



Trusting What Ought to Happen

Jonathan Tallant¹ 

Received: 24 April 2019 / Accepted: 8 August 2022
© The Author(s) 2022

Abstract

This paper introduces a new account of trust and distrust. The core aim of the paper is to introduce an account of trust that places trust and what ‘ought’ to happen as close conceptual companions. Over the course of the paper, I develop the account and compare it with certain rival accounts.

This paper argues for a new account of trust and distrust. With a small amendment to be introduced later, the central claim to be argued for is that:

x trusts y to ϕ if and only if x believes that y ought to ϕ , and relies upon y to ϕ .

An account of distrust is also offered.

1 Introduction and Background

The kind of trust that I have in my sights in this paper is 3-placed trust; trust of the form x trusts y to ϕ . This is to be contrasted with related concepts, such as 2-placed trust— x trusts y —and trustworthiness. The paper focuses on 3-placed trust, though also offers up a theory of distrust; specifically, a theory of what it is for x to distrust y to ϕ .

✉ Jonathan Tallant
jonathan.tallant@nottingham.ac.uk

¹ Department of Philosophy, University of Nottingham, Nottingham, United Kingdom

There has been a good deal written about 3-placed trust of late.¹ In adding to this literature, I wish to defend a novel analysis of 3-placed trust. Specifically, the following view will be developed (with a small amendment to be added a little later on):

x trusts y to ϕ if and only if x believes that y ought to ϕ , and relies upon y to ϕ .

To illustrate: Hilary trusts Nasser to look after his plants iff Hilary believes that Nasser ought to look after his plants, and relies upon Nasser to do so.

In the discussions that follow, ‘ought’ will take up a chunk of our time, so little needs to be said about it here.² However, to illustrate the general idea: once Nasser has promised to water the plants, he ought to water them (*ceteris paribus*). Given this analysis, just so long as Hilary is also relying upon Nasser to look after his plants, Hilary trusts Nasser to do so.

The stepping-off point for this paper is a 2014 work by Katherine Hawley.³ Hawley’s general line of argument is that her commitment account of trust is the account best-placed to make sense of the close connection between trust (on the one hand) and distrust (on the other). I will not repeat her arguments. Hawley’s arguments show that the accounts of trust offered by the likes of Baier (1986), Hardin (2002), Holton (1994), and Jones (1996, 2004) are at a disadvantage from the get-go and, as such, I don’t need a compare-and-contrast style argument to show that my Ought Account is to be preferred. Instead, I accept Hawley’s conclusion and then show that the Ought Account is an improvement on Hawley’s.

As such, I will show that the Ought Account performs as well as Hawley’s commitment account when it comes to explaining the close connection between trust and distrust. I’ll then move on to showing how three arguments (Unification (Sect. 3), Motive (Sect. 4) and ineffective commitment (Sect. 5)) leave the Ought Account better placed than Hawley’s commitment account.

If we’re to see Hawley’s account of trust as the ‘field leader’, then something needs to be said about exactly how it is that her account successfully accounts for the close connection between trust and distrust, which is central to her assessment of the commitment account. To help get a feel for Hawley’s view, Sect. 1.1 walks through her reasons for rejecting rival accounts (giving a quick example to illustrate her strategy), before (1.2) moving on to her preferred account.

1.1 Hawley’s On the Importance of Distrust

Hawley (2014: 1) begins by noting that to come to a full and proper understanding of trust, we must also come to an understanding of distrust. Hawley then notes that, typically, distrust is treated as a second-class citizen or else misunderstood altogether.

From this point, Hawley’s approach is to show that existing accounts of trust do not allow us to properly account for the distinction between trust and distrust. To give the reader just a sense of some of Hawley’s reasoning, we can consider her discus-

¹ A non-exhaustive list of items discussing the notion includes Baier (1986), Faulkner (2007), Hardin (2002), Jones (1996, 2004), Hawley (2014), and Tallant (2017).

² See Sect. 3 onwards for the bulk of the discussion.

³ As many readers will know, Katherine passed in 2021 and she is sorely missed. Friends and colleagues offer remembrances of her here: <https://katherinehawley.wp.st-andrews.ac.uk>.

sion of Hardin's (2002) account of trust. Hawley begins with the following account of Hardin's view of trust.

Hardin argues that when we trust someone, we expect the trustee to encapsulate our interests within her own because she has an interest in maintaining or strengthening her relationship with us. In trusting you to look after my vase, I take it that you will do so because you have incorporated my interest in preserving the vase amongst your own interests: looking after the vase is now in your own interest. (2014: 5)

As she goes on to point out, however, this account looks like it will face problems when we try to deal with distrust. The most obvious way to make sense of distrust on this model of trust would be to treat it as nonreliance and a negative attitude towards the motives of the distrustee. But, thinks Hawley, such an account will not do. Hawley invites us to consider a case in which an honourable person campaigns to have her (Hawley) imprisoned for her imagined crimes. Hawley could not rely upon that person to help her. She might also know that the honourable person bears her ill-will and is doing their best to frustrate her goals. Nonetheless, it need not be the case that Hawley distrusts her honourable foe; her foe is entirely honest and above board in their efforts (2014: 5–6).

This motivates Hawley's rejection of Hardin's account.⁴ Her general approach when considering opposing views is largely similar: identify the account of trust and explain the difficulties that we will face in developing an appropriate account of distrust.

1.2 Hawley On Trust

Hawley's own view of trust and distrust then comes to this:

To trust someone to do something is to believe that she has a commitment to doing it, and to rely upon her to meet that commitment. To distrust someone to do something is to believe that she has a commitment to doing it, and yet not rely upon her to meet that commitment (2014: 11).⁵

Thus, and to return to my earlier example: Hilary trusts Nasser to water his plants iff Hilary believes Nasser has a commitment to doing watering Hilary's plants, and relies upon Nasser to look after the plants.

We now have in hand Hawley's account of both trust and distrust, and at least an outline sense of why some opposing views should be rejected.

⁴ To be clear, Hawley has a slew of arguments against Hardin's position, of which this is only one. I do not claim to reproduce them all here. I offer one argument simply to give the reader a feel for the kinds of argument that Hawley offers.

⁵ 'Commitment' here is to be understood such that 'commitments can be implicit or explicit, weighty or trivial, conferred by roles and external circumstances, default or acquired, welcome or unwelcome' (Hawley, 2014: 11). In short, and as she acknowledges, (2014: 11) Hawley takes a very broad conception of commitment.

2 Counter Examples and the Ought Account

As powerful as Hawley's account is, it fails. In Sect. 2.1 I introduce a counterexample already available in the literature. I draw out some specific features of this *type* of counterexample. I refer to counter examples of this type as 'No reliance' counterexamples. In 2.2 I introduce a novel kind of counterexample that I refer to as 'diachronic'. In Sect. 2.3 I re-introduce my preferred Ought Account, and explain how (suitably modified) it can handle these two different kinds of counter example. In 2.4 show how the Ought Account can deliver an account of distrust.

2.1 No Reliance Cases

What I call 'No reliance' counter examples have already been introduced in the literature. Tallant (2017) invites us to consider a case in which a shopper, Hilary, is seeking out a Christmas present. Hilary considers both online retailers and a local, in-person store. In the case, Hilary is depicted as someone who knows the online retailer well. Hilary believes that the staff of the online retailer have good intentions and also that they *will* deliver in time for Christmas, if Hilary chooses to place an order with them. Nonetheless, in the scenario described, Hilary does not make the purchase with the online retailer. Rather, Hilary makes the purchase from a local, in-person store out of a desire to support local companies and to save a little money. This is a scenario in which Hilary believes that the online store and its employees have a commitment to delivering in time for Christmas, but Hilary does not rely upon them to do so.

As Tallant points out, if we use Hawley's account of distrust, this yields the result that Hilary distrusts the company to deliver the present in time for Christmas; Hilary recognises the intention of the company to deliver Christmas presents in time for Christmas, but *does not rely upon them to do so*. (Thus my referring to this kind of counter example as a 'no reliance' case.) But of course, it does not follow that Hilary *distrusts* the company to deliver in time for Christmas. For reasons wholly disassociated with her trusting of the company, and wholly to do with a concern for supporting the local economy, Hilary might end up taking the decision to make the relevant purchase elsewhere.

To deal with this case, Tallant considers the following modest revision to Hawley's position:

To trust someone to do something is to believe that she has a commitment to doing it, and: to either rely upon her to meet that commitment, or, if they are not called upon to meet that commitment, to believe that they could be relied upon to meet that commitment, if so called upon

To distrust someone to do something is to believe that she has a commitment to doing it, and yet not rely upon her to meet that commitment *because* of reservations about whether or not they can be relied upon to do it.

What this does for the Commitment account of trust is allow that trust can be present even if it is not called upon, and that distrust requires there to be a close connection

between not relying upon an agent to meet their commitments and concerns about whether the agent *can* be relied upon.

2.2 Diachronic Cases

Here is a different, novel kind of counter-example.

Elizabeth has important guests coming over to her home for a work dinner. She strongly suspects that future promotion prospects may turn on how well the evening goes. She has ordered food to be delivered. She has no time to make alternative arrangements. At the point she submitted the order, she thought that the company from whom she'd ordered had a commitment to timely delivery and she relied upon them to deliver on time, firmly believing them to be competent. However, since placing the order, Elizabeth has been informed via a text from a friend of the company's dire track record. They are well-meaning and sincere, but incompetent. She acknowledges that they have a commitment to delivering on time but now firmly believes that they will not do so. However, so much time has passed that Elizabeth does not pursue another course of action; she simply *hopes* that the company manages to deliver, despite now holding this belief that they will not do so.

The commitment accounts (either that defended by Hawley, or the refinement due to Tallant) require that, even though she now believes that the company will *not* deliver in time for dinner, Elizabeth trusts the company to deliver on time. After all, Elizabeth believes that they have a commitment to delivering on time and she is relying upon them to do so (she perhaps could make other plans, but has resigned herself to her fate).

But this seems quite wrong. Elizabeth started out, at time t , trusting the company. A little later, at t^* , she became aware of information that led her to believe that the company is incompetent and now, a little later still, believes that they will not deliver on time—they will not deliver by t^{**} . In that case, Elizabeth does not trust the delivery service to deliver on time even though she is relying upon them to deliver on time and believes that they have a commitment to doing so. She now has additional information. Elizabeth believes that their incompetence is such that they will not deliver on their promise. It's quite reasonable to think that Elizabeth still believes that the company have a commitment to delivering. She does not doubt their commitment, after all, merely their competence.

Structurally, this case reveals that the commitment accounts wrongly 'lock us in'. According to the commitment accounts, once we start to trust, provided we continue relying on the trustee, and provided we do not revise our views of what they have a commitment to doing, we *continue* to trust them, even in the face of evidence that leads to our forming the belief that they are not to be trusted—that they should not be relied upon. Psychologically, that seems implausible in an account of trust. We *can* change in the way I describe Elizabeth changing. We can shift from trusting to

not trusting in the fashion described. This is the diachronic problem for accounts of trust.⁶

2.3 Patching the ‘Ought’ Account

The Ought Account was introduced as follows.

x trusts y to ϕ if and only if x believes that y ought to ϕ , and relies upon y to ϕ .

The Ought Account thus *appears* prone to the same kinds of concern that Tallant thinks infect Hawley’s position. It *seems* prone to No Reliance cases. Hilary believes that the company *ought* to deliver presents in time for Christmas, but does not rely upon them to do so. The Ought Account also *appears* prone to diachronic cases. Elizabeth believes that the company ought to deliver on time and is relying upon them to do so. Nonetheless, she does not trust them. Thus, the Ought Account will require amendment.

In the No Reliance case, what is relevant to whether Hilary trusts the company to deliver in time for Christmas is not whether she *does* rely upon them, but whether her view of the company is such that she believes that were they called upon to deliver on time, then they would do so. In the diachronic case, what is relevant to whether Elizabeth trusts the company once she has discovered the shockingly bad delivery record is that, were she to now reconsider her original decision in light of her new information, it’s not the case that she would believe that they would deliver on time.

If that is correct, then it suggests that trust turns, not on whether we *do* rely upon other agents to perform as we wish, but on whether we believe that those agents *would* perform as we wish, *were we* to (now) take the decision to rely upon them in the manner specified. This general lesson yields the following, refined version of the Ought Account:

x trusts y to ϕ if and only if x believes that y ought to ϕ , and, believes that, were y called upon to ϕ , then y would ϕ ;

Adopting that approach, we correctly return the verdict that Hilary trusts the delivery company to deliver in time for Christmas in the No-Reliance case. She believes that they ought to deliver on time and believes that, were they called upon to do so, then they would deliver on time.

In the diachronic case, Elizabeth can correctly be said to *not* trust the delivery company at the later time. She does not (at the later time) believe that, were they called upon to deliver on time, then they would do so. Thus, she does not trust then. So far, so good.

2.4 Distrust

As we saw in 1.1, we require an account of distrust. A natural suggestion given the Ought Account is:

⁶ Notably, Elizabeth does have another option: she could cancel the event. To be sure, this will have negative consequences for her. But she does still have a choice to continue as described.

x distrusts y to ϕ if and only if x believes that y ought to ϕ , and, believes that were y called upon to ϕ , then y would not ϕ ;

But that won't quite do.

Suppose that Jack believes that Marcy should not lie about having poisoned Sam. However, Marcy has told Jack that she will lie about it—Marcy does not wish to go to prison. It seems wrong here to say that Jack distrusts Marcy to tell the truth. He knows full well that she will not. There should be no distrust here. Rather, this is a case where we simply *lack* trust. Nonetheless, Jack believes that: (i) Marcy ought to tell the truth (she ought not to lie) about having poisoned Sam; (ii) Marcy will lie if asked about having poisoned Sam. The analysis therefore wrongly returns the result that Jack distrusts Marcy. This seems the wrong result.

This suggests a modification to the account of distrust:

x distrusts y to ϕ if and only if x believes that y ought to ϕ , and believes that were y called upon to ϕ , it is not clear whether y would ϕ .

This delivers the correct account in the poisoning case. Jack does not distrust Marcy. He believes that it is clear that she will lie. If he believes that it is clear that Marcy will lie, then it is not unclear to him whether she will lie; Marcy is, thus, not distrusted in this case.

Before moving on, we should draw out one feature of this view of distrust—the element of uncertainty. It is built into the analysis that distrusting someone to ϕ requires that we not be clear whether they will ϕ . And this strikes me as quite right—at least, the phenomenology does. If I believe that it's clear that an agent will ϕ , then I might trust them to do it. I might also not trust them to do it. However, if I am distrustful of them, then (or so it seems to me) that brings with it a sense of uncertainty—a creeping sense of unease. A 'Will they? Won't they?' phenomenology. This seems to me to be *essential* to distrust. If I am certain that you will not buy me flowers, then I do not distrust you to buy me flowers. That being so, quite apart from getting the right result in the specific case, the analysis delivers against the phenomenology in a way that we should find appealing.

In closing this section, I should note that I do not have a view as to whether the commitment account can also be modified to deal with the considerations raised here. Arguments to prefer the Ought Account to the commitment account are deferred to Sect. 3. Before we move to that, however, I want to consider the *difference* between the Ought Account and the commitment account.

2.5 Commitment vs. Obligation

A key claim in this paper is that obligation is key to trust. Hawley says that commitment is key to trust. Obligations and commitments are closely related notions. Are they so close that we can get no purchase on the difference?

No. We can appreciate the difference between obligations and commitments. In exploring this matter, we can also respond to an argument due to Hawley against a version of the Ought Account. This will pave the way for Sect. 3, where we can turn

the tables and show that the Ought Account is to be preferred to the commitment account.

As Hawley (2014: 18) notes, commitments *typically* bring obligations. If I have a commitment to performing an action, then most of the time I will acquire an obligation to do the same. However, as Hawley also points out, there are cases where we have obligations but no commitments. Hawley takes this to favour the commitment account, since she believes that this generates the correct result in the specific case that she considers.

I agree that there is clear water between obligations and commitments. I disagree with Hawley that consideration of the case tells in favour of the commitment account.

Here is the case, borrowed from Hawley:

Suppose we are in the Wild West. In town, there is an uneasy truce, a semblance of law and order. Out in the desert, as everyone knows, there are no holds barred; you take your life into your hands if you venture there. You and I meet by chance in the desert, and you see that I am armed (of course). Before we exchange words, you may try to predict whether or not I will let you live. But no matter how confident you are in that prediction, it's not appropriate for you either to trust or to distrust me in this respect. Out here, no one has promised, implicitly or explicitly, to respect life and limb. I have no commitment to let you live, and you should neither trust or distrust me in this respect. If I let you live, I do not demonstrate trustworthiness; if I shoot you, I display many vices, but not that of untrustworthiness. But of course I still have an obligation to let you live, whether or not I recognise this. (2014: 18–9)

Hawley is correct that we have an obligation, but not a commitment. In the circumstances, Hawley has an obligation to let her armed opponent live, regardless of whether she recognises that obligation.

However, the implication that Hawley seeks to take from this is that, if trust were to be tied to obligations, then we thereby get the case wrong. This does not follow from the argument offered.

As she puts matters, emphasis mine:

it's not appropriate for you either to trust or to distrust me in this respect. Out here, no one has promised, implicitly or explicitly, to respect life and limb. I have no commitment to let you live, and you *should* neither trust or distrust me in this respect.

Let us grant Hawley is correct about exactly what is quoted. Trust would not be *appropriate* here.

Crucially, however, the propriety of trust isn't salient to whether trust *is granted*. We can agree with Hawley that no appropriate trust can be given in this situation, but that does not show that no trust *is* given. The question of whether Hawley's armed opponent *should* trust her isn't germane to the conceptual analysis of whether they *do*. Once we dig into the case, we can exploit this fact to demonstrate that what's doing the work here is the *beliefs* in play.

Let us hold fixed all of the details of Hawley's case, but add the supposition that Hawley's armed opponent forms the belief that Hawley has a commitment not to shoot her. This is caused by being a fellow-traveller in this wilderness. To be sure, no promises have been made (implicit or otherwise). Nonetheless, people sometimes form beliefs that they should not. Perhaps Hawley's opponent simply believes that fellow travellers in the Wild West *shouldn't* shoot one another.

Suppose, then, that the armed opponent acts on that belief, not drawing their weapon. The armed opponent thereby *trusts* Hawley, even given her own account, for they believe that she has a commitment to not shoot and they rely upon Hawley not to do so. Equally, they *may* believe that they ought to not be shot by Hawley (presumably, if they believe that she has a moral commitment to not shoot, then they believe that she ought not to shoot, after all).

However, it is equally important to note that the imagined armed opponent may have none of those beliefs. They may—as is implied in the original case—*not* form those beliefs about what Hawley is committed to. There is also nothing conceptually problematic about *that*. And, if they don't form those beliefs, then they don't trust her not to shoot. For, given the analysis, they don't believe that she has the relevant commitments.

What this case thus makes plain is that whether the commitment account generates the result that the armed opponent trusts Hawley turns on whether they form particular beliefs about her commitments. But it is entirely reasonable that we make the same point with respect to obligations. If, for instance, the armed opponent *believes* that Hawley has an obligation not to shoot, and then relies on her not to do so, then there is nothing wrong with saying that they trusted her (*mutatis mutandis* if they do not form that belief). Thus, though Hawley is correct that commitments and obligations come apart, what is doing the work in this case is the presence of (or otherwise) of relevant beliefs.

3 On Obligation

Before shifting to the arguments in favour of the Ought Account, I must say something about obligation. Obligation, as expressed by 'ought' is a key part of the puzzle. In adopting the proposed Ought Account, I don't intend to require the adoption of any particular theory of the nature of obligation, or of what (if anything) grounds or moral obligations. I hope to remain theoretically neutral. Nonetheless, I can still do some work to make clear what I take to be the key features of obligation.

The first aspect of moral obligation to draw out is its normative force. If I have a moral obligation to carry out some action, then I must do it; I am required to. For instance, when he discusses the 'purported normative force' of moral obligation, Darwall (2006: 26) describes them as being 'supremely authoritative'. The second aspect of moral obligation to draw out is the closely related notion of what Scheffler calls 'the thesis of overridingness'. To wit: 'it can never be rational knowingly to do what morality forbids' (Scheffler, 1992: 52). The idea here is that the only thing that can trump, or override, a moral obligation, is some *other* moral obligation. There is, of course, much more that could be said about the nature of obligation. However,

to do so would take us much further afield. Let us now turn to the question of how, with this notion of obligation in hand, we can look to argue for an obligation based account of trust.

3.1 The Argument from Unification

A feature of the literature on 3-placed trust is that it requires us to differentiate two kinds of trust. On the one hand, we have trust in a rich, morally loaded sense, such that a breach of that trust will lead to moral disapprobation. On the other hand, we have trust in a sense that carries with it no such moral overtones. For instance, I might reasonably be said to ‘trust a ladder to take my weight’. If the ladder collapses under me, it would be quite wrong to (morally) blame the ladder for my subsequent injury.

Carving the conceptual terrain in this way has led to the latter sense of trust being described as ‘mere reliance’ and the former sense of trust being the focus of discussion. Nonetheless, it is acknowledged that we use the term ‘trust’ for both, and that we have two related concepts here.

This leaves us with a choice-point.

Option 1: there is some overarching notion, *trust*, of which morally-loaded-trust and mere-reliance-trust are but two kinds.

Option 2: just as we have two genuinely distinct concepts of what it is to be a bank (river and financial), so we have two genuinely distinct notions of trust.

Option 2 is problematic for two reasons. First, we might reasonably think that when providing a theory about the world, there is some theoretical pressure to minimise the number of concepts that we deploy (cf. Schaffer, 2015). If so, that would give us some reason to prefer an account whereby we have but one concept of trust that comes in two flavours, rather than two genuinely distinct notions of trust. Second, though related, we have no reason at all to think that there are two distinct kinds of trust, as opposed to one kind of trust that comes in two flavours. Whilst it’s true that some instances of trust appear to bring with them moral commitments, while others do not, that serves only to show that some instances of trust are morally loaded and that others are not. It does not, on its own, suffice to show that we have two different concepts of trust.⁷

⁷ By way of analogy, we might consider parsimony itself. It is quite true to say that parsimony comes in different kinds. There is quantitative parsimony (a pressure to reduce the number of things), qualitative parsimony (a pressure to reduce the number of kinds of thing), ideological parsimony (a pressure to reduce the number of concepts in our theory), and so on. But these are all instances of parsimony; a pressure to reduce the number of xs, for some value of x. In contrast, we might consider the ways in which metaphysicians and scientists use the term ‘substance’. In metaphysics it is a term of art, typically associated with whether or not some x is metaphysically fundamental; in physics, it most typically refers to a chemical kind. In that case, there is no overarching theory of substance of which the metaphysical and scientific uses are two instances; these are two genuinely distinct notions. And, to repeat, in the case of trust, I do not yet see that we have reason to think that the fact that there is moral trust and non-moral trust gives us any reason to think that there are two distinct kinds of trust (akin to the case of substance), as opposed to two different forms of trust that fall under the same theoretical banner (akin to the case of parsimony).

In sum: we have pressure to reduce the number of genuinely distinct ideological kinds in our best theory, and no reason here to think that there really are two distinct kinds here. We also have some pressure from the use of terms like ‘trust’ in both moral and non-moral contexts, to provide a unificatory account of trust. This leaves us with the first horn: there is some overarching notion, *trust*, of which morally-loaded-trust and mere-reliance are but two kinds.

It is hard to see how to recover that from Hawley’s account. By way of reminder, Hawley maintains that to trust someone to do something is to believe that she has a commitment to doing it, and to rely upon her to meet that commitment. And, as Hawley (2014: 5) has it, it’s fairly common in the literature on three-placed trust to treat it as reliance *plus* some extra factor. On reliance itself, little is said. Hawley (2014: 4) approvingly gives the following gloss on Holton’s (1994) account, whereby ‘to rely on someone to X is to act on the supposition that she will X: this can fall short of believing that she will X, though it is incompatible with outright belief that she will not X’.

If it were the case that reliance is non-moral trust, then it would also need to be the case that to non-morally trust someone to X would be to act on the supposition that she will X. However, as we can see from the following section, it’s clear that reliance is not the same as non-moral.

3.2 (Non-moral) Trust is not Reliance

The first very general point is that we can rely on objects that we do not trust (in any sense at all). This is easy to illustrate. I may rely upon my historically break-down-prone car to get me home. That is, I may act on the supposition that my car will get me home. I will get into it and attempt to drive it home. I will not pursue other travel options or seek to delay subsequent appointments. Nonetheless, in this situation I would be most unlikely to assert that I *trust* my car to get me home (even in any non-moral sense). This is a situation where I’m crossing my fingers and hoping for the best—I am relying on my car to give me just one last journey.

Similarly, I may rely on my rickety windows to keep out intruders, though I do not trust them to do so. Suppose, for instance, that though aware of the potential threat to my belongings and personal safety if I do not replace my windows, despite living in an area with high prevalence of crime I choose to spend my money on a holiday. Quite reasonably, I do have some misgivings about this choice. I am concerned about whether my windows will do their job. They *look* pretty fragile. They don’t look at all reassuring—though, to be sure, they do offer at least some minimal protection. In this situation, I certainly don’t *trust* my windows to keep me safe. Nonetheless, I continue to act on the supposition that they will protect me and my belongings. I rely on them.

Second, as in the No-Reliance cases described in 2.1, I can *also* non-morally trust an entity without actively relying upon it. Suppose, for instance, that I am considering whether to place a vase upon a shelf. The vase is extremely valuable, but not very large and not very heavy. The shelf is firmly fixed to the wall. Nonetheless, I decide to place the vase somewhere else. My partner asks me, somewhat teasingly, ‘what’s the matter, don’t you trust the shelf to hold the weight of the vase?’. Not rising to the bait, I respond that I do, but that I think the vase would look better elsewhere. There is

nothing conceptually challenging about this interlude. I (non-morally) trust the shelf to take the weight of the vase, but choose not to rely on it to do so. For these reasons, then, non-moral trust does not seem to be mere reliance. I (non-morally) trust the shelf to take the weight of the vase, but do not rely on it to do so.⁸ Thus, it's not merely that we have reliance and (moral) trust. Rather, we have reliance, non-moral trust, and moral trust.

3.3 Unification Through the Ought Account

We can do better if we pursue the Ought account. By way reminder, the Ought account states that:

x trusts y to ϕ if and only if x believes that y ought to ϕ , and, believes that, were y called upon to ϕ , then y would ϕ ;

The key point here is that 'oughts' come in many flavours. There are robust, morally loaded and normative 'oughts': "you ought not to eat babies"; "you ought to give to charity"; "you ought to be nice to your sister". There are also non-moral 'oughts'. To give but two examples: pragmatic 'oughts' ("you ought to leave now to catch the train") and predictive 'oughts' ("the ladder ought to take your weight").⁹

My unification is therefore simple: different *oughts*, different kinds of trust. The account of trust that I have given is perfectly general. What makes some instances of trust instances of moral trust is that they include a moral ought. What makes some instances of trust non-moral trust is that they include a non-moral ought. Hilary trusts Nasser to water their plants because Nasser has made a promise to water the plants; Nasser's making of a promise to water the plants means that (*ceteris paribus*) Nasser ought (morally) to water the plants. Susan trusts the ladder to take her weight because she thinks that a well-designed ladder that is robust and maintained as hers, ought (in the predictive sense) to take her weight.

Moreover, the Ought Account also deals with the concerns discussed in reference to unreliable objects. I do not think that my rickety windows *ought* (in the predictive sense) to keep out an intruder; I think they may well cave in at the first shove. I do not think that my dodgy car *ought* (in the predictive sense) to get me home; I think it may well fail. Simply, and in both cases, the predictive 'ought' that I introduced above, is *not* one that holds in these cases. As such, it is not the case that I trust these objects—in any sense. That being so, the Ought Account of trust delivers the desired unification. This is a mark in its favour.¹⁰

⁸ It's also hard to see how to recover non-moral trust on any other account of moral trust that treats moral trust as reliance plus some other factor, F. After all, if moral trust is reliance plus another factor, then once we remove the other factor, we no longer have trust—we have mere reliance.

⁹ I take 'ought' to be a weak necessity modal and thus as different from 'must', a strong necessity modal. See Silk ([forthcoming](#)) for illuminating discussion.

¹⁰ Even if we allowed that every commitment brings with it an obligation (something stronger than Hawley (2014: 18) herself claims) we wouldn't thereby account for non-moral trust on Hawley's analysis, for windows and cars make no commitments.

3.4 Unification and Obligation vs. Commitment (Redux)

I argued above that Hawley's Wild West case didn't tell in favour of a commitment account. I want to revisit the case, here. When we reflect on the Wild-West case, it's interesting to note that non-moral trust is relevant here.

For instance: let us suppose that I come across you, Hawley's armed opponent, in the desert. Setting to one side questions of what it's moral to do, I focus on matters of survival: *what ought I to do in order to survive?* I reason as follows. I know nothing about you. We're a long way from town. For you to be brave enough to be out this far, alone, it stands to reason that you're a good shot; confident enough in your abilities to hold your own. However, you've not *immediately* tried to shoot me. So, I conclude as follows: I think that you're a good shot who believes and that you believe you could probably successfully engage me in a fight, but you've not fired. This tells me that you might be open to negotiation before gunfire. In those circumstances, I think I can trust you not to shoot—in a sense. This is not trust in a rich, morally loaded sense of the word. This is kind of predictive trust, based on my best guesses as to our relative psychological states and our sharpshooting abilities.¹¹ The Ought Account makes space for this. This is a predictive trust not a moral one.¹²

4 The Argument from Motive

Consider the following case. Alice is a detective. She has been trailing Smith and Jones. Jones is the crime King-pin who is blackmailing Smith into doing terrible things. Alice has been painstakingly building her case against Jones, but has not made any contact (direct or indirect) with Smith. She learns, via her sources, that Jones is demanding that Smith rob a local convenience store and that Smith has agreed. This is the big break she has been waiting for. She believes that Smith has a commitment to robbing the store and, in setting up her sting operation, relies upon him to do so. However, in a twist, Smith decides that enough is enough and stands up to Jones, at the last-minute refusing to go through with the robbery. Alice's operation fails. She later learns of Smith's change of heart.

I think that it would not be appropriate to call this an instance of (broken) moral trust between Alice and Smith. For instance, this does not seem to be a case where Alice would feel (morally) betrayed or let down by Smith. He made her no promises. He has never spoken to her.

¹¹ Indeed, it's a form of trust—sometimes called 'process based'—frequently discussed in the non-philosophical literature. See Zucker (1986) for the distinction between (1) process-based trust, which is connected to past or expected exchange; (2) characteristic-based trust, which is connected to people and based on social characteristics; and (3) institutional-based trust, connected to formal societal structures, based on individual or firm-specific attributes or on intermediary mechanisms.

¹² This insight also enables us to deal with other seemingly difficult cases. For instance, I might predictively trust an assassin to murder a child once they have been contracted to do so, whilst judging that, morally, the assassin ought not to murder the child. Accounting for this kind of case is easy enough once we recognise that a predictive ought can be very different from a moral ought.

Nonetheless, it also seems that there would be nothing wrong with Alice believing that, *in a sense*, she was trusting Smith to rob the store. Certainly, in explaining herself and her failed operation to her superiors, there would be nothing conceptually problematic about Alice noting that she was trusting Smith to rob the store in order that she could catch him in the act, and trace this back to Jones. This is, once again, trust in a predictive sense.

On Hawley's account, we have no explanation of this. Alice either believes that Smith has a commitment to robbing the store, or she does not. If Alice does believe that Smith has a commitment to robbing the store and she relies on him to do so,¹³ then on Hawley's account Alice (morally) trusts him to rob the store. However, that would be the wrong verdict. She didn't trust him *in the moral sense* of the word. So, to the other option: Alice does not believe that Smith has a commitment to robbing the store. In that case, we don't get the incorrect verdict that 'Alice (morally) trusted Smith to rob the bank'. But we are still at a deficit here, for we have no explanation from Hawley's account as to why it's quite correct to say that 'Alice trusted Smith to rob the bank' in the predictive sense of the word 'trust'.

The Ought Account delivers the desired explanation. There is a sense of 'ought' relevant to the intelligence that Alice has received, such that Alice thinks that Smith ought to rob the store. It seems a predictive ought (and a predictive form of trust). To give colour: Alice may think that 'Smith ought to rob the store at about 1:00am tomorrow, so we ought to be in place by 09:30'. This seems a predictive 'ought' and leads to trust in the non-moral sense. Again, a non-moral 'ought' leads to non-moral trust.

5 The Argument from Ineffective Commitments

The final kind of case to be considered is one where sincerely held commitments are causally ineffective. A case we can use to illustrate is a high stakes athletics race. Ash has been in decent form for the season and makes the final at a prestigious competition, but does so in last place. Nothing in his form suggests he has a hope of winning—a decent season is not a great season and he's been consistently beaten by all other seven competitors in the final. In turn, this means that very good odds are available on Ash winning. He's not thought to have much of a chance, and so the bookies are encouraging people to take a punt. As it goes, Mark is nearly broke due to his gambling habit. Mark is aware of just how unlikely it is that Ash will win: he's had an eye on Ash's form throughout the season. He hears Ash interview just after the semi-final and hears Ash express his entirely sincere *commitment* to winning the final: rolling out the clichés, Ash claims that he is not there to take part, he's there to win. Mark takes what remains of his funds and gambles on Ash winning, figuring that the odds are good enough that this is the bet that could give him the financial respite that he so desperately needs. Mark thereby believes that Ash has a commitment to winning the race and, with his high-stakes bet that will wipe him out if he loses, and

¹³ Hawley defends the view according to which assertions generate commitments, so Smith's assertions that have been reported to her suffice for Alice to form the view that Smith has such a commitment.

make his fortune if he wins, is now relying on Ash to win. Nonetheless, it's clear in this scenario that Mark is unlikely to win. It would require a hugely implausible upset for the rank outsider to best their competitors.

However, nothing in this scenario suggests that Mark trusts Ash to win. The chances are so low, that, notwithstanding Ash's commitment, and Mark's reliance on him, the claim that Mark thereby needs to trust Ash appears far-fetched (so, unless we held a theoretical commitment to the Commitment Account, I can't see why we would think of this as trust).

Indeed, there is *no* sense in which we should say that Mark trusts Ash to win—moral or non-moral. Mark has never spoken to Ash. Aside from his interview, Mark has never heard Ash speak.

We can't very well say that this is an instance of trust. So, what can we say? One option for Hawley would be to say that this is just a case of Ash speaking imprecisely, or over enthusiastically—in an effort to psych themselves up, perhaps. We thus take what Ash says and re-interpret him as meaning something quite different.

We could do this, of course. But: (i) it's not clear that would be true to Ash's psychology; (ii) it's not clear why we would want to reinterpret Ash in this way. On the first point: elite athletes typically frequently and sincerely believe that they will win the events they enter. Athletes are also often funded with an expectation that they *will* win (or at least will deliver a medal). Thus, it's reasonable to think both that Ash believes that he will win and that his role (as a funded elite athlete) demands it of him. On the second point: we can grant that we should paraphrase Ash in this way if we wish to preserve the commitment account of trust. That's fine. But why *should* we do so? Ash has uttered a sentence. Why not take him at his word? Aside from a wish to preserve a theory of trust, what reason do we have to offer a reinterpretation like this? I'm aware of no reason, so I'm aware of no pre-theoretic reason not to take his utterance at face value. I don't want to deny that the proponent of the Commitment account *can* offer such a reinterpretation, but we have no reason to go along with it.

Note, also, that this is not a problem that is faced by the Ought Account. On the Ought Account, Mark does not believe that Ash *ought* to win—for any sense of 'ought'. In terms of prediction: as before, Mark acknowledges that Ash's victory is unlikely (it's precisely this that leads to him getting such good odds). He does not sincerely entertain the thought that: 'Ash's form is *so good* that he really ought to win this'. In terms of morality: there is no moral failing that Mark believes Ash would make if he failed to win the race—there is no moral 'ought' here. That being so, the 'Ought' account correctly delivers the result that Mark does not trust Ash to win the race (in any sense of the word 'ought').

6 Conclusion

I've presented a novel account of trust and argued it fares better than Hawley's commitment account. I've not tried to show that the Ought Account outperforms all other options in the philosophical market-place. Nonetheless, I think I've done enough to show that the view is a going concern.¹⁴

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Baier, A. (1986). "Trust and Antitrust. *Ethics*, 96, 231–260.
- Darwall, S. (2006). *The Second-Person Standpoint: Morality, Respect, and Accountability*. Cambridge: Harvard University Press.
- Faulkner, P. (2007). On Telling and Trusting. *Mind*, 116, 875–890.
- Hardin, R. (2002). *Trust and Trustworthiness*. New York: Russell Sage Foundation.
- Hawley, K. (2014). Trust, Distrust and Commitment. *Nous*, 48, 1–20.
- Holton, R. (1994). Deciding to Trust, Coming to Believe. *Australasian Journal of Philosophy*, 72, 63–76.
- Jones, K. (1996). Trust as an Affective Attitude. *Ethics*, 107, 4–25.
- Jones, K. (2004). "Trust and Terror," in *Moral Psychology*, edited by Peggy DesAutels and Margaret Urban Walker, Lanham MA: Rowman and Littlefield.
- Schaffer, J. (2015). What not to multiply by necessity. *Australasian Journal of Philosophy*, 93, 644–664.
- Scheffler, S. (1992). *Human Morality*. New York: Oxford University Press.
- Silk, A. forthcoming. 'Weak and strong necessity modals: On linguistic means of expressing "a primitive concept ought"', in *Meaning, decision, and norms: themes from the work of Allan Gibbard*. Dunaway, B. & Plunkett, D. (eds.). Michigan Publishing (see here for the pre-print: <https://research.birmingham.ac.uk/en/publications/weak-and-strong-necessity-modals-on-linguistic-means-of-expressin>).
- Tallant, J. (2017). Commitment in cases of Trust and Distrust. *Thought*, 6, 261–267.
- Zucker, L. G. (1986). 'Production of trust: Institutional sources of economic structure, 1840–1920'. *Research in Organizational Behavior*, 8, 53–111.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

¹⁴ Many thanks to the referees for comments on previous versions of this paper.