaker is potential evidence that p , and thus epistemically warrants the audience's belief that p (154-159).

I have worries about the plausibility of both parts of this explanation. First, I have doubts about Faulkner's suggestion that trust provides A with an epistemic reason to believe S 's testimony because it makes it subjectively probable that the testimony is true. Faulkner's idea is that in affectively trusting S , A accepts, or presumes, that S will be motivated to tell the truth because S recognizes A 's dependence on S . And accepting this makes it probable for A that S 's testimony that p is true. The problem with this is that Faulkner does not explain how we are to understand this notion of presumption, and how presumptions can justify beliefs. If we are to conceive of a presumption as a kind of belief backed by a presumptive epistemic right, then it would perhaps be clear how such a presumption can justify a belief. But were this Faulkner's position then his position would not seem to differ from that of the non-reductionist. If, on the other, the mental state ascribed to an audience who presumes that the speaker is trustworthy is similar to that of a judge who presumes that the defendant is innocent because the law requires this of him, then it is not at all clear how such a presumption can make any belief *epistemically* reasonable, not even in a minimal sense.

In any case, as Faulkner is well aware, even if trusting the speaker would have made A 's belief in the truth of S 's testimony epistemically reasonable in the above minimal sense, this might not suffice to render it epistemically warranted. Faulkner argues, however, that when S is in fact trustworthy, then the fact that A trusts S epistemically warrants A 's belief

in *S*'s testimony. For in such a case, he argues, the fact that *A* trusts *S* is potential evidence for the truth of *S*'s testimony. And if so, then the fact that *A* trusts *S* provides *A* with potential evidence that the testimony is true, and hence renders *A*'s belief epistemically warranted. Faulkner argues for the first claim based on an account of potential evidence developed by Achinstein (1978). Essentially, the claim is that *A*'s trust is potential evidence for the truth of *p*, if two conditions hold—both of which are satisfied when *S* is trustworthy: If the fact that *A* trusts *S* makes it sufficiently objectively probable that *p* is true; and if it is sufficiently objectively probable that there is an explanatory relation between *A*'s trust and the truth of *p*.

However, the question is whether the fact that her trust is potential evidence for *p* suffices to warrant *A*'s belief that *p*. The fact that it is potential evidence for *p* would surely warrant *A*'s belief that *p* were this belief well-based on this evidence. However, 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, let alone well based. Suppose, for the sake of illustration, that the following is true of a certain type of cancer patients: if patients of this kind have a desire to survive their illness, then this desire not only causes them to believe that they will survive it; it also sufficiently increases their chances of survival, and would play a role in explaining their survival should they survive, so that having the desire is potential evidence for the proposition that they will survive. Arguably, that such a patient desires to survive does not provide her with an epistemic reason for believing that she will survive, and does not suffice to warrant such a belief. While different epistemologists would make different suggestions about what else is required to make such belief warranted, most would agree that the following two facts do not suffice to make *A*'s belief that *p* warranted: That *A*'s having mental state *M* is potential evidence that *p* and that *A*'s having *M* is the cause of *A*'s belief that *p*. More than that is required for epistemic warrant. Some might suggest that *A* must also be justified in believing that *M* is evidence for *p*. Others, that *A* must in some sense be reliably sensitive to *M*'s being potential evidence for *p*, or causally related to *p*. The worry is that both types of additional conditions might not be satisfied in the case where an audience affectively trusts a trustworthy speaker.

It seems to me that this is a genuine worry in the case of affective trust. If the audience trusts the speaker because she has internalized norms of trust (regardless of whether her trust is evidence for the speaker's trustworthiness), and if trust invariably leads to belief, then even in cases where the audience's trust happens to be evidence for the truth of the belief, that it is evidence is not the audience's reason for belief. It is not the case that the audience believes that p because she has this evidence, or because she believes she has evidence; and she would believe the same regardless of whether her trust is evidence for her belief. It therefore seems that more needs to be said in order to argue that the audience's trust epistemically warrants her belief.

Moreover, it is not clear that this worry can be set aside by pointing to a parallel with perceptual appearances. True, appearances might be misleading. And it is an epistemological commonplace that thinkers cannot always distinguish between misleading perceptual appearances and veridical ones; and that the fact that a thinker would form a perceptual belief regardless of whether a perceptual appearance is potential evidence for the truth of the belief does not entail that perceptual appearances do not provide epistemological reasons for belief. However, the case of affective trust seems to be significantly different. Thinkers do not invariably form the belief corresponding to their perceptual appearances; trusting a speaker, in contrast, invariably involves believing the speaker's testimony. So the sense in which affective trust seems to involve forming a belief regardless of whether such trust constitutes evidence for the belief is not the sense in which perceptual appearances may lead us to form a belief regardless of whether they constitute evidence for it.

It therefore seems to me that Faulkner's arguments against alternatives to his trust theory are more convincing than his positive account of how knowledge is obtained through trust. Faulkner is correct in arguing that reductive and non-reductive approaches in the epistemology of testimony fail to adequately explain how knowledge is obtained through trust. However, I doubt if his account of affective trust provides us with a successful explanation of this. It is doubtful whether his account explains why we believe the testimony of speakers when we affectively trust them; and it is doubtful whether it explains why beliefs thus formed often constitute knowledge. Further discussion is required

to determine whether the gaps in the explanation can be filled. However, given the worries raised here, I doubt if any non-doxastic account of trust which like the account of affective trust, suggests that trusting a person need not involve belief in her trustworthiness, can successfully meet the challenge.

Arnon Keren

University of Haifa

akeren@research.haifa.ac.il

References

- Achinstein, P. (1978) 'Concepts of Evidence'. In Achinstein, P. (ed.) *The Concept of Evidence*. Oxford: Oxford University Press.
- Cohen, L. J. (1992) *An Essay on Acceptance and Belief*. Oxford: Clarendon Press.
- Faulkner, P. (2007) 'Telling and Trusting', *Mind*, 116: 875-902.
- Faulkner, P. (2011) *Knowledge on Trust*. Oxford: Oxford University Press.
- Fricker, E. (2006) 'Second Hand Knowledge', *Philosophy and Phenomenological Research*, 73: 592-618.
- Hinchman, E. S. (2005) 'Telling as Inviting to Trust', *Philosophy and Phenomenological Research* 70: 562-587.
- Holton, R. (1994) 'Deciding to Trust, Coming to Believe', *Australasian Journal of Philosophy*, 72: 63-76.
- Keren, A. (2007) 'Epistemic Authority, Testimony and the Transmission of Knowledge', *Episteme*, 4: 368-381.
- McMyler, B. (2011). *Testimony, Trust, and Authority*. Oxford: Oxford University Press.
- Moran, R. (2005) 'Getting Told and Being Believed', *Philosophers' Imprints*, 5: 1-29.
- Williams, B. (2002) *Truth and Truthfulness*. Princeton: Princeton University Press.