



This is a repository copy of *Giving the benefit of the doubt*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/126485/>

Version: Accepted Version

Article:

Faulkner, P.R. (2018) Giving the benefit of the doubt. *International Journal of Philosophical Studies*, 26 (2). pp. 139-155. ISSN 0967-2559

<https://doi.org/10.1080/09672559.2018.1440952>

This is an Accepted Manuscript of an article published by Taylor & Francis in *International Journal of Philosophical Studies* on 23 Feb 2018, available online:
<http://www.tandfonline.com/10.1080/09672559.2018.1440952>.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Giving the Benefit of the Doubt

Faced with evidence that what a person said is false, we can nevertheless trust them and so believe what they say — choosing to give them the benefit of the doubt. This is particularly notable when the person is a friend, or someone we are close to. Towards such persons, we demonstrate a remarkable epistemic partiality. We can trust, and so believe, our friends even when the balance of the evidence suggests that what they tell us is false. And insofar as belief is possible, it is also possible to acquire testimonial knowledge on those occasions when the friends know what they tell us. This paper seeks to explain these psychological and epistemological possibilities.

1.

Walking home alone at night you come across someone badly injured. Going to their aid you find, to your horror, it is your estranged partner and they have been stabbed with what appears to be your kitchen knife. You call the ambulance and police, but your partner dies before they arrive. When they do arrive, the police arrest you. The evidence against you is significant. Your prints are all over the knife and you had a motive. In this

nightmare scenario, you know that you are innocent. You know that all the evidence indicating otherwise is misleading and has some other explanation. What you need to do is convince someone of these facts so that this alternative explanation is sought! The problem is that all the evidence suggests that what you say is false. And if one had such seemingly good but in fact misleading evidence that an instrument is broken, there would be no epistemic options: this evidence would mean that one could not but reject its readings. However, in protesting your innocence to your friend, you do not take your friend to be similarly epistemically bound by the evidence. It seems possible for your friend to give you the benefit of the doubt. It seems possible for your friend to *trust you*, and so to learn what really happened.

The aim of this paper is to give an account of what makes this possible. In the next section, I develop the epistemological issues that situate the problem.

2.

Evidence can be misleading. So it is possible for an audience to have evidence that a speaker is lying when in fact the speaker is being truthful and telling the truth. Since the possibility under investigation is a matter of what is possible for an audience, imagine now that you are in the position of the audience who is the friend.

The murder case. The evidence suggests that your friend is guilty of murdering their estranged partner. Your friend comes to you and protests their innocence. In fact, your friend is innocent and has been very cleverly framed.

Three things, I think, are true of this case.

First, it is *psychologically possible for you to believe your friend*. One way of making sense of this possibility is to imagine that, in considering whether your friend is innocent, you bring to bear everything you know about your friend — their history, circumstance, character, values, the fact that they value your friendship and care about truthfulness — and so come to judge that this evidence, in conjunction with their testimony, outweighs the evidence pointing to their guilt. However, this possibility should be excluded: the *murder case* should be read as a case where the *balance of evidence*, when impartially considered, points to your friend’s guilt. The first claim is that even in this situation, you could believe your friend.¹ The claim, here, is not that belief is probable but possible. The way that things would probably play out would be that you would listen to your friend’s account, weigh the evidence, and at a certain point either trust your friend, giving them the benefit of the doubt — or not.² Moreover, that you *should* trust your friend — that you *should* put your faith in them, and so believe them — is a claim that has been much discussed.³ A presupposition of this discussion is that belief is a psychological possibility.

Second, *it is epistemologically possible for you to get to know that your friend is innocent on the basis of believing her when she tells you this*. That this should be taken as an epistemological datum is shown by imagining the *murder case*, as section 1 does, from

¹ For an actual case of giving the benefit of the doubt, consider Peter Hyman’s testimony that he believed O.J. Simpson when the latter told him, before the trial and in the face of the evidence, that he didn’t kill Nicole. (From Ezra Edelman’s documentary *O.J.Simpson: Made in America*, 2017).

² See Audi’s (1997: 407) example of talking to a stranger on a plane and coming to believe them.

³ See Baker (1987), Keller (2004), Stroud (2006), Kawall (2013) and Hawley (2014).

the speaker's perspective. In seeking a friend, so that you can tell them what happened to you, you hope to be believed not merely because this is good in-itself, but because you want your friend *to learn what really happened*. In telling your friend that you are innocent you take it that you have put them in a position to know this, if they can only believe you. Moreover, suppose you were wrong in this and that your telling does not put your friend in this epistemic position; to assert so much would be to deny that trusting someone who knows offers a unique route to knowledge. The second starting point of this paper is the contrary claim that trusting someone who knows does identify a unique route to knowledge.⁴ It is the (epistemic) intuition that your friend could learn what happened, if only they could believe you, which, I think, is as robust as the (psychological) intuition that your friend could believe you.

Nevertheless, third, this epistemological datum is *problematized* by the fact that the balance of evidence supports your believing that your friend is guilty. So while it might be allowed that trusting someone who knows identifies a unique way of getting to know, it can still be denied that it is a way of getting to know *in this case*. That is, in this case the possibility of acquiring knowledge seems to be blocked. This idea might be put in terms of defeat: any possibility of getting to know that your friend is innocent seems to be *normatively defeated*.⁵ The relevant defeater here is normative because given that the evidence of your friend's guilt is in fact misleading, there are no undefeated *objective* defeaters. And if you do manage to believe your friend, trusting her for the truth, it

⁴ For support, see Vendler (1979), Ross (1986), Moran (2005), Hinchman (2005), Faulkner (2007), Keren (2007), McMyler (2011), and Fricker (2012).

⁵ For this classification of defeaters see Lackey (2008: 44-46).

follows that there are no undefeated *subjective* defeaters. Of course, the worry is that epistemically you *should not* trust your friend for the truth given the evidence you have, which, again, is to say that the relevant defeaters are *normative*: the issue is the *epistemic reasonableness* of believing your friend is innocent when the *balance of evidence*, impartially considered, points to your friend's guilt.⁶ It follows that a requirement on any explanation of how you get to know that your friend is innocent is that it must show how believing your friend can be epistemically reasonable.

The philosophical challenge is then to account for these three data-points: (1) to explain how belief in the face of the evidence is a psychological possibility; (2) to explain how the acquisition of testimonial knowledge in the face of the evidence is an epistemological possibility; and (3) to explain how believing your friend can be epistemically reasonable when the balance of evidence, impartially considered, suggests that it is not epistemically reasonable, and thereby implies that the possibility of acquiring knowledge is normatively defeated. In the next three sections, I sketch three possible explanations of these data-points, before arguing for a fourth explanation in the final two sections.

⁶ Similarly, Keller's (2004: 331-332) central case of Rebecca and Eric has the balance of evidence favour disbelief.

3.

One explanation of how belief is possible in the *murder case* starts from the fact that in this case it is *a friend* that tells you she is innocent. “Friendship”, Sarah Stroud (2006: 499) observes, “places demands not just on our feelings or our motivations, but on our *beliefs* and our methods of forming beliefs”. This is to say that friendship “positively demands epistemic bias, understood as an epistemically unjustified departure from epistemic objectivity” (Stroud 2006: 518). Similarly, Simon Keller (2004: 329) concludes that “epistemic norms sometimes conflict with the requirements of friendship”. In listening to these requirements of friendship, we are then led to demonstrate epistemic bias; the balance of evidence might suggest that our friend is guilty, but we do not consider this evidence impartially. First, we question any evidence that puts our friend in a bad light and attempt to undermine its evidential value. If one bit of evidence against your friend is that their estranged partner was stabbed with their kitchen knife, you would try and imagine alternative explanations of this fact under which your friend is innocent; the knife was stolen, for instance. That is, we “try to discredit the evidence” (Stroud 2006: 509). Second, we try to put a different interpretation on what evidence cannot be discredited in order to charitably conceive a friend’s action and motives. If you could not deny that your friend did the killing, you might, for example, then conceive it as an act of self-defence. In this way, we seek interpretations of our friends’ conduct that is as favourable as the evidence will allow. In doing so, Stroud (2006: 508) jokes, “we become spin doctors”.

The idea that we treat testimony from friends differently — that we are epistemically partial and so more likely to believe it — explains the psychological datum,

namely that belief is a possibility in the *murder case*. Moreover, late in her paper Stroud (2006: 522) suggests that this fact should force us to “rethink the assumption that epistemic rationality requires the kind of epistemic objectivity or impartiality from which friendship seems necessarily to depart”. Here Stroud questions the assumption that epistemic partiality entails epistemic irrationality; rather, she suggests that the demands of friendship can make it *epistemically* reasonable to believe your friend. Assuming that this is true of the *murder case*, Stroud both endorses the third epistemological data point — that believing your friend can be epistemically reasonable — and offers a suggestion as to why this is so: the demands of friendship make it so. However, what is left mysterious is *how* the demands of friendship do this. In the next section, I consider one way of developing this explanation and removing this mystery. But first I want to raise a general worry about this explanatory strategy.

The general worry is that *believing-your-friend-out-of-friendship* seems to fall short of simply *believing your friend*; and if it does, then it would leave your friend frustrated. Pamela Hieronymi (2008: 219) makes this point when discussing her own murder case and in criticism of non-cognitive notions of trust: what the friend demands is trust *in the sense of belief*; “your friend wants you to trust that she is innocent; she wants you to believe her”. Hieronymi’s worry is that unless trust entails belief, the friend’s demand to be trusted is unsatisfied. For example, merely accepting what your friend says, and thereby treating your friend *as if* she were innocent, offering her shelter and so on, would not be enough. It would not be enough because you could do this all the while believing, based on the evidence, that your friend is guilty. Similarly, unless *believing-out-of-friendship* entails *believing* — and so entails crediting your friend with knowing

that she is innocent — your friend’s demand to be believed is unsatisfied. The following explanation of how friendship supports belief then goes some way to appeasing this worry.

4.

What is needed is a defence of “the rationality of partiality”, and such an argument is provided by Berislav Marušić (2015: 174). Consider first what reason one has for believing that someone, S, will ϕ . In considering this question, one must distinguish between the perspectives of the *agent* and an *observer*. An observer, if disinterested, will look to the evidence and here S’s track record will be crucial. But from the agent’s point of view whether ‘S will ϕ ’ is a matter up to the agent, S, to decide. The decision to ϕ , or not, is something S resolves by practical reasoning. And in supporting the decision to ϕ these practical reasons equally support S’s belief that she will ϕ ; this belief is a *practical belief*: a belief “formed in light of, and made rational by ...practical reasons” (Marušić 2015: 137).

Two consequences follow from this idea that practical reasons can support belief. First, it would be an instance of *bad faith* and a denial of her agency for S to look to the evidence in forming a belief as to whether she will ϕ . For S, these theoretical reasons for belief are the *wrong kind of reasons* for belief. Second, S can rationally decide to ϕ and so believe that she will ϕ even when the evidence suggests that she will fail to ϕ . For example, S can decide to quit smoking on the basis of good reasons for doing so, and thereby believe that she will do so, even if she has tried and failed before. That is, the fact

that a practical belief “goes against the evidence does not entail that the belief is epistemically irrational ... [because e]vidential considerations are ... the wrong kind of reasons to evaluate such beliefs” (Marušić 2015: 138). Rather, it is S’s practical reasons for quitting that justify S’s belief that she will quit; they do so because these reasons make it likely that this belief is true.⁷

This distinction between the perspectives of the agent and an observer then has an analogue in the epistemology of testimony: it is the distinction between the perspectives of a *participant* and an *observer*. And this distinction is relevant when considering what reason one has for believing that S will ϕ when S tells one that she will ϕ . An audience A takes the observer’s perspective when he treats S’s telling him that she will ϕ as piece of evidence and proceeds to judge whether S will ϕ on the basis of the evidence. Citing Richard Moran, Marušić claims that this response to S’s testimony introduces “disharmony” into the testimonial exchange because S’s expectation of A is that A will believe she will ϕ not on the evidence of her telling but on the promise that her telling amounts to.⁸ The correct response to S’s telling is then to take S at her word and to see it as a promise to ϕ , or in Moran’s terms S’s *assurance* that she will ϕ . Were A to respond

⁷ While S can believe she will quit smoking even if she has tried to do this and failed before, S must recognise her ‘facticity’, which is to say any track record of failure, or other evidence that she will fail. However, these facts introduce practical considerations; they do not defeat the rationality of S’s belief, and so this need not force S to retreat to the belief that she will merely try to quit, see (Marušić 2015: 125-28). This account of practical belief is inspired by Velleman’s (1996) idea that we should distinguish beliefs that ‘create the truth’ and beliefs that ‘reflect the truth’.

⁸ Moran (2005: 23) quoted Marušić (2015: 182).

in this way, A would take the *participant stance* and *trust* S.⁹ To trust S would then be for A to regard S's testimony as offering an answer to the question of whether S will ϕ , and to accept S's answer as his own.

It follows, Marušić (2015: 198) claims, that in trust, “[t]he identity of the speaker and the nature of the speaker’s relationship to the addressee are of crucial importance if telling consists in offering an answer ... [because w]hen we are offered something it matters to us not just what is offered but also who is offering. For instance, if you are offered a ride, it very much matters who is offering — whether it is a stranger or an acquaintance, a friend or lover, a parent or supervisor”. The identity of the speaker is thereby an essential part of the reason that trust provides: in believing because of trust, A’s reason is *that S* told him. Trust based reasons are thereby different in kind to evidential reasons. And since “the weight of the reasons of trust does not correspond to their evidential weight, it is possible to rationally trust against the evidence” (Marušić 2015: 202). So, to continue the last example, S’s friend could believe her when she says that she will quit smoking while someone taking the observer’s perspective might not be able to do so, given S’s track record. Like practical belief, *trusting belief* is thereby contrasted with theoretical belief.

Marušić could then give the following account of the *murder case*. Given that it is your friend who tells you that she is innocent, and so offers her answer to the question of her innocence, you will be prone to trust your friend and so take her answer as your own.

⁹ The term ‘participant stance’ comes from Holton (1994), who takes it from Strawson (1974).

In believing your friend, you “maintain the participant stance and refuse to take the objective stance” (Marušić 2015: 202). And in doing this your belief is supported by the reason — that *your friend* told you that she was innocent — and reasonable because of this.

Suppose that this is all accepted. It follows that an explanation of the first datum is given: you can believe your friend simply because you resolutely stay within the participant stance and so trust your friend. And an explanation of the third datum is given: you are epistemically reasonable in believing your friend is innocent because this is a trusting belief. But what about the second datum? Does this account of the reasonableness of trusting belief explain how you can learn that your friend is innocent?

It seems not. On Marušić’s account it turns out that the acquisition of testimonial knowledge can be defeated *even when* the reasons got from trust render this testimonial belief reasonable. Suppose that the reason stated by ‘*my friend* told me that she is innocent’ is strong enough to fix testimonial belief and make it reasonable. The problem is that in the *murder case* the *balance of evidence* still points to your friend’s guilt. This point deserves emphasis: given that it is *your friend* who tells you that she is innocent, you have particularly good epistemic reasons to believe this — reasons based on your knowledge of your friend’s track record, her character, her values, that she cares for your friendship etc. — but the hypothesis is that no matter how good these epistemic reasons, the *balance of evidence* points to your friend’s guilt.¹⁰ And when this is so, on Marušić’s

¹⁰ As such, the reasonableness of your believing your friend cannot be explained by reference to the additional epistemic reasons that friendship provides, contrary to Goldberg (2017).

account, this balance of evidence prevents your learning from your friend that she is innocent: when “you have evidence that there is a significant chance of not- p , then you do not know that p through testimony, because this evidence is a defeater for your knowledge” (Marušić 2015: 189).

Does the evidence which defeats the acquisition of testimonial knowledge also defeat the epistemic reasonableness of trust based belief? Given that one can reasonably trust against the evidence, one might think not. However, trust isn't always reasonable: a reason of trust is a “*prima facie* conclusive reason” (Marušić 2015: 193, n.31) which is defeated by evidence of untrustworthiness.¹¹ In the *murder case*, the evidence of your friend's guilt implies that your friend is lying when she tells you that she is innocent. It follows that this evidence can equally be taken to be evidence of your friend's untrustworthiness. However, this kind of evidence of untrustworthiness would be present whenever there was counter-evidence to testimony, and it cannot be that *all* counter-evidence is evidence of untrustworthiness because Marušić wants to allow that trust can be reasonable in the face of the evidence. Thus, he states, “it is possible to rationally trust against the evidence. And if it is possible to rationally trust against the evidence, then we can rationally trust without having knowledge” (Marušić 2015: 202). It follows that the evidence in the *murder case* would normatively defeat the acquisition of testimonial knowledge even though it need not defeat the reasonableness of trust.

Thus, Marušić's account of the rationality of partiality, which develops Stroud's idea that friendship can make belief epistemically reasonable, explains two out of three of

¹¹ Hinchman (2014) equally takes trust to be a “*prima facie* conclusive reason”.

the starting data points. It fails to explain how it is that you can learn that your friend is innocent from her telling you this.

5.

How is it, then, that you can get to know that your friend is innocent from her telling you this? In this section, I want to consider the idea that knowledge and epistemic justification are sensitive to practical and moral matters. This will allow the argument that in the *murder case* the evidence pointing to your friend's guilt does not allow you to know that your friend is guilty; but the evidence of your friend's testimony does allow you to know that she is innocent.

Differences in *practical facts*, Jason Stanley argues — that is, differences in facts about the costs associated with one's beliefs being right or wrong — can constitute differences in knowledge. Knowledge is partly “a matter of *how much is at stake*” (Stanley 2004: 6). Stanley illustrates this, and argues for it, by reference to the well-known *bank case*.¹² You need to deposit a cheque and go to the bank after work on a Friday to find a long queue. You recall that the bank was open on a Saturday morning two weeks ago. There are then two cases to compare. In the *high stakes* case, you have a bill due and insufficient funds to cover it if your cheque is not deposited before Monday. In the *low stakes* case, you have no such worries. When considering these cases, the intuition is meant to be that you know the bank is open on Saturday in the low stakes

¹² Owed to DeRose (1992)

case, but do not know this in the high stakes case, even though your evidence is the same in both cases. This intuition, Stanley argues, then “provides a prima facie case for the thesis that knowledge is not just a matter of non-practical facts, but is also a matter of *how much is at stake*” (Stanley 2004: 6). The domains of theoretical reason and practical reason rather than being distinct, as Marušić argues, in fact encroach on one another.

This “prima face case” for *pragmatic encroachment* can be supported by reference to the principles linking knowledge and rational action.¹³ If you know that *p*, then you are practically justified in acting on *p*.¹⁴ And, it follows, if you are not practically justified in acting on *p*, then you don’t know that *p*. In the low stakes case, it would be practically rational to delay depositing the cheque (given the queue and your recollecting that the bank was open two Saturdays ago). So there is no barrier to supposing that you know that the bank is open this Saturday. However, in the high stakes case, it would be practically irrational to delay depositing the cheque until tomorrow (given the cost associated with a failure to deposit it). So it cannot be the case that you similarly know that the bank is open this Saturday. Thus, the idea is that practical facts encroach on our knowing things through raising and lowering the epistemic standards for knowledge.

Moreover, it is arguable that moral facts similarly encroach on knowledge and justification. Michael Pace (2011) illustrates this claim with the example of a shipbuilder

¹³ Term from Ichikama and Steup (2014)

¹⁴ Compare “*one should act only on what one knows*”, Stanley (2004: 9); and “If you know that *p*, then *p* is warranted enough to justify you in ϕ ing, for any ϕ ” Fantl and McGrath (2009: 67).

who, because of cutting costs in manufacture, has evidence, but not good evidence, that his latest vessel is seaworthy. Given the risk for the passengers of this belief being false, practical costs raise the standards for justification. The result is that the cost cutting means the shipbuilder no longer knows, as he should, that his vessel is seaworthy. But what, Pace asks, if the shipbuilder does not have the practical concerns that morally he ought to have? One would not want to conclude that he does, after all, know the vessel to be seaworthy. Rather, this is a case where moral considerations raise the epistemic bar. Equally, but conversely, “committing to be in certain relationships — e.g. becoming ... a friend — seems to generate moral obligations to have low evidential standards for certain propositions. These obligations are naturally described as ‘giving the benefit of the doubt’ or ‘thinking charitably of others’” (Pace 2011: 258).

Now consider the *murder case* in the light of these claims. Confronted with evidence of your friend’s guilt, the demands of friendship — be they moral demands or practical concerns — are such that this is a high stakes case. It is a high stakes case because if you don’t trust your friend when she is trustworthy, you risk your friendship. Your valuing your friendship and the obligations associated with friendship itself (if there are any) then make this risk a significant one. The epistemic requirement on knowing that your friend is guilty is thereby raised; and one might suppose it is raised to the point that you don’t know that your friend is guilty. But if your evidence does not put you in a position to know that your friend is guilty, then it is not true that you ought to believe this. In Pace’s terms (2011: 251), you should follow the *counsel of caution*: “don’t believe *p* so that you avoid believing a falsehood”; that is, you should rather *not* believe that she is guilty. It follows that balance of evidence does not give you a normative

defeater for the belief that she is innocent. So you can get to know that she is innocent on the basis of trusting her when she tells you this.

The problem with this strategy is that there must be limits on how high the demands of friendship can raise the epistemic bar. For example, the epistemic requirement could not be raised to certainty because to raise it so high would be to declare that practical facts and moral considerations can render observational truths unknowable. But if the upper limit falls below certainty, it will always be possible to imagine a case where there is sufficient evidence to cross the raised threshold for knowledge, and yet the evidence still be misleading. The problem is then: this is what could be imagined in the *murder case*. To add some colour consider the *Hollywood murder case* which is the same as the *murder case* except that it could be a plot written for Hollywood, which includes a powerful and nefarious government agency setting out to frame your friend. It follows that the evidence of your friend's guilt might be so stacked that even given the demands of friendship the balance of evidence would allow you to cross whatever epistemic threshold is needed for knowledge. It follows that you ought to believe that your friend is guilty and so the possibility of learning that your friend is innocent is normatively defeated. As such, while this strategy might be able to explain the three data points in some cases, it cannot explain how you can learn that your friend is innocent in the beefed-up *Hollywood murder case*.

6.

In this section and the next, I propose a solution to the starting problematic. Friendship, I will argue, provides practical reasons for you to reconsider the question of your friend's guilt. And it motivates you to deliberate over this question until the point that you are epistemically justified in suspending your belief that your friend is guilty. Insofar as this belief has been justifiably put to one side, it ceases to be the case that you ought to believe that your friend is guilty. As such, in trusting your friend when she tells that she is innocent, you are in a position where you can learn that she is innocent.

To develop this proposal, let me return to the challenge facing Marušić's position. To learn that your friend is innocent, Marušić (2015: 179) claims, the reason given by her telling "would have to *defeat* all [the] evidence" you have that she is guilty. Suppose that this is taken to be a requirement on your justifiably suspending your belief that your friend is guilty: that your friend's testimony provide sufficient defeating reason. That is, suppose that *one is justified in suspending belief only if one has sufficient defeating reason*. To propose this requirement on the justified suspension of belief would be to introduce, what Matthew Soteriou (2013: 360) calls, *epistemic traps*. These are situations where one is unjustified in belief — because one forms the belief in an epistemically unjustified way — but possessing no counter-evidence one is nevertheless unjustified in abandoning this unjustified belief. When epistemically trapped, the epistemic demands conflict: one's belief is unjustified but one is not justified in abandoning it. The murder case itself is not an epistemic trap since you are justified in believing your friend is guilty. Nevertheless, this justification is not knowledge supporting, since it based on misleading evidence, so it is a case where there is a justificatory failure, but where, on the

condition proposed, suspending belief is still unjustified. Thus, like the possibility of epistemic traps, it motivates Soteriou's proposal of a different account of when the suspension of belief is justified.

“S is justified in suspending judgement over p ’ can be interpreted as meaning different things”, Soteriou (2013: 363) says,

either (a) given S’s evidential situation — i.e. given the evidence S does and does not possess — S ought to suspend judgement over p , or (b) S has attempted to determine whether p and she has not (yet) done anything that justifies her judging that p or judging that not- p .

The first interpretation, (a), applies to situations where S doesn’t already believe that p . For instance, I doubt that anyone believes that there are an odd number of ducks, and per (a), one ought to suspend judgement over this. The second interpretation, (b), applies to situations where S already believes that p , but then withdraws assent from p , suspending belief. To illustrate this situation, consider the *case of the student mathematician*. From a recent lecture S knows that p and is given the exercise of proving that p . In attempting this proof, S brackets her knowledge that p and proceeds to try and work out whether p . Suppose she fails to come up with a proof and after long frustration at trying to come up with one she begins to doubt that she copied the theorem down correctly. In this position, the two routes available to her for justifiably judging that p — namely the production of a successful proof, or recollecting what the lecturer said — are both blocked. In this position, she cannot do anything to justifiably judge that p , and is thereby justified in suspending judgement over p .

Consider, then, the application of this theory of when the suspension of belief is justified to the *bank case* and the *murder case*. In the two bank cases, one knows that the bank is open on Saturday because one can recall going there a couple of weekends ago. However, if one brackets this belief and tries to work out whether the bank is open on Saturday, as one does in the high stakes case, this recollection is undermined by the thought that Saturday open hours are irregular.¹⁵ One has no information about the opening hours *this* Saturday, so nothing now allows one to justifiably judge that the bank is open tomorrow. In this position, one is justified in suspending one's belief that the bank will open tomorrow. As such, Soteriou (2013: 367) proposes that

we should not regard such cases as ones in which the epistemic standards required for knowledge have shifted. Rather, we should regard them as cases in which two options are epistemically available to the subject. Both suspending judgement over *p* and continuing to believe that *p* are epistemically permissible. The differences in practical circumstances between the two situations affect which epistemic option it is more *practically* rational for the subject to take.

The suggestion here is not that practical facts encroach upon epistemic justification (as was proposed in section 5) but that the *bank case* is one where there are epistemic options — continue to believe the bank is open on Saturday or suspend this belief — and practical facts then determine which epistemic option it is reasonable to take.¹⁶

¹⁵ Or, as Stanley (2004: 4) has it, by one's partner pointing out that "banks do change their hours".

¹⁶ Similarly, see Brown (2008).

The same then goes for the *murder case*. Given all the misleading evidence, you are justified in believing that your friend is guilty of murder. But because of friendship when your friend protests her innocence, you reconsider whether she is guilty or not. This sees you bracketing your belief that she is guilty and attempting to work out whether she is. In this situation, it might be that other explanations of the evidence can be given such that you can conclude that nothing allows you to justifiably judge that she is guilty. Suppose for the moment this is so, and that the *murder case* is just like the *bank case*. It follows that you would be justified in suspending your belief that your friend is guilty, and, given your deliberation, it is thereby no longer true that you ought to believe her guilty. That is, insofar you are *justified* in suspending your belief that she is guilty, where this belief was based on all the evidence against your friend, there is no further normative defeater that would block your learning that she is innocent. So in trusting her when she tells you this, you could get to know that she is innocent.

The *murder case* is then different from the *bank case* because it can be imagined in a ‘Hollywood’ style, where the evidence of your friend’s guilt is very convincing. With respect to this *Hollywood murder case*, when you bracket your belief that your friend is guilty and consider whether she is, the evidence upon which you originally acquired your belief in her guilt might be such that you can think of no other plausible alternative explanation of it than your friend’s guilt. So even after reflection, you would still be in a position of being able to justifiably judge your friend guilty, and would thereby be unjustified in abandoning this belief. However, reflection does not need to end at this point. Reflection can be extended. Even if you conclude that the evidence does allow you to judge that your friend is guilty, you can then bracket your belief in the evidence, or

some part of it, and reconsider what grounds you have for believing some bit of evidence holds true. In this regard, Soteriou (2013: 368-9) says,

Even if, upon reflection, one realizes that one is justified in judging that one does not know that p , it does not follow that one does not know that p . For suspending judgement over q (when the evidence that q may have played a role in the reasoning leading to one's justified judgement that one does not know that p) may be an epistemically permissible option.

Thus, it seems that through critically reflecting on whether your friend is guilty and the evidence you have for this, it is possible for a route to open to the epistemically justified suspension of your belief that she is guilty. Insofar as this is possible, the conclusion of the previous paragraph can be restated: were you able to justifiably suspend this belief, there would be no further normative defeater blocking your learning from your friend that she is innocent.¹⁷

The obvious worry about this proposal concerns where the epistemic limits lie. There must be limits as to when we can justifiably suspend belief, and so there must be

¹⁷ This account of when the suspension of belief is justified entails that different beliefs can be epistemically reasonable on the same body of evidence. See Kelly (2008) for further support. Hawley (2014: 2040) then argues that such “epistemic permissiveness” — like Pace (2011), she focuses on how different people differently “strike the balance between pursuing true belief and avoiding false belief” — “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.” This section then elaborates this claim through explaining how believing a friend in the face of the evidence can be epistemically reasonable.

versions of the *Hollywood murder case* where the possibility of your learning of your friend's innocence is normatively defeated. Even when arguing that friendship involves departures from epistemic rationality as it is standardly conceived, Stroud (2006: 506) observes that friendship cannot require "wilfully believing the false" or "involve a total imperviousness to damning evidence". The question is then: where do these limits lie? And one answer might be: one could not justifiably suspend one's belief that p if one knows that not- p . It would follow that one could not believe the friend who told one something one knew to be false. However, this condition would not change the result in the *murder case* since you do not know that your friend is guilty, and cannot know this since it is false. Moreover, this is not where the limits lie as is illustrated by the *case of the student mathematician*. This student starts from a position of knowing that p but is nevertheless justified in suspending her belief that p .¹⁸

A stronger proposal would be that the epistemic limits lie at the point of *epistemic obligation*. For instance, if the evidence against your friend were such that it obliged you to believe that she is guilty, then were you to bracket this belief and consider whether she is guilty, the reflective result would be that you are epistemically obliged to judge that she is so. Withdrawing assent from your belief that your friend is guilty would then not be justified and the route to learning of your friend's innocence would be normatively defeated.

¹⁸ In addition, to put the limit at this point would be to succumb to the *dogmatism paradox*, or the claim that if one knows that p , one can ignore all evidence that not p . See Lewis (1996) and Soteriou (2013: 364).

However, one is rarely under epistemic obligation. Observational evidence is fallible, and so consistent with the falsity of what it is evidence for. And it would be a sorry state of affairs if you were epistemically obligated to believe the falsehood that your friend is guilty. To be epistemically obligated to believe a falsehood would be akin to being caught in an epistemic trap: it would amount to some deep-rooted failure of justification. But insofar as the evidence does not oblige belief, there will be a route available that ends in the justified suspension of belief. The more extravagant the *murder case* is imagined to be — the more like a Hollywood movie it is — the greater the degree of reflection that would be needed to open this route up. For this reason, your friend turns to you — *her friend* — since only friendship, with its accompanying demands, would make it practically rational to engage in this kind of reflective endeavour. But on the assumption that the evidence doesn't oblige belief, as it is very unlikely to, the reward of this reflective endeavour would be that you could justifiably abandon your belief that your friend is guilty. And since this belief was based on all the evidence against your friend, there is no further normative defeater that would prevent you learning of her innocence in trusting her when she tells you this.

7.

The last section focussed on the issue of normative defeat. It sought to explain how it need not be the case that normative defeat prevents your learning your friend is innocent. However, to argue so much is not yet to show that you can learn that your friend is innocent. What is needed for this is some positive statement as to how you can believe

that your friend is innocent and be epistemically reasonable in believing this. Once this positive account is in place, all three of the starting data-points will then be accounted for.

The account I would propose, (Faulkner 2011), as to why you believe that your friend is innocent and can be epistemically reasonable in doing so is similar in key respects to that offered by Marušić (2015) and outlined in section 4. In particular, the key idea is that your belief that your friend is innocent arises through taking the *participant stance* and so *trusting your friend*. As I would characterise things, to trust your friend for the truth is to believe what your friend tells you, and to believe what she tells you just *because* you trust her for the truth in this matter. Your having this attitude of trust involves taking a participant stance because it involves viewing your friend's testimony in the way that your friend expects you to view it, which is as offering her assurance that what she tells you is true. That you can, in trust, view your friend's testimony this way explains how belief is psychologically possible. This is the first data-point: belief is psychologically possible because trust is.

Your attitude of trusting your friend for the truth explains your willingness to view her telling as assurance because trust involves you making the presumption that your friend tells you that she is innocent just because she knows that this is so. Insofar as trust involves this presumption your belief that she is innocent is thereby epistemically rationalised by your attitude of trust, (Faulkner 2011: 150-9). Moreover, in trusting your friend for the truth, and so making this presumption, you would thereby recognise your friend's epistemic authority and so regard her telling you that she is innocent as giving you a *preemptive* reason to believe this, which is also a reason to “disregard” any

counter-evidence, (Keren 2007: 373). To “disregard” the counter-evidence, I suggest, is to suspend belief in it. Given that your friend knows what she tells you, whether you get to know too then depends on whether your disregarding the counter-evidence is justified. And this hinges on whether you engage in the kind of deliberative process described in the last section. Thus, and in sum, the possibility of trust then further explains how belief can be epistemically reasonable, and the possibility of justified suspension of judgement explains how the knowledge that you can acquire on trust need not be defeated. These are the other two data-points.

In conclusion, the final question to consider is what role friendship plays in these explanations. You believe that your friend is innocent because you trust her for the truth in this matter. And your belief that she is innocent is epistemically reasonable because you trust her the truth in this matter. Trust plays both a causal and justificatory role. Friendship figures in these explanations — it is your friend that you trust — but it does not figure essentially: what makes belief possible and epistemically reasonable are facts to do with trust rather than friendship. Specifically, and put in general terms, it is that in trusting a speaker who tells one that p one makes the presumption that the speaker tells one that p because they know that p . For this reason, in trusting a speaker one views their telling as the assurance it purports to be. On the speaker’s side the same is true: friendship figures in the explanation — it is you, her friend, that she tells — but does not figure essentially. The expectation that your friend has of you — that you will trust her for the truth, and so presume she knows what she says, and so view her telling as giving you a distinctive reason for belief — is the general expectation that one has, as a speaker, in telling someone something.

Then with respect to your getting to know that your friend is innocent, what makes this possible is your being justified in suspending any belief in your friend's guilt as a result of rationally deliberating whether your friend is guilty. Friendship figures in this explanation — it is the demands of friendship that prompt your deliberation — but it does not figure essentially insofar as what makes the justified suspension of belief possible are general facts about the agency we enjoy in thinking. Maybe the most that could be said is this: friendship is *practically necessary* because while critical reflection on the question of your friend's guilt is open to any audience, it is only because of friendship that you are willing to push your deliberation to the point of being justified in giving your friend the benefit of the doubt. Equally, it is only because of friendship that your friend can expect this of you. For this reason, your friend seeks you out, her friend, because there is no point in telling another that she is innocent if they could not get themselves into a position where they could learn from her that this is so because what your friend wants, crucially, is for someone else to know of her innocence.¹⁹

Audi, R. (1997). "The Place of Testimony in the Fabric of Knowledge and Justification", *American Philosophical Quarterly*, 34(4): 405-422.

Baker, J. (1987). "Trust and Rationality" *Pacific Philosophical Quarterly*, 68: 1-13.

¹⁹ Thanks to Sandy Goldberg, Arnon Keren, Matt Soteriou, Bob Stern and audiences at Universities of Frankfurt, Haifa, London and at the 2017 Summer School in Social Epistemology.

- Brown, J. (2008). "Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning", *Noûs* 42(2): 167-189.
- DeRose, K. (1992). "Contextualism and Knowledge Attributions", *Philosophy and Phenomenological Research*, 52(4): 913-929.
- Fantl, J., & McGrath, M. (2009). *Knowledge in an Uncertain World*. (Oxford: Oxford University Press).
- Faulkner, P. (2007). "On Telling and Trusting", *Mind* 116(464): 875-902.
- . 2011. *Knowledge on Trust*. (Oxford: Oxford University Press).
- Fricker, M. (2012). "Group Testimony? The Making of a Collective Good Informant", *Philosophy and Phenomenological Research* 84(2): 249-276.
- Goldberg, S. (2017). "Against Epistemic Partiality in Friendship: Value Reflecting Reasons", *unpublished paper*.
- Hawley, K. (2014). "Partiality and Prejudice in Trusting", *Synthese* 191: 2029-45.
- Hieronymi, P. (2008). "The Reasons of Trust", *Australasian Journal of Philosophy*, 86(2): 213-236.
- Hinchman, E. (2005). "Telling as Inviting to Trust", *Philosophy and Phenomenological Research* 70(3): 562-587.
- . (2014). "Assurance and Warrant", *Philosophers' Imprint*, 14: 1-58.

- Holton, R. (1994). "Deciding to Trust, Coming to Believe," *Australasian Journal of Philosophy*, 72(1): 63-76.
- Ichikama, J. J., & Steup, M. (2014). "The Analysis of Knowledge", in E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy (Spring 2014 Edition)*.
<<http://plato.stanford.edu/archives/spr2014/entries/knowledge-analysis/>>.
- Kawall, J. (2013), "Friendship and Epistemic Norms", *Philosophical Studies*, 165: 349-70.
- Keller, S. (2004). "Friendship and Belief", *Philosophical Papers*, 33(3): 329-51.
- Kelly, T. (2008). "Disagreement, Dogmatism and Belief Polarization", *The Journal of Philosophy*, 105(10): 611-633.
- Keren, A. (2007). "Epistemic Authority, Testimony and the Transmission of Knowledge", *Episteme*, 4(3): 368-381.
- Lewis, D. (1996). "Elusive Knowledge", *Australasian Journal of Philosophy*, 74: 549-567.
- Marušić, B. (2015). *Evidence and Agency: Norms of Belief for Promising and Resolving*. (Oxford: Oxford University Press).
- Moran, R. (2005). "Getting Told and Being Believed", *Philosophers' Imprint*, 5(5): 1-29.
- Pace, M. (2011). "The Epistemic Value of Moral Considerations: Justification, Moral Encorachment, and James' 'Will to Believe'", *Noûs*, 45: 239-268.
- Ross, A. (1986). "Why Do We Believe What We Are Told?", *Ratio* 28(1): 69-88.

Soteriou, M. (2013). *The Mind's Construction: The Ontology of Mind and Mental Action*.
(Oxford: Oxford University Press).

Stanley, J. (2004). *Knowledge and Practical Interests*. (Oxford: Oxford University
Press).

Strawson, P.F. (2008). "Freedom and Resentment", in Strawson (Ed.), *Freedom and
Resentment and Other Essays*, (London: Routledge), 1-28.

Stroud, S. (2006). "Epistemic Partiality in Friendship", *Ethics*, 116: 498-524.

Velleman, D. (1996). "The Possibility of Practical Reason", *Ethics*, 106(4): 694-726.

Vendler, Z. (1979). "Telling the Facts", in French (Ed.), *Contemporary Perspectives in
the Philosophy of Language, Midwest Studies in Philosophy*, (1979): 220-232.