

Partiality and Prejudice in Trusting¹

You can trust your friends. You *should* trust your friends. Not all of your friends all of the time: you can reasonably trust different friends to different degrees, and in different domains. Still, we often trust our friends, and it is often reasonable to do so. Why is this? In this paper I explore how and whether friendship gives us reasons to trust our friends, reasons which may outstrip or conflict with our epistemic reasons. In the final section, I will sketch some related questions concerning trust based on the trustee's race, gender, or other social identity.

1. Trust, Reliance and Belief

There are rival philosophical accounts of what trust is. But most have a common underlying form: trusting someone to do something involves relying upon her to do it, plus some extra factor which distinguishes genuine trust from the attitude of (mere) reliance we may take to inanimate objects like climbing ropes, cars or crash helmets. So perhaps trust involves reliance plus a tendency to certain reactive attitudes; or perhaps trust involves reliance plus an expectation that the trustee will have certain sorts of motives, such as goodwill, or concern for you, or concern for your mutual relationship; or perhaps trust involves reliance plus an expectation of moral integrity. (McLeod 2006 includes a good critical overview of various accounts of trust.)

Fortunately, for the purposes of this paper I do not need to settle on a particular account of trust, since most of what I want to say turns on their common factor, the notion of reliance. This is spelt out in different ways by different authors, but I will use a version borrowed from Richard Holton (1994). For Holton, to rely upon someone to do something is to work the supposition that she will do it into your plans. Working the supposition that p into your plans need not involve believing that p , even if we think of belief as a disposition. We may recognise that working on the supposition that p is the best option in the circumstances, and thus decide to do this, but this need not involve believing that p ; if the circumstances change, we might abandon the supposition that p . Moreover, as Holton argues, 'Those who see belief

¹ This paper was presented at the Universities of Aberdeen, Edinburgh, Nottingham and St Andrews. I am grateful to audience members for some very helpful comments and suggestions, and also to two anonymous referees for this journal.

as a disposition must have some way of distinguishing belief from perfect pretence' (68).

The gap between relying on someone to do something and believing that she will do it is not a mere quirk of Holton's account of reliance. Instead, it is a crucial feature of any notion of reliance suited to form the basis of a notion of trust. This is because in some circumstances trust seems to be a matter of decision: we can sometimes directly control whether we trust, even though we lack that kind of direct control over our beliefs. Trusting someone may eventually lead to belief, but it is possible for the trust to come before the belief. So trusting someone to do something need not involve belief that she is trustworthy, nor belief that she will do what she is trusted to do, nor even belief that it is likely she will do it.

I have distinguished between trusting someone to do something and believing that she will perform the relevant task. These can come apart in both directions: you can trust someone to do something without believing that she will, so long as you do not believe that she will *not* do it. And you can believe that someone will do something without trusting her to do so, if this is a situation in which only reliance is appropriate, not genuine trust. So for example, you might rather selfishly rely on your neighbour to tidy your shared garden, not because she is concerned for your well-being, nor because she promised to tidy the garden, nor because you have done your share of tidying, but because she herself wants a tidy garden. In such a situation, you believe that your neighbour will tidy the garden, but you do not trust her to do so (and nor do you distrust her in this respect): you would feel disappointed but not resentful if she didn't do the job. After all, you could hardly complain.

For simplicity's sake, I will ignore cases like that of your attitude to your neighbour, in which reliance but not trust is appropriate. To investigate these properly would require me to discuss the differences between rival theories of trust, and their rival stories about when trust, as opposed to mere reliance, is appropriate. Instead, I will focus on situations in which a decision to rely is a decision to trust, and a belief that someone is reliable amounts to a belief that she is trustworthy.

One terminological note before we go on. I am concerned with different sorts of reasons to trust – roughly, epistemic and non-epistemic reasons (reasons of friendship in particular). But this distinction amongst reasons does not coincide with the distinction between trusting someone as a speaker or source of knowledge, and trusting her as an actor, even though the former is sometimes called ‘epistemic trust’, in contrast with ‘practical trust’. Trusting someone as a source of knowledge involves relying upon her to speak knowledgeably (or to speak truthfully); this might in principle be done for epistemic and/or for non-epistemic reasons. (As in other cases, trusting in such contexts will also involve something additional to merely relying.)

2. Trust and Knowledge are Compatible

I aim to explore situations in which epistemic reasons to trust or distrust may conflict with non-epistemic reasons to trust or distrust. But a certain misconception might threaten this project. The misconception is that trust and knowledge are somehow incompatible with one another, or at least in tension: if I rely upon you only after I have conducted a cold-eyed assessment of the evidence, then this doesn’t seem a very trusting attitude. Moreover searching for evidence is often taken as an indication of mistrust: if I find that my boss keeps checking my work, or that my husband keeps checking my whereabouts, I will resent their lack of trust in me.

Certainly, trust does not require knowledge: sometimes we trust and are betrayed, sometimes our trust is vindicated though we did not know it would be, perhaps did not even believe it would be. But trust does not require ignorance either: you can trust someone to do something whilst knowing that she will, indeed your trust may be based, in part, on this knowledge. Only in part, because trust goes beyond mere reliance: you can know that someone will do something for fun, for its own sake, or out of self-interest, without its being appropriate for you to trust her to do this. (Different accounts of trust make different claims about what more is required for trust to be appropriate, over and above the appropriateness of reliance.)

Why think that trust requires ignorance? One possible source of confusion is that when we trust people, we rely upon them to act in certain ways as a matter of free choice. You neither trust nor distrust me to metabolise the food in my digestive system, because you see that I have no choice in the matter. But free choices can be

entirely predictable: I know you won't choose steak because you're a vegetarian, and I know you will choose strawberry ice-cream because that's your favourite. Fatalists deny that any of us are free, the sceptic casts a shadow over all of our empirical knowledge, and human beings are complex systems whose behaviour can be hard to predict. But by ordinary standards we often do know what others will freely choose to do, and so can sometimes know that others will fulfil our trust.

What about checking? If I claim to trust you in some respect, but continue to seek evidence about your behaviour, you would be justified in concluding that I do not really trust you. However, there is a pragmatic explanation for the fact that checking is evidence of lack of trust. Checking whether p – seeking evidence with respect to whether p – is unnecessary when I already know whether p . So my checking up on you indicates that I take myself not to know what you will do. None of this shows that trust and knowledge are incompatible. Checking suggests (perceived) lack of knowledge, as well as lack of trust, and does not discriminate between these two.

Finally, unintentional acquisition of further evidence about someone's reliability need not undermine trust. Suppose I trust you to take my cash to the bank, and I happen to bump into you while you're doing it: this doesn't make me trust you any less.

Trust may sometimes occur without knowledge, or even belief, but knowing that someone will do something is entirely compatible with, may even be part of the rationale for, trusting her. So there certainly can be epistemic reasons to trust; my question is how these relate to non-epistemic reasons, and reasons of friendship in particular.

3. Stroud and Keller on Friendship

We stand in different relationships to different people, some chosen, and some not. And some of these relationships seem to underpin special obligations, or at least to make it permissible to treat certain people better than others, whether they are our friends, our parents, our colleagues, or our fellow-citizens. Although the nature and existence of these obligations and permissions is controversial, they are usually discussed in connection with behaviour rather than belief: for example, it seems I have obligations to care for my parents which go beyond any obligations I have

towards your parents, and it seems permissible for me to buy modest gifts for my children, whilst buying nothing for yours.

Friendship can underpin partiality of this kind. Moreover Sarah Stroud (2006) argues that we owe our friends a form of partiality in our beliefs as well as in our actions:

Friendship places demands not just on our feelings or our motivations, but on our beliefs and our methods of forming beliefs...this epistemic partiality is contrary to the standards of epistemic responsibility and justification held up by mainstream epistemological theories. (2006, 499)

Simon Keller argues independently for a similar view, concluding that ‘epistemic norms sometimes conflict with the requirements of good friendship’ (2004, 329).

Suppose you are told that your friend was rude and obnoxious at a recent social event; how should you react to this news? According to Stroud, ‘We need to ask whether...as a good friend, you ought to react differently to this information that you otherwise would, or than a detached observer would. I shall suggest an affirmative answer.’ (2006, 504) Stroud argues that you should go to greater lengths to think of alternative explanations of the reported behaviour, demand higher levels of evidence that your friend did in fact behave as reported, question the motives of the person reporting this to you, and so on. You are not required to believe the very best of your friend come what may, to defy even the strongest of evidence. But you are required to be unusually hesitant to believe the worst.

For Stroud, ‘friendship requires epistemic irrationality’ (2006, 518), at least if we understand epistemic rationality as it is characterised by mainstream epistemological theories; Stroud briefly discusses a less mainstream, more socialised approach which might accommodate friendship. I will return below to the possibility of such non-standard characterisations, but for now will use phrases like ‘epistemic rationality’, ‘epistemic reasons’ and ‘epistemic norms’ to refer to the mainstream notions, i.e. the kind which, according to Stroud, may conflict with friendship.

Neither Stroud nor Keller maintains that favouring our friends in our beliefs is always the right thing to do. Given that we already think relatively well of our friends, there are many epistemically-disreputable reasons why we might resist bad news about

their characters or behaviour. As in any enquiry, we are prone to confirmation bias, the tendency to notice evidence which confirms our existing beliefs, and ignore evidence which undermines these (Mercier and Sperber 2011). Moreover we may resist bad news about a friend if it retrospectively undermines our shared history of friendship: part of the pleasure and value of friendship is a sense of mutual understanding and respect, making it particularly painful to discover we have been wrong all along about a friend. And if we accept bad news about a friend's character, we then face difficult choices amongst unappealing options: should we quietly drop the friend, provoke a confrontation, or accept a continuing friendship contaminated with doubts?

These are genuine phenomena, which indicate how likely we are to be led astray epistemically with respect to our friends. But these phenomena do not show that we have reasons of friendship to resist bad news in this way, or that norms of friendship require us to do so. The considerations are all in some sense selfish – they play on our wish to be right, to have been right, to be a good judge of character, and to avoid difficult situations. None of these are really reasons of friendship, and none are directed towards our friend's interests as opposed to our own. Instead, they are like the many factors, from tiredness to stereotyping, which can disrupt epistemic rationality in any context, leading us to reason badly or misjudge evidence. We do not have 'reasons of tiredness' to be forgetful or sloppy in our thinking, even though tiredness can bring this about, and nor is there a 'norm of tiredness' that we should think poorly.

In contrast, for Keller

A norm of friendship is a truth about what you should do, insofar as you are a particular person's good friend. If you and I share a close friendship, then there may be norms of friendship telling me to keep in touch, to come and pick you up when your car breaks down... These are all things that I might not do for just anyone, but they are things that a good, true friend would do for you. (2004, 330)

Keller argues that the norms of friendship include truths about what you should believe in certain circumstances, insofar as you are a particular person's good friend, as well as truths about what you should do.

Friendship can distort our thinking about our friends. But Stroud and Keller contend that friendship also gives us good, though epistemically dubious, reasons to resist bad news about our friends. Friendship is interestingly different from tiredness, confirmation bias, stereotyping, and so on, since it is a central, valuable part of human flourishing, not a weakness or flaw to be resisted. It is not just that friendship causes us to distort our reasoning; rather, friendship demands such ‘distortion’, and moreover the demands of friendship deserve to be taken seriously. Only a cartoon villain would seek to avoid all reasoning whilst under the influence of friendship. This is why it is of particular interest to discover whether norms of friendship can indeed conflict with epistemic norms: if there is such a conflict, it is far from clear how we ought to respond.

4. Trust, Belief and Epistemic Norms

How do these claims about friendship apply to the particular case of our beliefs and behaviour around trusting? Trustworthiness is an admirable trait, thus, if Stroud and Keller are correct, it is one we should be quick to ascribe to our friends, and slow to withhold. On this picture, friendship can give us reasons to believe our friends trustworthy, reasons which can go beyond whatever epistemic reasons we have to adopt this belief; these are reasons which do not apply to beliefs about the trustworthiness of nonfriends. Belief in our friends’ trustworthiness is thus a special case of the phenomenon discussed by both Stroud and Keller.

But, as I discussed earlier, not all trust involves belief: trusting someone to do X is compatible with lack of belief that she will do X. Where trust does not involve belief, it might seem obvious that trust is therefore not governed by epistemic norms. This would then leave space for friendship – or prudence, or morality, or etiquette – to give us reasons to trust someone to do something even when we do not have epistemic reasons to believe she will fulfil our trust. We do not need epistemic reasons to believe, because we do not need to believe. But although this sort of trust-without-belief can be reasonable without positive epistemic support, it can nevertheless be made unreasonable if it is undermined by epistemic factors. Although it can be reasonable to trust someone without good epistemic reason to believe that she is

trustworthy, trust becomes unreasonable when there is good epistemic reason to believe the person is untrustworthy.

In Holton's terms, this is because working on the supposition that p is rationally incompatible with belief that not- p . That is because if you believe that not- p , you should work the supposition that not- p into your plans. And if you work both the supposition that p and the supposition that not- p into your plans, you will be unable to plan coherently. You might act in public as if supposing that p , whilst believing that not- p – this might be the best option politically – but Holton argues that belief that not- p will manifest itself in contingency plans, anxieties, and other private behaviours. If you have overwhelming evidence that not- p , then you ought to believe that not- p . And if you ought to believe that not- p , then you ought work the supposition that not- p into your plans, and thus you ought not to work the supposition that p into your plans.

So it is epistemically unreasonable to rely on someone to do something if you have overwhelming evidence that she will not do it. Again, this is not simply a quirk of Holton's account of reliance. Instead, it is a crucial feature of any notion of reliance suited to form the basis of a notion of trust. This is because, although we can in some circumstances choose to trust, in other circumstances we simply cannot, and this is sometimes because of the weight of evidence that the person is not trustworthy. In such circumstances we may try to help the person, or pretend to trust, or reconcile ourselves to being let down, but we cannot just decide to trust in the face of the evidence.

Suppose, however, that we are in a situation in which, epistemically speaking, we are neither required to believe someone to be trustworthy, nor required to believe her to be untrustworthy. We are epistemically permitted to choose whether or not to trust-without-belief. By hypothesis, in such situations epistemic reasons do not govern our choice, though they help determine whether we are in such a situation in the first place. How, then, might reasons of friendship help determine our choice? Recall Keller's claim that if you and I are friends, this gives me reason to come and pick you up when your car breaks down, for example. If I do this, you benefit. Likewise, if I trust you, you are likely to benefit: being trusted brings freedoms and opportunities not otherwise available. Moreover, coming to pick you up when your car breaks

down is somewhat onerous for me, and I don't seem to be obliged to do this for strangers (supposing they're merely inconvenienced by the breakdown, not endangered). Likewise, trusting you can be somewhat onerous for me, especially where the evidence isn't strong enough to require belief that you are in fact trustworthy. Trust often involves risk, and I don't seem to be obliged to take these risks for strangers.

So there are reasons of friendship to trust our friends when this is epistemically permissible, even where we do not believe our friends will fulfil our trust. And in such situations we also have self-interested reasons to trust our friends – trusting your friend will help deepen the friendship, and encourage her to trust you in return, bringing future rewards for both of you.

5. (When) do Friendship and Epistemic Rationality Conflict?

For the remainder of this paper, I will focus on cases in which epistemic reasons do have some bearing on whether we should trust, either by seeming to require us to believe that the person in question is untrustworthy (so seeming to forbid trust), or by seeming to forbid us to believe that the person is trustworthy, because we do not have enough evidence, for example (so seeming to require our trust to be at best trust-without-belief). I will consider what friendship requires or demands in such cases. I will argue that in fact we do not have good grounds for thinking there is a conflict between the demands of friendship and the demands of (mainstream) epistemic rationality when it comes to beliefs about our friends' trustworthiness. This is in part because of some special features of trustworthiness, and in part because of the limitations of Stroud's and Keller's arguments.

Both Stroud and Keller support their views by describing cases in which the demands of friendship seem to conflict with the demands of epistemic rationality, as measured by a counterfactual difference between what we seem permitted or required to believe when a friend is involved, and what we would have seemed permitted or required to believe had a non-friend been involved in an otherwise-similar case. I do not object to this case-driven methodology in principle. However, I will provide a number of reasons to think that any intuitions about the counterfactual differences in what we are

required/permitted to believe can be explained without supposing that the demands of friendship conflict with those of epistemic rationality.

My arguments are not entirely conclusive, because the verdict about each case depends on its particular details, and I rely upon quasi-empirical generalisations about what we tend to know about our friends and others, how people are often motivated, and so on. As Stroud and Keller say themselves, different friendships have different characters, and make different demands upon us. (Keller mentions a friendship between ‘two competent, self-assured philosophers’ who expect of each other that their ‘belief-forming mechanisms...aim unflinchingly at the truth’, in some domains at least (2004, 339).)

So I do not take myself to have demonstrated that reasons of friendship never conflict with epistemic reasons. But I do hope to show that such conflicts can arise, if at all, only under very specific circumstances, more specific than Stroud and Keller seem to allow, and such circumstances are all the more rare whether beliefs about trustworthiness or untrustworthiness are concerned. My strategy is first to show how epistemic rationality may permit or require more than Stroud and Keller suggest, and second to show how friendship may require less than Stroud and Keller suggest.

5.1 What Epistemic Rationality Requires

Stroud and Keller are surely right that we ought sometimes to respond differently to new evidence depending upon whether a friend is involved. But for this to provoke a conflict with epistemic rationality, this differential response needs to go beyond anything which can be explained by our differential epistemic situation with respect to our friends. Put simply, we already know a lot about our friends, and this can give us good epistemic reason for treating new information about our friends differently from new information about our non-friends. We already have lots of beliefs about our friends, and these beliefs count as evidence to be weighed in the balance with new stories about our friends, giving us some reason to reject those stories which do not fit our existing picture.

This breadth of evidence is not exclusive to friendship, for we can come to know work colleagues, carers, patients, or even our students better than we know our

friends, in some respects at least. Stroud discusses a biographer who becomes expert on a particular historical figure without ever meeting her; we might also consider an obsessive fan who acquires more information about her idol than even the idol possesses. And where we do know a great deal about non-friends, we should be cautious about accepting new stories which conflict with the mass of existing evidence, just as we are with our friends: I have never known my colleague to lose his temper in a meeting, and this makes me sceptical about the claim that he did so yesterday, even though we are not friends.

However, sheer familiarity alone cannot explain the asymmetry in our judgements about trusting our friends. The more we think we know about some subject matter, the more cautious we are likely to be in accepting evidence which seems to fly in the face of what we already believe. This consideration explains why we are slow to change our minds about people we think we know well, whether they are friends or not. But this explains a tendency to trust our friends only if we already believe that our friends are trustworthy, and this tendency is epistemically reasonable only insofar as the pre-existing belief in trustworthiness is epistemically reasonable.

Although she does not explicitly discuss trust or trustworthiness, Stroud does argue that our existing beliefs about the positive characteristics of our friends do not really have the evidential status we attribute to them. After all, those beliefs themselves will have been formed under the pressure of friendship, which pushes us to put a positive spin on the words and behaviour of our friends.

But think how much of that “knowledge” of your friend’s character may already have been corrupted and slanted by the interpretive heuristics of the good friend. It is simply not clear that your beliefs about your friend’s character constitute evidence in a sense that supplies epistemic justification. (2006, 516)

Stroud claims that friendship-driven beliefs are ‘corrupted and slanted’, so cannot in turn justify further beliefs. Her opponent denies the premise, seeing the earlier beliefs as epistemically justified by earlier experience of the friend’s character. Both sides agree that friendship does not license or require us to believe whatever we like about our friends. And both should acknowledge that we sometimes find it difficult to think

clearly about our friends, just as we find it difficult to think clearly when tired, distracted, or drunk. Whilst this may reduce the evidentiary value of our existing beliefs about our friends, it also emphasises the difficulty of showing, as Stroud must, that friendship really provides good reasons for belief.

Moreover, Stroud's objection here seems to ignore the extent to which becoming friends is typically a gradual process. We have not always been friends with our friends, and only gradually take on the obligations – whatever they are – of the good friend. So earlier beliefs about our friend's behaviour may have been less 'corrupted and slanted', and thus can provide some grounds for our current opinions.

In the case of trust, in particular, the very fact of friendship can provide evidence of trustworthiness. This is because there is often a two-way causal interaction between friendship and trustworthiness. Roughly, people are more likely to behave in a trustworthy manner towards their friends, and we are more likely to form friendships with people we consider to be trustworthy. (Again, I take these to be quasi-empirical claims about typical behaviour, not strict generalisations; friendships in which this is not the case, or indeed cultures in which this behaviour is not typical, would make different degrees of trust epistemically rational.)

Many people are more motivated by the trust of their friends, and by the commitments they make to their friends, than they are with respect to strangers. This difference may be generated by noble considerations of friendship, or by more pragmatic concerns: there can be more to lose from letting down a friend. This differential behaviour is not entirely admirable – promises should be kept, and lies should not be told, no matter who the audience is – though in situations where we cannot keep all of our conflicting promises, it may be acceptable to prioritise our friends, other things being equal. Nevertheless, it is often more reasonable to believe your friends will prove trustworthy to you than to believe this of strangers, since your friends are objectively more likely to fulfil *your* trust, however they behave to others. (In Stroud's main case, where you hear a story about your friend's behaviour towards some third party, this consideration does not apply: perhaps *you* can trust your friend, but others should not.)

There are, as ever, exceptions. There are circumstances in which your friend might let you down instead of disappointing someone else, because your friend hopes that you will understand and forgive her. She may also find it more tempting to lie about some matters to you than she would do to a therapist, doctor, or even a stranger on a train, because she is concerned to maintain your good opinion of her, or because she doesn't want to worry you. But if she takes these liberties too often, you will feel you have been taken for granted, and come to resent your friend. Friendship requires mutual respect and openness as well as forgiveness.

So it will often be more reasonable – epistemically more reasonable – for you to believe that your friends will be trustworthy to you than to believe this of strangers. And we can have a good sense of where the exceptions to this rule may lie, in part through knowing our friends well: we know their weaknesses as well as their strengths.

I have argued that we often have good reason to expect our friends to prove trustworthy. But this point should be distinguished from the superficially similar claim that trust can encourage trustworthy behaviour, at least where the trust is made manifest to the trustee (Paul Faulkner explores related issues in his (2007)). This is not the universal situation, but again it may be more common with respect to friends: our friends may be more likely to recognise and respond to our belief in them. Similar considerations count against distrusting our friends: friends are more likely to be insulted by distrust, or even suspension of judgement, and step back from trustworthy behaviour, whilst a stranger might not give this a second thought.

Do such phenomena create additional epistemic reasons to believe our friends will fulfil our trust – and to make this belief manifest? In this respect, believing someone trustworthy can be a self-fulfilling prophecy: if you believe, that will help make it true. Given that we want our friends to be trustworthy, this clearly creates a prudential reason to believe them trustworthy. But it is less clear that it creates an epistemic reason to believe. After all, these reasons count equally well in favour of distrust: if you believe your friend to be untrustworthy, your believing this will help make it true. There is a prudential asymmetry here – we have good reason to

encourage trustworthiness rather than untrustworthiness – but no epistemic asymmetry.

I have been suggesting that many of us may have particularly good epistemic reasons to think that our friends will prove trustworthy to us, because in fact people are motivated by friendship to be trustworthy; this makes it more difficult to show that the demands of friendship take us beyond what is epistemically reasonable in such cases. Moreover, making judgements about trustworthiness is an important element in developing a friendship in the first place, or continuing to maintain a friendship.

Keller notes a more general point:

If I say ‘That Steven, he certainly isn’t a selfish, lying scoundrel’, and you say, ‘You only believe that because you’re his friend’, then you are probably getting things the wrong way around. I am friends with Steven partially *because* I do not think that he is a selfish, lying scoundrel. (2004, 337, Keller’s italics)

It may be too idealistic to think that we always seek out moral virtue in forming friendships, and indeed a priggish insistence on this would damage many relationships. But certain vices place practical obstacles in the way of friendship, even if we hold back from moralising about this. It is difficult to maintain a friendship with someone you consider untrustworthy, even if you manage to forgive this, for example because friendship involves joint planning and cooperation, and a sense of mutual respect.

Sometimes we differentiate, trusting one friend sober but not drunk, trusting another friend to keep a secret, but not to repay a loan. But this guides our reaction to negative stories about them: no surprise that Belinda behaved badly after drinking all evening, but it’s hard to believe that she stole a purse at the gym. Trustworthiness is not the same virtue as honesty or sincerity, since people can be untrustworthy through ignorance or incompetence: someone who sincerely relates a bunch of falsehoods, or sincerely undertakes to complete a task she is incapable of performing, is untrustworthy though she is well-intentioned. Nevertheless, the self-knowledge required for trustworthiness in the fullest sense is something we value in our friends,

and attempt to demonstrate to them. So we take ourselves to have good evidence that our friends are trustworthy in at least certain respects.

The situations described by Keller and Stroud are familiar to us; we often favour our friends in our beliefs, including beliefs about their trustworthiness, and this is often the right thing to do, not a failing or weakness to be overcome. However it is especially difficult to identify cases in which friendship gives us good reason to believe someone will fulfil our trust, when we ought not to believe the same thing of a non-friend even given the same evidence. This is because it is especially difficult to construct a case in which we would have just the same evidence of a non-friend's trustworthiness to us as we have of a friend's trustworthiness to us. So the picture is messy, and the case for non-epistemic reasons for belief is as-yet unproven.

But it would be rash to conclude that we are never justified in believing better of a friend than we do a non-friend except where evidence demands this. Instead of insisting on this conclusion, I will explore the idea that even if there are such differences, this does not show that friendship norms sometimes conflict with epistemic norms, or that friendship can require epistemic irrationality.

How can we judge what epistemic norms, or epistemic rationality, demand of a subject in a given situation? Both Keller and Stroud operate a kind of counterfactual test of epistemic reasonableness. They consider what we would (or should) conclude about a non-friend, given the evidence available, then compare this to what we do (or should) conclude about a friend, given the same evidence. For example Keller writes that, in a situation where your friend denies an accusation 'You might find yourself believing his story while having the nagging realization that you would not believe him did he not happen to be your friend' (2004, 331). Stroud asks 'whether, as a good friend, you ought to react differently to this information [which casts your friend in a bad light] than you otherwise would, or than a detached observer would' (504), and she answers that you should indeed react differently.

The thought is that friendship *per se* is an epistemically-irrelevant factor, and so if this factor makes us adopt a doxastic attitude we would not adopt in its absence, then it tips us into epistemic irrationality. In itself, the counterfactual difference between

what we would believe of a stranger and what we do believe of a friend looks symmetrical. So we judge our friends differently from our non-friends: why suppose that it is our treatment of our friends which is epistemically irrational? Perhaps we're being too mean to our non-friends. To avoid this concern, the counterfactual test sometimes turns on direct comparison between what we *ought* rationally to believe of a stranger, and what we *ought* to believe of a friend, rather than what we do and would believe.

Let's concede that if we must choose a single epistemic standard, then it is plausible that the standard we apply to non-friends is the epistemically rational one. So if friendship requires us to deviate from the standards we ought to apply to non-friends, then the norms of friendship conflict with epistemic norms. But why think that we must choose a single standard for all cases? This is to assume that there is always a single epistemically permissible or reasonable attitude to take to a given proposition, on the basis of a given array of evidence. But this assumption is dubious. Sometimes the evidence clearly tells us to believe that p (or that not- p); sometimes it is clear that we do not have enough evidence to arrive at a reasonable belief either way. But in many circumstances, a number of different options are permitted.

Different people – and the same person in different moods – can reasonably differ in their doxastic policies. Some people are somewhat quicker to belief than others are, some put more weight on the evidence of their own senses than others do, some are more sceptical about testimony than others are. There are many reasonable attitudes to epistemic risk, many acceptable ways to strike the balance between pursuing true belief and avoiding false belief. We are all familiar with situations in which there is more than one conclusion which could reasonably be drawn, even against the same background assumptions. For example, disagreements between colleagues about what grade a student paper deserves are not always reducible to disagreements about what qualities a student paper ought to exhibit. One can be an atheist, and regard oneself as epistemically reasonable in this respect, without regarding theists as epistemically unreasonable. Scientists can reasonably disagree about the merits of a given theory, even though all have access to the relevant evidence.

If, in a given situation, more than one set of standards is epistemically reasonable, then the fact that we employ different standards with respect to our friends than we do with non-friends does not entail that we are doing something epistemically unreasonable in either case. This picture explains why friendship does not give us reason to believe that our friends are absolutely marvellous, only to see them in what Keller describes as ‘the best possible light’ (330). ‘Possible’ here may be governed by epistemic permissibility, not psychological capability; friendship requires us to give our friends the benefit of the doubt, but no more than this.

I am suggesting a form of epistemic permissiveness, a denial of the doctrine which Roger White (2005) calls ‘Uniqueness’: given one’s total evidence, there is a unique rational doxastic attitude that one can take to any proposition (see also Feldman 2007). The permissive, anti-Uniqueness thought is that, under certain circumstances, more than one doxastic attitude is epistemically permissible, so that a difference between how we treat our friends and how we treat our non-friends, even given the same evidence, need not amount to epistemic irrationality in either case. White raises a number of important concerns about the denial of Uniqueness, which I cannot address here. Nevertheless, we should note that epistemic permissiveness, if it proves to be coherent, places yet another obstacle in the way of those who argue that non-epistemic reasons, such as reasons of friendship, can conflict with epistemic reasons, especially where trust is concerned.

Recall that Stroud is committed only to a conflict between friendship and epistemic rationality as it is usually characterised: she remains open to the possibility that we can develop a characterisation of epistemic rationality which defuses this conflict. Stroud’s very interesting suggestion, which she does not explore in detail, is that from a social perspective there may be overall epistemic reason for us to have the practice of taking different epistemic attitudes to different people, depending upon our relationship to them. Perhaps the truth is more likely to come out about an individual when there are at least some people disposed to take a very charitable view of her, as well as others who are disposed to take the ‘objective’ view: the interplay between these perspectives may be of epistemic benefit to all concerned. (It would be worth exploring the role of enmity, as well as friendship, in such a system.) Moreover the adoption of epistemic permissiveness – with or without the social perspective Stroud

considers – might be another way of shifting away from ‘mainstream’ epistemological theorising in response to the apparent demands of friendship in this realm. If we take these concerns seriously, this may provide some grounds for endorsing permissiveness.

We have established that we often have better epistemic reason to believe that our friends will fulfil our trust than we do for the same belief about non-friends. Thus the belief that our friends will fulfil our trust is often better justified than the analogous belief about non-friends. Even where we do not have this clear difference, given epistemic permissiveness, it may be epistemically permissible to believe as Stroud and Keller say that friendship requires us to do, to believe that our friends will fulfil our trust on evidence which would not compel us to believe that non-friends will do so.

Each of these points make it incrementally harder to establish a conflict between epistemic norms and friendship norms, although I have not demonstrated that such conflicts cannot possibly arise. Stroud and Keller proceed by describing cases in which we seem required to think better of our friends than of our nonfriends, and they argue that these differences in what we are required to do cannot be explained on (mainstream) epistemic grounds. I have argued that the epistemic resources at our disposal are richer than Stroud and Keller suggest – especially where beliefs about trustworthiness and untrustworthiness are concerned – and thus that it is correspondingly harder to show that friendship requires us to go beyond mainstream epistemic rationality.

5.2 What Friendship Demands

I have argued that epistemic rationality often permits us to treat our friends in a distinctive way when it comes to trust or distrust. In this section I will argue that friendship often demands less than Keller and Stroud suggest, again making it more difficult to show that the demands of friendship outstrip what is permitted by epistemic rationality.

Keller’s central case is explicitly designed to circumvent some of the arguments I presented above. Eric’s friend Rebecca invites him to attend her poetry reading, at a

venue where he has heard lots of terrible poetry; Eric has no previous experience of Rebecca's poetry. Keller argues that, on purely evidential grounds, Eric should believe that Rebecca's poetry will be bad: he would believe this of any unfamiliar poet about to read in this venue. But, as a friend, Eric should not believe this of Rebecca. Moreover, during the reading, Eric ought to form beliefs about Rebecca's poetry more charitably than he would regarding a stranger's poetry.

Significantly, Keller argues that neither Rebecca's being Eric's friend, nor their shared history, give him evidence that she is a good poet. Eric does not choose his friends for their poetic talents, yet it seems his friendship with Rebecca requires him to think well of her in this respect. Thus some of the arguments I presented above, which played on the connection between trustworthiness and friendship, do not apply here.

I will argue, however, that in such cases friendship demands less than Keller suggests. We are prompted to agree that friendship can give us reason to believe something we would not believe otherwise. A more deflationary description of these situations would concede that friendship gives us reason to act as if we believe the proposition in question, but that this falls short of having a reason to actually believe it. Thus we should protest out loud when someone tells us a negative story about a friend, but we should privately judge on the evidence. And Eric's friendship with Rebecca obliges him to behave in supportive ways towards her, maybe helping her find further opportunities to read her poetry but perhaps it does not oblige him actually to believe that her poetry is good.

Keller responds that what Rebecca wants is that Eric believe her poetry is good, not just that he act as if he does. 'I doubt, however, that this kind of pretend approval is what we really want in a good friend. It is not likely to be what Rebecca wants when she looks for a supportive person in the audience.' (335) But what our friends want of us is not a straightforward guide to what friendship gives us reason to do for our friends. There may be a conflict between acting as your friend wishes, and acting in her best interests, though concerns about paternalism should prevent us invariably opting for the latter. Moreover, what Rebecca wants is for it to be the case that her poetry is good. So at least one reason why she wants Sam to believe that her poetry is

good is that this will provide evidence that the poetry is indeed good; she might care less about what Sam believes if she takes him to be a poor judge of poetry. But if Sam's belief is guided by non-epistemic factors, it is not evidence of its own truth, and thus lacks the instrumental value Rebecca hopes it will have.

Likewise, Stroud argues that 'A good friend does not defend her friend outwardly...while inwardly believing the worst of her friend' (505). But again there are a number of possible explanations available. First, we can concede that a good friend should not believe the *worst* of her friend, i.e. should not lean towards the least charitable extreme of what is epistemically permissible; this is compatible with the claim that a good friend need not believe what is epistemically impermissible. Second, the person described here is behaving insincerely, saying one thing whilst believing another, and this provides at least some reason to condemn her behaviour.

These difficulties are heightened when we consider trust. I argued above that friendship gives us reason to trust our friends, where this is epistemically permissible at least, because of the benefits our friends will reap from our trusting them. But many of these benefits are available so long as we trust, and do not depend upon us also believing that our trust will be fulfilled. Moreover the risks we take when we trust without believing are not always less onerous than the risks we take when we believe that our trust will be fulfilled. Friendship gives us reason to take these risks in order to offer our friends these benefits, but this falls short of giving us reason to believe our friends trustworthy where this belief is not supported by the evidence.

Likewise, where epistemic rationality requires us to believe that our friends are untrustworthy (rather than suspending judgement), friendship may nevertheless give us reasons to behave as if we trust, even if genuine trust is not an available choice. Behaving as if we trust involves risk, both practical and emotional, and this may be something we are required to do for our friends. This rings true: even when we know we have generated overwhelming evidence of our own untrustworthiness, we might still ask of our friends – but not others – that they give us one last chance to mend our ways, to prove that things can be different from now on.

So the arguments which Stroud and Keller present, and the types of case they discuss, have not given us good reason to think that friendship requires us to neglect epistemic rationality either in our beliefs about (un)trustworthiness or in our trusting. To some extent, this is due to special features of trust, distrust, trustworthiness and untrustworthiness, and to that extent Stroud and Keller's more general arguments emerge unscathed. But not entirely: several of the objections I have raised apply to beliefs about our friends more generally.

6. Rational Trust, Epistemic Injustice, and Forbidden Base Rates

In this final section, I want to outline some connections between these issues about trusting on the basis of friendship, and issues about trusting – or distrusting – on the basis of race, gender, class, or other social identities. I do not have space to explore these matters in the depth they deserve, so will limit myself to raising some questions which I hope to explore in later work.

Miranda Fricker in her *Epistemic Injustice* (2007), explores the notion of 'testimonial injustice'. The central case of testimonial injustice is identity-prejudicial credibility deficit – to a first approximation, this occurs when a speaker is not taken seriously as an informant because of the audience's prejudices about her race or gender, for example. Fricker has all sorts of interesting things to say about testimonial injustice, and the distinctively epistemic harms to which it gives rise. But I will focus on the fact that credibility deficit is defined both in counterfactual terms, and by reference to what is epistemically required of the audience. Counterfactual: 'the prejudice...results in [the speaker] receiving less credibility than she otherwise would have – a credibility deficit' (2007, 17). Epistemic requirements: '...there is no puzzle about the fair distribution of credibility...the hearer's obligation is obvious: she must match the level of credibility she attributes to her interlocutor to the evidence that he is offering the truth' (19).

Like Stroud and Keller, Fricker uses counterfactual considerations as a guide to what epistemic rationality requires. In situations where (e.g.) a speaker's being female is *per se* an epistemically irrelevant factor, then if this factor makes us adopt a doxastic attitude we would not adopt in its absence, i.e. if the speaker were not female, then it tips us into epistemic irrationality. As with Stroud and Keller, the possibility of

epistemic permissiveness derails this argument: if more than one doxastic attitude is epistemically permissible, then a difference in our responses between two evidentially-equivalent situations need not show that we are epistemically irrational in either case.

But this pushes us to make explicitly moral evaluations of these counterfactual differences in judgement. Setting aside large issues about the compatibility of morality and partiality (Baron 1991), considerations of friendship seem a morally acceptable reason to give someone the benefit of the doubt, at least where this is epistemically permissible. But it does not seem morally acceptable to use considerations of race, gender or class in this way. Even if it is epistemically permissible to be somewhat mean-minded, and epistemically permissible to be fairly charitable, it is not morally permissible to switch between these doxastic policies on grounds of race, gender, or class. Thus epistemic permissiveness – if it proves viable – opens up new questions about epistemic injustice, showing that this can arise even without epistemic irrationality.

Tamar Gendler takes these matters a step further. Like Fricker, she raises many more interesting issues than I can mention here, and I will focus on the notion of ‘forbidden base rates’ (Gendler 2011, section 5, where she discusses Tetlock et al (2000)). We do not live in perfectly egalitarian societies, and race, gender, class and other identities can significantly affect how our lives work out. For example, in the US in 2004, only 10.2% of engineers were female (source: Society of Women Engineers). You may not have known this exact statistic, but the ballpark figure probably doesn’t surprise you.

Now suppose you’re at a reception for engineers and their spouses, and you’re introduced to a male-female couple about whom you know next to nothing. Odds are, he’s the engineer. But if you have anti-sexist instincts, you may feel pulled towards keeping an entirely open mind about which of these two strangers is the engineer, rather than allowing your statistical knowledge to incline you towards the man. If you do ‘slip’ into assuming the man to be the engineer, and this turns out to be a mistake, you’re likely to be more embarrassed than you would be had you wrongly assumed the couple to live in the local area, on the grounds that most guests at the reception

live locally. And a female engineer from out of town is likely to get more annoyed by people assuming she's 'the spouse' than by people assuming she lives locally, even though both assumptions have the same statistical support.

The forbidden base rate here is the statistical information about the relative numbers of male and female engineers. Discussing a different, race-oriented case, Tetlock and his co-workers argue that 'the primary obstacle to using the putatively relevant base rate is not cognitive but moral' (2000, 854). There are many issues to unpack before we can draw any determinate conclusions (Egan 2011 has useful discussion, and also makes the connection to Keller and Stroud). However this at least opens up the possibility that moral reasons – indeed, reasons connected quite directly with issues of fairness and impartiality – can direct us towards doxastic states which are not epistemically permissible.

Neither Gendler nor Tetlock et al are specifically concerned with trust, and Fricker's main concern is with trusting speakers, rather than trusting people to act. Nevertheless, these intriguing arguments have clear connections to debate about trust more generally. But exploring these connections must remain a task for another occasion.

References

- Baier, Annette (1986): 'Trust and Antitrust', *Ethics*, 96: 231-60.
- Baron, Marcia (1991): 'Impartiality and Friendship', *Ethics*, 101.4: 836-57.
- Egan, Andy (2011): 'Comments on Gendler's 'The epistemic costs of implicit bias'', *Philosophical Studies*, 156: 65-79.
- Faulkner, Paul (2007): 'On Telling and Trusting', *Mind* 116: 875-902,
- Feldman, Richard (2007): 'Reasonable Religious Disagreements', in Louise Antony (ed.) *Philosophers without God: Meditations on Atheism and the Secular Life*, Oxford: Oxford University Press.
- Fricker, Miranda (2007): *Epistemic Injustice*, Oxford: Oxford University Press.
- Gendler, Tamar (2011): 'On the Epistemic Costs of Implicit Bias' *Philosophical Studies*, 156: 33-63.
- Holton, Richard (1994): 'Deciding to Trust, Coming to Believe', *Australasian Journal of Philosophy*, 72.1: 63-76.

- Keller, Simon (2004): 'Friendship and Belief', *Philosophical Papers* 33.3: 329-351.
- McLeod, Carolyn (2006): 'Trust', *Stanford Encyclopedia of Philosophy*.
- Mercier, Hugo and Sperber, Dan (2011): 'Why do Humans Reason? Arguments for an Argumentative Theory', *Behavioral and Brain Sciences*, 34, 57-111.
- Stroud, Sarah (2006): 'Epistemic Partiality in Friendship', *Ethics*, 116.3: 498-524.
- Tetlock, P.F, Kristel, O., Elson, B., Green, M., and Lerner, J. (2000): 'The Psychology of the Unthinkable: Taboo Trade-Offs, Forbidden Base-Rates and Heretical Counterfactuals', *Journal of Personality and Social Psychology* 78.5: 853-870.
- White, Roger (2005): 'Epistemic Permissiveness', *Philosophical Perspectives* 19: 445-459.