

Forthcoming in *Philosophical Psychology*

Therapeutic Trust

J Adam Carter

adam.carter@glasgow.ac.uk

This paper develops and defends a new account of *therapeutic trust*, its nature and its constitutive norms. Central to the view advanced is a distinction between two kinds of therapeutic trust—*default therapeutic trust* and *overriding therapeutic trust*—each which derives from a distinct kind of trusting competence. The new view is shown to have advantages over extant accounts of therapeutic trust, and its relation to standard (non-therapeutic) trust, as defended by Hieronymi (2008), Frost-Arnold (2014), and Jones (2004).

1.

Suppose you are leaving town for the weekend and need someone to watch your house, feed your pets and water your plants. Now imagine two choices you might make. You might trust a reliable friend who has an established track record of responsibility and loyalty. But, you might instead trust your 16-year-old nephew with no such track record to speak of. In the latter kind of case, suppose trust is undertaken with the intended aim of bringing about (or increasing) trustworthiness.¹ Philosophers of trust often use the term ‘therapeutic trust’ to refer to this latter species of trust, in order to distinguish it from more standard cases of (non-therapeutic) interpersonal trust.

The matter of how exactly to characterise the relationship between non-therapeutic and therapeutic trust is contested.² Here is the problem in a nutshell. Philosophical accounts of the nature of trust attempt to say what trusting someone with something essentially involves³, typically by focusing on how exactly to characterise the kind of trusting attitude one has towards her trustee. Once such accounts are made precise, it looks like therapeutic trust—given how the attitude we have in such cases about the trustee’s reliability is usually much *less* optimistic than in non-therapeutic cases—either (i) simply doesn’t get ‘ruled in’ as genuine trust on the account, or (ii) the account gets modified—perhaps stretched quite thin—in order to fit therapeutic trust in.

¹ See, for example, Horsburgh (1960) and Jones (2004)

² See, e.g., Horsburgh (1960) and McGeer (2008). For an overview, see McLeod (2015, 1).

³ The kind of trust that is principally at issue in debates about therapeutic trust is *three-place* trust, e.g., with an infinitival component (schematically: A trusts B to X). As Baier (1986, 236) puts it, philosophers of trust—and not just those interested in therapeutic trust—are concerned centrally with ‘one person trusting another with some valued thing’; likewise, as Hawley (2014, 2) puts it: trust is ‘primarily a three-place relation, involving two people and a task’. On one popular way of thinking about the relationship between three-place trust and two-place trust, latter can be explained in terms of the former, which is the comparatively more fundamental notion. Explicit defences of ‘three-place fundamentalism’ can be found in, e.g., Holton (1994) and Jones (1996). For some recent resistance to this idea, see Domenicucci and Holton (2017).

Here is the plan. In §§2-4, I discuss three notable ways philosophers of trust have attempted to deal with the tricky issue of therapeutic trust and its relationship with ordinary non-therapeutic trust: (2) Pamela Hieronymi's pure/impure approach⁴; (3) Karen Frost-Arnold's 'unity' approach⁵; and (4) Karen Jones's 'normative difference' approach⁶. Each is shown to be problematic. 5-9 then develops a new way of thinking about therapeutic trust which avoids the problems facing the other three views while at the same time offering its own additional advantages. 10 concludes by canvassing some potential objections and replies.

2.

According to Pamela Hieronymi (2008), therapeutic trust is not 'pure' or 'full-fledged' trust. Trust is full-fledged (alternatively: pure) only if one actually *believes* that the person in question will do the thing in question.⁷ One can risk betrayal by entrusting something to someone without believing they'll⁸ actually do what they've been entrusted to do. But this is not 'pure' trust.

In support of this way of thinking about therapeutic trust, Hieronymi offers the following case-pair involving the betrayal of a secret.

SECRETS: Consider two cases. In one, I fully believe you are trustworthy; in the other, I have doubts about your trustworthiness, but, for other reasons (perhaps to build trust in our relationship, perhaps because I think friends should trust one another, or perhaps simply because I have no better alternative), I decide to tell you my secret. Suppose that, in both cases, you spill the beans, and that you do so in the same circumstances, for the same reasons.(2008, 230)

According to Hieronymi, once we thus hold fixed both (i) the 'importance of the good entrusted' (2008, 230); and (ii) 'the wrongness of the violation' (2008, 230), then:

[...] it seems plausible that one's degree of vulnerability to betrayal tracks one's degree of trusting belief ... further, this seems to be because, in the second case, there was less trust to betray (2008, 230–1).

There are, however, two problems with this diagnosis of SECRETS. The first is that it's not at all clear that one's degree of vulnerability to betrayal really tracks one's degree of trusting belief, *even when* the importance of the good entrusted and the wrongness of the violation are held fixed. To see why, just suppose we run a variation on Hieronymi's SECRETS case-pair where, in the first case, my belief that you are trustworthy is full (stipulate: credence 1) but, at the same time, completely *irrational*. By Hieronymi's reasoning, the betrayal is more greater by degree in the first case simply because of the irrationally ratcheted up belief. But it's not. I might, due to having this strong albeit

⁴ (Hieronymi 2008).

⁵ (Frost-Arnold 2014).

⁶ (Jones 2004).

⁷ For further support of the idea that non-therapeutic trust requires belief that the trustee will prove trustworthy, see Adler (1994), Keren (2014) and McMyler (2011). For criticism, see Jones (1996), McLeod (2002), McGeer (2008), Faulkner (2007); Faulkner (2011); Baker (1987).

⁸ The singular pronouns 'they' and 'them' are used throughout whenever gender is unknown or irrelevant.

irrational credence that you are trustworthy, be even more inclined than otherwise to *think* that the betrayal is serious. But it wouldn't in fact be a worse betrayal simply *on account of* the ratcheted up irrational credence.

But suppose, for the sake of argument, we grant that one's degree of vulnerability to betrayal tracks one's degree of trusting belief. It is worth noting that *even if* this were true, it could be explained without recourse to the idea that there is—in the therapeutic case—“less trust” to betray. For example, it might be that vulnerability to betrayal is one among various features of trust, and it's a feature that lines up with (e.g., by closely tracking) trust's doxastic component, whereas other features of trust (for example: whatever features makes it resilient to certain kinds of monitoring⁹) might track some *non-doxastic* (e.g., affective) component of trust.

If something like this were right, then we couldn't move simply from the idea that one's degree of vulnerability to betrayal tracks one's degree of belief to the conclusion that there is less trust in cases with less belief. After all, such cases might feature more prominently some other (e.g., affective) aspect of trust, and in virtue of the presence—perhaps, surfeit—of which there is not, on the whole, “less trust.”

But Hieronymi has a second argument for relegating therapeutic trust the ‘impure’ category. The second argument has to do not with vulnerability but with the legitimacy of certain kinds of complaints. This second line of reasoning goes as follows. People can *legitimately complain* about not being trusted fully when they are trusted in the absence of belief, which occurs only when other people lack confidence in them but trust them nonetheless (2008, 230). For example, imagine the 16-year-old from our opening case saying: “But you don't *really* trust me”, upon finding out that the rationale for the trust was largely trust-building, in the absence of a belief that they'd prove trustworthy. The felicitousness of such a complaint is, for Hieronymi, meant to support the idea that therapeutic trust is not pure or full-fledged trust.

This reasoning is also problematic, though, in so far as it's supposed to motivate relegating therapeutic to a ‘second tier’. Just as the teenager could complain in this scenario, they could also felicitously praise or thank you for trusting them *despite* lacking confidence. “Wow, you trusted me without believing—you must have *really* trusted me!”¹⁰ This is not to say that Hieronymi's example complaint is *not* felicitous, nor that praise or gratitude for trusting despite lacking confidence is any more felicitous than is complaining that one has trusted in the absence of belief. Rather, the point is that it is not clear that complaining about trust in the absence of belief is in any way *more* felicitous than praising or thanking a trustor who trusts one in the absence of it.¹¹

In sum, Hieronymi's arguments from vulnerability and complaint legitimacy don't give us good reason to think that the difference between non-therapeutic trust and therapeutic trust is a difference in ‘purity’ of trust.

⁹ For discussion on this point, see, e.g., Baier (1986) and Wanderer and Townsend (2013).

¹⁰ For a defence of the idea that trust's most pure form involves the absence of belief, see Mollering (2006).

¹¹ Moreover, the felicitousness of such a reply gains some support in the literature on the psychology of gratitude (see, e.g., Emmons and McCullough 2004); for instance, gratitude is a predictable response by one to another who has placed a kind of ‘faith’ in their good will or competence.

3.

Let's look now at an attempt—due to Karen Frost-Arnold (2014)—to 'broaden' an account of trust so that it is wide enough to rule both in both varieties of trust. On Frost-Arnold's proposal, A trusts B to ϕ iff the proposition that B will ϕ is part of A 's 'adjusted cognitive background' (2014, 1963–4), where one's adjusted cognitive background includes all and only those propositions that one *accepts* for the purposes of practical reasoning—where acceptance does not entail positive belief¹² (e.g., positive belief of the sort that is generally lacked in therapeutic cases, even if often present in non-therapeutic cases). This kind of 'unity' view does not relegate therapeutic trust to a second-tier, as Hieronymi's proposal does, but rather 'brings it in to the first tier'—e.g., by subsuming it within a wider account of trust *simpliciter*.

There are two main problems with Frost-Arnold's 'unity'-style approach. The first is that the acceptance requirement needn't be satisfied in all cases of therapeutic trust. Suppose you trust your teenager to drive your car for the weekend and return it safely. Suppose further that, upon doing this, you purchase additional insurance, just in case. By purchasing this additional insurance, you are not accepting the proposition that the teenager will return the car safely—viz., to do what you'd trusted them to do—in the course of your practical reasoning. You act instead on the proposition that they might realistically enough not do so.¹³

But the existence of this mitigating back-up plan doesn't preclude the case from having been a case of therapeutic trust in the first place. That is, you don't suddenly cease to be therapeutically trusting the teenager with whom you aspire to build trust once you buy the insurance. It's not as though the vulnerability to betrayal one subjects oneself to is eliminated by one taking steps to mitigate damages against the risk occurring.¹⁴

Some philosophers of trust have pressed back on this point. For example, Keren (2019) holds that trusting involves declining to take precautions against the trustee's failing to come through.¹⁵ This

¹² The idea that accepting a proposition, understood as taking it for granted in one's practical deliberations—alternatively: acting 'as if' the proposition is true—does not entail that one believes the proposition to be true has been defended variously in epistemology, the philosophy of science and elsewhere. For some representative discussions of how belief and acceptance come apart, see Cohen (1989), Bratman (1992), and Buckareff (2010).

¹³ This is the case, to be clear, even though we needn't suppose that you positively believe that the teenager *won't* return the car as entrusted.

¹⁴ This point, it is worth noting, is compatible with the widely accepted idea that *monitoring* is incompatible with trusting, either of a non-therapeutic or therapeutic variety. See, e.g., Baier (1986, 260).

¹⁵ Keren (2019, 121) actually formulates his position as follows, with the qualifier 'every': 'If you rely on a person to ϕ but take *every* precaution against the possibility that she might not ϕ —by seeking evidence that might indicate that she might fail to ϕ and by acting in order to minimize the harm caused in case she fails to ϕ —then you do not trust her to ϕ . You might rely on her to ϕ , but you do not trust her to do so' (my italics). As formulated, this is not controversial, as this is tantamount to the statement that trusting is incompatible with certain kinds of monitoring. What is at issue in the example I am discussing above, involving an insurance policy, is rather whether mitigating at all against the risks of the damage that would be incurred by the trustee's betrayal would be compatible with nonetheless therapeutically trusting that person. My contention that it is is thus compatible with granting Keren's point that some kinds of monitoring—i.e., such as those that involve taking every precaution against the possibility the trustee won't come through—are incompatible with trust (therapeutic or otherwise). That said, Keren also makes claims about taking precautions, in the specific case of epistemic trust, which appear to go beyond the statement of his view noted above, and which appear to imply that it is essential to epistemically

idea seems *prima facie* plausible in the epistemic case, specifically, where what the truster trusts the trustee to do is to tell them the truth. What ‘taking precautions against the trustee’s failing to come through’ would amount to in this case would be finding additional evidence that bears on the truth of the proposition. But *then*, having sought sought such evidence, it doesn’t look as though you are trusting the person’s word at all.

To the extent that trusting (therapeutic or otherwise) *does* involve declining to take precautions against the trustee’s coming through, this might very well be idiosyncratic to the epistemic case where there is a constitutive tension between relying on one’s word and acquiring the kind of evidence one would acquire by taking precautions. Crucially, we find no such tension though outside the epistemic case, at least when we hold fixed that the precautions are (as in the example of the insurance policy) precautions that are solely designed to mitigate damages *if* the trustee does not come through. Compare: this is importantly different from, and does not imply, taking precautions designed to lower the likelihood that the trustee will *fail* to come through — as one might do by hiring a team to accompany the teenage driver. Accordingly, the attempt to reply to the objection raised to Frost-Arnold by way of appealing to Keren’s insights about epistemic trust looks to come up short.

The upshot is that Frost-Arnold’s unified account of trust which frames trust in terms of acceptance is still too narrow to do what she wants it to do, which is to rule in all cases of trust, non-therapeutic and therapeutic alike.

Even more, the proposal faces a second problem. The second problem concerns the evaluative normativity of trusting. Generally speaking, evaluative norms – unlike prescriptive norms, which prescribe conduct – regulate what it takes for a token of a particular type of thing to be good or bad with regard to its type, where the ‘goodness’ or ‘badness’ here is *attributive* in Geach’s (1956) sense – viz., the sense in which a sharp knife is a good knife, *qua* knife, regardless of whether it is good or bad *simpliciter*. (Likewise, in this sense, a known belief is a good belief, regardless of whether it would be good or bad *simpliciter* – viz., as it would be were the content of the knowledge instructions for igniting a terrible bomb.¹⁶) The worry is that the view lacks the resources to account for why reasonable therapeutic trust isn’t just *bad* as an instance of trusting.

Continuing with the teenager car case: let’s suppose you have no trust-building objectives in mind, and simply want someone dependable to drive your car for the weekend and bring it back safely. Foolishly, you choose the teenager with a record you know is patchy at best. This looks like *bad* trust, even if it would *not* be so with therapeutic purposes in play.¹⁷ But it’s hard to see how we’d explain this normative difference on Frost-Arnold’s unity-style account. One might try to begin to

trusting someone that you decline entirely from taking precautions against their not coming through. Because this thesis, at least if applied generally and not just in the epistemic case, is in tension with my assessment of the insurance policy case, I focus in the main text on it as opposed to on Keren’s less contentious formulation quoted above.

¹⁶ For a helpful overview of the prescriptive/evaluative norm distinction, with reference to attributive as opposed to predicative goodness, see McHugh (2012, 22) and, as this distinction applies to belief specifically, Simion, Kelp, and Ghijzen (2016, 384–6).

¹⁷ One sense in which the trust here is bad is that it is not likely to be *successful*, in that trusting a teenager involves incurring a relatively higher risk of betrayal than normal. For discussion on this point, see Carter (2020; forthcoming).

tell such a story by appealing to the ‘epistemic constraint’ that Frost-Arnold places on the kind of acceptance that matters for a proposition’s being ruled-in the adjusted cognitive background. But the epistemic constraint she places on acceptance is really a very minimal one. It precludes just one thing: positive belief that the person will *not* do the thing in question. This kind of constraint won’t help us in any way to adjudicate the normative question we’re interested in—viz., how to distinguish at least some cases of good therapeutic trust from plain old bad trust.

There’s another thread to this point. Just as picking out an unreliable person is ‘bad’ trusting (likely to lead to one’s trust being betrayed¹⁸) even if trusting that unreliable person could have been reasonable were therapeutic purposes suitably in play, it doesn’t follow that *simply stipulating* a therapeutic purpose suffices to make any instance of therapeutic trust *good* therapeutic trust. One can surely be better or worse at therapeutic trusting, just as one could be better or worse at trusting more generally—and indeed, very plausibly in light of different kinds of skill sets. None of this looks explicable (at least, in any straightforward way) if we embrace a unity-style view like Frost-Arnold’s.

4.

The foregoing discussion suggests that what’s needed is an account of therapeutic trust which explains clearly how it features some kind of *normative difference* with respect to ordinary, non-therapeutic trust. This is exactly what Karen Jones’s (2004) account of therapeutic trust offers. Unfortunately, she identifies the *wrong kind* of normative difference.

According to Jones (2004), therapeutic trust involves the normative attitude that the trustee *ought* to do what one trusts them to do, rather than optimism that they will do it. With reference to our opening case pair: when you trust the reliable friend to watch your house for the weekend, you are optimistic that they will do this as you’ve entrusted them. While you’re not optimistic that the teenager will do the same when you trust them with the task, you nonetheless think in trusting them that they *ought* to do what you’ve entrusted them to do.

There are three main problems with this proposal. First, the normative attitude that the trustee ought to do what one trusts them to do is not necessary for therapeutic trust. Consider a case where a CEO, with the aim of striking up a romantic relationship with a low-level employee, entrusts that employee with an inappropriately enormous responsibility—hoping that doing so will help generate a trusting relationship between them as a precursor to such a romance. If the CEO is not blind to their exploitative reasons underlying the trust they are placing in this inexperienced employee, then they will not have the view that the trustee *ought* to actually do what they are entrusted to do. Quite the contrary, the CEO might well know that that the employee’s succeeding in doing what they’ve been entrusted with is beyond reasonable expectations. But this is therapeutic trust nonetheless, in

¹⁸ See Carter (2020; forthcoming) for a defence of this way of thinking about bad trust.

that it is—albeit for morally dubious reasons¹⁹—aiming to bring about and strengthen a trust relationship.²⁰

A second objection to Jones’s proposal—in so far as it purports to distinguish therapeutic from non-therapeutic trust—is that some cases of non-therapeutic trust involve not only optimism that the trustee will do what they are entrusted to do, but *also* the normative attitude that the trustee ought to do what they are entrusted to do. Suppose someone is drowning. I am nearby with a life vest—my expensive life vest—but my arm hurts, and so I can’t throw it off the boat to help. I trust my able-bodied and reliable friend to throw it. In this case, where I trust my friend to throw my vest, the trust isn’t therapeutic in any interesting sense.²¹ And yet, I have the (strong) normative attitude that the trustee *ought* all-things-considered to do what I’ve entrusted them to do. Thus, believing that the trustee ought to do what one entrusts them to do isn’t distinctive of therapeutic trust but not ordinary non-therapeutic trust.

A third objection to Jones’s proposal is that some cases of therapeutic trust positively *do* involve optimism that the trustee will do what they are entrusted to do, even if this optimism persists along with some serious doubts. To see why, it will be instructive to first consider how optimism comes apart from belief in both directions. In the literature on the psychology of optimism (e.g., Carver, Scheier, and Segerstrom 2010), an optimistic attitude, with respect to some situation X, is often characterised in terms of a kind of attention profile directed at favourable features of that situation. For example, if my car breaks down and I’m stranded on a highway, then an optimistic attitude might lead me to focus my attention on how doing certain things under my control (e.g., walking to the nearest petrol station) could better my situation.

Coming back to therapeutic trust: one can distribute one’s attention patterns in ways that line up with optimism (with respect to a trustee proving trustworthy) *without* having any positive belief that the trustee will prove trustworthy. (Compare: I can be optimistic when stranded without actually having the *belief* that I will be saved). For example, being optimistic that the teenager will look after the house properly or return the car might involve focusing on the teenager’s good traits, feeling pride in remembering past times they’ve exceeded expectations, etc. This is all compatible with a

¹⁹ No assumption is being made, to be clear, that the norms of trust are moral norms. I have suggested (in responding to Frost-Arnold) that we should expect an account of therapeutic trust to be reconcilable with plausible claims about the evaluative normativity of trusting; but this commitment is a very general one—viz., to there being norms (however we best articulate them) that regulate what it takes for a token of trusting to be good or bad with regard to its type. This is at most a commitment to attributive (rather than predicative) goodness of trusting in certain cases.

²⁰ The same kind of point can be made with reference to a more paradigmatic kind of ‘teenager trust’ case. Suppose a mafioso with a conscience but a weak will accepts a hit job and, rather than to do it himself, entrusts his unreliable teenager to carry out the hit—hoping that doing so will build trust. Assume the target of the hit is known by the mafioso to be completely innocent. It is entirely plausible here that the mafioso, bering this in mind, will appreciate that what he’s entrusted the teenager to do is not something the teenager ought to do. Yet, this fact (as in the CEO case) does little to change the fact that the trust here is of a therapeutic variety. Granted, this—as well as the CEO case—relies on a weak assumption in moral psychology, which is that one can desire to bring about some state of affairs while acknowledging that its being brought about would violate one (or more) prescriptive norm. For discussion, see Stocker (1979).

²¹ As we’ll see in 7, though, there is plausibly an uninteresting kind of ‘default’ therapeutic trust in play here that is implicated by most cases of non-therapeutic trust.

lack of belief that they will in fact bring the car back. Note, furthermore, that belief and optimism come apart in the other direction as well. You could believe someone will bring a car back without being optimistic simply because your attentional profile does not line up with what you believe. You might be irrationally paranoid, given a pessimistic perspective that does not line up with your belief that the trustee will prove trustworthy. These considerations in favour of the idea that optimism can float freely of one's doxastic attitudes supports that (*contra* Jones) the kind of doubts one has in the case of therapeutic trust aren't doubts that, as such, would preclude optimism that the trustee will prove trustworthy.

In sum, Jones is mistaken that therapeutic trust involves the normative attitude that the trustee ought to do what one trusts them to do, rather than optimism that they will do it. This is because it neither requires the normative attitude that the trustee ought to do what one trusts them to do—as per the CEO case—*nor* does it preclude optimism that they will do what they are entrusted to do, at least in so far as optimistic attitudes are plausibly demarcated by their attentional profiles.

5.

So far, we've seen that prominent extant accounts of therapeutic trust run into various kinds of problems. A presupposition common to each of the three views considered is that therapeutic trust is a univocal kind, and this is a presupposition we'd be better off rejecting.

There are two importantly different species of therapeutic trust—*default therapeutic trust* and *overriding therapeutic trust*. Each species of therapeutic trust interacts with ordinary (non-therapeutic) trust differently. And, each is *normatively constrained* differently from each other. Appreciating how this is so, we can—in addition to avoiding the kinds of problems considered—make sense of something other views can't, which is what makes therapeutic trust (of a philosophically interesting sort) *good* when it is.

6.

Any kind of performance with an aim internal to that performance-type can be evaluated along three dimensions: it can be evaluated for *accuracy* (i.e., did the performance succeed in attaining its aim), *adroitness* (i.e., was the performance skilful), and *aptness* (i.e., was the success *because* of the skill?).²² Take, for example, archery. An archer's shot is accurate if and only if it hits the target.²³ Regardless of whether it hits the target, it might be adroit—viz., it might be fired in a manner that would usually (enough²⁴) result in a successful shot, in normal conditions, and regardless of whether it in fact did. If the shot is both accurate and adroit—and further, if the accuracy is *because of* the adroitness—then the shot is *apt*.²⁵

²² This three-fold model for evaluating performances is due to Ernest Sosa (e.g., Sosa 2007, 2010a, 2015). For other applications of this 'AAA' model of performance evaluation, see for example, Turri (2016), Carter (2021), Kallestrup (2016), Kelp et al. (2017), and (ed.) Fernández Vargas (2016). For criticism, see Chrisman (2012).

²³ For an initial use of this analogy in order to illustrate performance normativity, see Sosa (2007).

²⁴ The threshold for reliability that is required for adroitness differs across performance domains. For example, to be adroit at shooting free-throws, one might require at least 70% reliability. The very best baseball-hitting averages, however, are barely over 30%. For discussion, see Sosa (2017, 75) and Greco (2010a, 77–78).

²⁵ Note that the conjunction of a shot's being accurate and adroit is compatible with it's failing to be apt. In the archery case, this might occur if a shot is fired adroitly while a chance gust of wind (i) first blows the arrow off

A key insight of recent virtue epistemology is that belief is a kind of performance with an aim, the aim of truth.²⁶ With reference to the ‘AAA’ model, we can ask three distinct questions for any given belief: is it true (accurate), is it formed in a manner that would reliably enough attain the truth aim (e.g., is it adroit), and thirdly, is the belief *apt*—viz., is it accurate (i.e., true) *because* of its adroitness.

Conveniently, non-therapeutic trust—as a kind of performance—can be subsumed within this kind of ‘AAA’ template just as belief can.²⁷ By placing it within the template, it will become clearer how we should think about therapeutic trust, or so I’ll argue. In order to get things off the ground, though, we need clarify the sense in which there is an aim internal to trusting in the standard non-therapeutic case. Such an aim, I submit, falls out of a more or less generic characterisation of (non-therapeutic) trust as an attitude.

The idea is as follows: just as an archer’s shot attains its aim if and only if it hits the target, and a belief succeeds in attaining its aim if and only if true, trust succeeds in attaining its aim if and only if *the trustee takes care of things as entrusted*.²⁸ If the trustee does not take care of things as entrusted, then that trust is not successful *even if* the trust is adroit—viz., even if one trusts in ways (e.g., by seeking out a reliable trustee, assessing risks of betrayal, etc.) that *ordinarily* would lead to successful trust, but was just unlucky on this occasion.

The ‘AAA’ model of performance assessment—extended to ordinary (non-therapeutic) trust—takes on the following template shape:

| | Accuracy/success | Adroit/skilful | Apt |
|---------------------|---|---|--|
| <i>archery shot</i> | hit target | shot issued from archery competence | shot successful (hit target) because of (archery) competence |
| <i>belief</i> | true | belief issued from epistemic competence | belief successful (true) because of epistemic competence |
| <i>trust</i> | trustee takes care of things as entrusted | trust issued from trust competence | trust successful because of trust competence |

target; and then (ii) a second gust of wind blows the arrow back on target, guiding it to the bullseye. In such a case, the accuracy + adroitness conjunction falls short of aptness because the accuracy is not *because* of the adroitness, but because of the second gust of wind. For discussion of cases with this kind of structure, see, along with Sosa (2007), also Greco (2010b) and Pritchard (2012).

²⁶ The canonical presentation of this idea is due to Sosa (2007). See also the essays in Fernández Vargas (2016). For a revisionist ‘knowledge-first’ approach to this model, which replaces truth with knowledge as the aim of belief, see, e.g., Kelp (2017) and Miracchi (2015).

²⁷ For a development of this idea, see Carter (2020; forthcoming).

²⁸ This phrasing is intentionally compatible with different kinds of substantive glosses. For example, according to Baier (1986, 234), taking care of things ‘as entrusted’ will include at least that the trustee do so out of goodwill toward the trustor. For criticism of this goodwill caveat, see Holton (1994, 65). Alternatively, for a very different way to gloss the idea of taking care of things as entrusted, consider Katherine Hawley (2014, 10) view that the trustee must be *believed to have* a commitment such that she is relied on to meet that commitment. For a more recent development of this idea, see (2019).

Throughout this analogy, adroitness is understood in terms of *competence*, which is an important concept in performance-theoretic evaluations. Before zeroing in on therapeutic trust specifically, it is worth clarifying a point about the structure of competences generally.

An archery competence is a disposition to reliably enough hit the archery-relevant aim (i.e., the target) when one tries. But there is an important qualification here, which is that—just as dispositions like flammability are indexed to manifestation conditions²⁹—competences are indexed to performance conditions (i.e., shape and situation) which are the pertinent ones to that performance type,³⁰ viz., the ones in which reliability is valued. It doesn't count against your archery competence, for instance, if you would reliably *miss* the target when releasing the arrow while drugged or placed in unusually high winds. Likewise, it doesn't count against a competence to trust successfully (reliably enough) if, too often, your trust would be betrayed when in improper shape (e.g., mentally incapacitated prior to selecting whom to trust with something) or improperly situated (e.g., unbeknownst to you, placed in an abnormal social-epistemic environment, viz., a city of liars, or a deception-driven flash mob).³¹

7.

The previous section has shown how we can situate paradigmatic non-therapeutic trust within the performance-theoretic framework that is familiar in other areas of philosophy. With non-therapeutic trust repositioned in this way, we now have a new vantage point to theorise about *therapeutic* trust—and in particular, about what I'll call *overriding* therapeutic trust—which is the most philosophically interesting variety of therapeutic trust.

But first it is worth making the following explicit: there is a kind of *default* therapeutic trust—viz., therapeutic insofar as it (trivially) aims at trust-building—that is *implicit* in paradigmatic cases where one trusts with the aim that the trustee take care of things as entrusted. We take for granted in trusting that trust will—apart from whatever else it does—play its normal social functions, functions that plausibly include the social function of strengthening trust relations.³² This is so even when we trust—as we do when we seek out someone reliable and trustworthy—with the basic aim that the trustee take care of things as entrusted.

In this respect, there is a minimal and trivial kind of therapeutic trust that is going to be implicit in garden variety (non-therapeutic) trusting, and this is so even in the absence of any explicit *intention*—

²⁹ For instance, flammability is a disposition a match has if and only if the match would struck if lit *under certain normal conditions*, which include, e.g., being suitably dry, being struck in the presence of ambient oxygen, etc. For discussion, see Sosa (2010a, 466). Note that the indexing of dispositions (and by extension, competences) to manifestation conditions fits snugly conditional approaches to analysing dispositions, but doesn't rely on any such commitment. For classical approaches to indexing dispositions to manifestation conditions, see, e.g., Ryle (1949), Goodman (1954), Quine (1960). For an overview, see McKittrick (2018) and Choi and Fara (2018).

³⁰ See Sosa (2010a, 466–7) and (2017, 195).

³¹ For a detailed discussion of the structure of competences, see, along with Sosa (2010a), and also Sosa (2017, 191–2).

³² For some representative defences of the role of trusting in trust-building, see Faulkner (2011 Ch. 1), Alfano (2016), Hall (2005), and Solomon and Flores (2003).

the kind of intention that is explicit in teenager-style cases—to satisfy the aim of building trustworthiness.³³

Moreover, the ‘implicit’ kind of therapeutic trust that accompanies ordinary trust as a default does not have its own constitutive aim. This is because default therapeutic trust is just *implicated by* normal, non-therapeutic trusting, which itself constitutively aims at the trustee’s taking care of things as entrusted. In this respect, the competences that are relevant—trivially—to default therapeutic trust are just those that matter for non-therapeutic trust that implicates it.

8.

The most interesting kind of therapeutic trust is not default therapeutic trust, but *overriding therapeutic trust*. This occurs when, as in our paradigmatic teenager cases, the aim of successful trust—given perceived vulnerabilities—isn’t itself enough to *motivate* one to risk trusting. Simply wanting your house to be watched over responsibly wouldn’t, from the perspective of a trustor, favour entrusting such a task to the teenager, as opposed to someone regarded to be more reliable; rather, the opposite would be the case. Necessary for bringing about overriding therapeutic trust is thus an ‘overriding’ and intentional aim—the aim of building or strengthening trust—that is distinct from the constitutive aim of ordinary non-therapeutic trust.

Unlike default therapeutic trust that is implicit in most normal trusting, *overriding* therapeutic trust is a distinct kind of performance from normal (non-therapeutic) trust, with its own constituent normativity. The constitutive aim of overriding therapeutic trust is not *merely* to trust successfully (viz., that the trustee take care of things as entrusted). But nor, it should be emphasised, is it *merely* to build trust. It is to *build trust through trusting successfully*.

Consider that just as ordinary (non-therapeutic) trust is defective when it misses its internal aim (that the trustee take care of things as entrusted), your choosing to trust your teenager to watch over the house has missed *its* mark if either (i) trust is not built (e.g., if a result of this trusting is not a strengthened trust relationship) *or* if (ii) trust is not successful (e.g., if the teenager throws a party, during which items from the house are stolen). *Even more*, though, your trust will have missed its mark even if the trust serves to build trust *and* the trust is successful, but (iii) if the trust built is not built *through* the successful trust, but for some reason disconnected with the therapeutic trust placed in them. This might be the case, for example, if the teenager watches over the house successfully, though—unaffected entirely by the trust you’ve placed in them—comes to trust you more nonetheless due to having, while watching over the house, spent some time reading false accounts of sacrifices you’ve made for them in the past, and only on this basis, develops toward you a stronger bond of trust.

Question: if overriding therapeutic trust constitutively aims not at mere successful trust, nor at the mere building of trust, but at building trust *through* successful trust, then what do (i) adroit and (ii)

³³ Consider, by way of analogy, one of the plausible social-epistemic functions of *assertion*, which is to generate knowledge in the hearer (e.g., Kelp 2018; Simion 2019; cf., Williamson 2000). On the assumption that assertion has such an aim, constitutively, it’s easy to see how asserters *implicitly*, in asserting, aim at other things (even if not intentionally), namely, whatever generating knowledge in a hearer generally involves, including playing roles that knowledge normally plays for the hearer. For example, one role that knowledge plausibly plays for a hearer who acquires it is that of being a possible premise in the hearer’s practical reasoning (e.g., Hawthorne and Stanley 2008).

apt overriding therapeutic trust consist in? Put another way: when we add overriding therapeutic trust and its constitutive aim to our ‘AAA’ template, how should we fill the rest in?

| | Accuracy/success | Adroit/skilful | Apt |
|-----------------------------|---|---|--|
| <i>archery shot</i> | hit target | shot issued from archery competence | shot successful (hit target) because of (archery) competence |
| <i>belief</i> | true | belief issued from epistemic competence | belief successful (true) because of epistemic competence |
| <i>trust</i> | trustee takes care of things as entrusted | trust issued from trust competence | trust successful because of trust competence |
| <i>trust_{O.T.}</i> | trust built through successful trust | ? | ? |

Adroit overriding therapeutic trust, on the ‘AAA’ model, will derive from a competence to attain the constitutive aim of *overriding therapeutic trust* reliably, which is the aim of building trust *through* successful trust. Given that competences are indexed to performance conditions, a clear view of the kind of competence that matters for overriding therapeutic trust requires an understanding of the conditions under which reliable performance matters for this *particular* kind of trusting. These conditions include (at least) the satisfaction of what I’ll call an *openness condition* and a *reciprocity condition*.

To appreciate the former condition, consider the following case:

DIANE: You need someone to babysit on short notice. There are a number of people you could ask, however, you choose a local teenager, Diane, whose parents you know. You have heard that Diane is troubled, and you have had a standing desire to take Diane under your wing in hopes of having a positive influence on her. A first step toward having such a positive influence, you think, will be to establish a bond of trust, a bond you hope to develop by entrusting her with the babysitting task despite her reputation. Unfortunately, and unbeknownst to you, Diane recently experienced a highly traumatic event, to which she has responded by closing off the possibility of developing a trusting relationship with *anyone*, at least until she has worked through this trauma. She succeeds in the task of babysitting, though at no point was she in a position where her being entrusted with this would have changed her distrusting stance of others.

In DIANE, the conditions for successful overriding therapeutic trust are simply not in place *ex ante*—and this is so even though the conditions in DIANE do not preclude her in any way from taking care of things as entrusted. With respect to the aim you have of building trust through successful trust, Diane is ‘closed’. She is not in a position where your trusting her with the task that you do could—even when that trust is fulfilled by her—contribute to building trust on account of that fulfilment.

Now consider a twist on this case:

DIANE*: You need someone to babysit on short notice. There are a number of people you could ask, however, you choose a local teenager, Diane*, whose parents you know. You have heard that Diane* is troubled, and you have had a standing desire to take Diane* under your

wing in hopes of having a positive influence on her. A first step toward having such a positive influence, you think, will be to establish a bond of trust, a bond you hope to develop by entrusting her with the babysitting task despite her reputation. Diane* is open in principle to building trust with someone who would entrust her with this kind of task. Unfortunately, and unbeknownst to you, Diane* bears a deep-seated grudge against you. Though she babysits the kids successfully (suppose, she needs the money)—and though her doing so successfully in fact contributes to making her more trustworthy *generally* speaking—it plays no role in establishing or strengthening any trust between you and her.

Diane* is *not* closed to building trust through successful trust, as Diane is, generally. However, the conditions in DIANE* are such that they prevent building her trust *with you* through successful trust. This is not to say that Diane*'s grudge would never subside so as to open up such a possibility later. The point is that the situation in which you encounter Diane* is not one in which, were she to come to establish and build trust with you, this could be achieved in the way you're attempting to do so here—viz., through facilitating successful trust via the babysitting task.

There are two interrelated points to draw from the DIANE and DIANE* cases. The first is that it doesn't count against one's overriding therapeutic trust competence, viz., one's disposition to attain the aim of overriding therapeutic trust reliably enough, were one to be *unreliable* at attaining this aim in cases like DIANE or DIANE*, where the conditions are, for different reasons, not suitably conducive to building trust through successful trust. Secondly, and relatedly: the kind of competence that *matters* for overriding therapeutic trust is, accordingly, a disposition to build trust through successful trust reliably enough when one is in conditions that *are* appropriate to doing so, conditions that include at least that openness and reciprocity are satisfied, as they are not in DIANE and DIANE*, respectively.

A further point is that *when* these conditions are met, some are disposed to achieve the aim of overriding therapeutic trust *more* reliably than others. And that is just to say that, when it comes to overriding therapeutic trust, some are *more competent* than others, some of whom simply lack this competence by not being suitably reliable in conditions that are favourable to this kind of trust.

What makes the difference? One factor that's worth noting explicitly is that we vary in the capacities we have to reliably assess trust-building payoffs. For example, recall our case of the CEO (4) who entrusted the low-level employee with a *disproportionately* large task, one which not easily the employee would have managed. The overriding therapeutic trust is unlikely to payoff here simply given that the difficulty of the task choice will make unlikely the building trust through *successful* trust. Conversely, entrusting *too small* a task, with therapeutic aims, is likewise unlikely to payoff, though, for a different reason. (Compare: suppose you were to, with trust-building aims, entrust a teenager not with looking over the house or the kids, but with looking after a small cactus for the weekend). The task is not certainly too difficult to undermine the likelihood that the trustee will take care of things *as entrusted*, but it is *so easy* that it undermines the likelihood that, through being undertaken successfully, it will play a (non-negligible) role in increasing any kind of trust bond with the trustee. In short: (i) a propensity to miss the mark too often in *either* direction will undermine one's reliability at attaining the aim of building trust through successful trust, and so (ii) a competence to hit this aim reliably (when appropriately situated to do so³⁴) requires a capacity for the kind of risk assessment

³⁴ That is: when the trust environment is such that the openness and reciprocity conditions described in this section are met.

that's needed to prevent one from too often 'over' or 'under' trusting (as in the CEO and cactus cases, respectively).

Let's return now to our 'AAA' table. *Accurate* or successful overriding therapeutic trust (i.e., attaining the first 'A') occurs when overriding therapeutic trust hits its constitutive aim, which is the aim of building trust through successful trust. *Adroit* overriding therapeutic trust issues from a *competence* to hit this aim reliably enough when one trusts with a therapeutic aim whilst appropriately situated—where being appropriately situated for this kind of trust requires at least the satisfaction of the openness and reciprocity conditions. *Apt* overriding therapeutic trust can now be defined in terms of accurate and adroit therapeutic trust—viz., apt overriding therapeutic trust is overriding therapeutic trust that is *accurate because adroit*, viz., when one's building trust through trusting successfully manifests one's competence to therapeutically trust successfully reliably enough in appropriate conditions.

Apt overriding therapeutic trust is a kind of *achievement*, just like any kind of aim attained through skill rather than by other means.³⁵ In this respect, apt overriding therapeutic trust stands to *mere* successful overriding trust as knowledge to lucky true belief, and to an archer's successful shot attained through skill to the same success attained any old way. However, as we've seen, the achievement of apt overriding therapeutic trust is a *different* achievement than the achievement of apt (non-therapeutic) trust, one that involves the attaining of a different aim through the manifestation of a different sort of competence.

9.

We began with a puzzle about therapeutic trust and its relationship to ordinary non-therapeutic trust. Three prominent attempts to address this puzzle were considered, and each was shown to be problematic for different reasons. One notable problem common to each of the three accounts was that none was well-suited to explain—given what each maintains, respectively, about therapeutic trust and how it differs from non-therapeutic trust—in virtue of *what* good therapeutic trust differs from plain old bad trust, including incompetent trust that just so happens to result in the building of trust, as well as successful *and* competent trust that builds trust for reasons that have nothing to do with the trust placed.

The account I've proposed has a number of advantages over these accounts. First, it avoids the traps that these other accounts were shown to fall into given their specific commitments. The key move proposed which helps to get things right involves the recognition of two kinds of therapeutic trust. There is a philosophically uninteresting species of therapeutic trust that is implicit in ordinary trusting—what I called *default therapeutic trust*. While default therapeutic trust (trivially) aims at building trust, it does so only because building trust is among the normal social functions of ordinary non-therapeutic trust, which aims constitutively at the trustee taking care of things as entrusted. *Overriding therapeutic trust*, by contrast, has its own constituent normativity—with reference to which we can normatively assess this kind of trust differently from how we normatively assess standard trust. In doing so, we can say *why* each kind of trust is good when it is, without reducing the goodness of either kind of trust to the goodness of the other. Moreover, the view can help us to make sense of how the the skills needed for reliable therapeutic trust come apart from the skills

³⁵ For some representative discussions of the value of achievements understood as having a success-through-ability structure, see, e.g., Bradford (2013); (2015); Sosa (2010b); Pritchard, Turri, and Carter (2018); Greco (2014); Carter and Gordon (2014), Pritchard (2009) and Zagzebski (1996).

needed to be good at trusting well more generally; the ‘AAA’ profiles of adroit trust and adroit therapeutic trust differ in clear ways. Finally, by distinguishing between ordinary apt trust and apt (overriding) therapeutic trust on the model proposed, we have a perspective from which to appreciate two different *achievements* in trusting and why neither of these achievements reduces to the other.

10.

(Objection 1).

On the view proposed, the constitutive aim of overriding therapeutic trust is meant to be distinct from the constitutive aim of standard (i.e., non-therapeutic) trust in that: (i) the aim of the former is that the trustee take care of things as entrusted; whereas, (ii) the aim of the latter is to build (or strengthen) trust *through* successful trust, viz., through the trustee’s taking care of things as entrusted.

However, the suggestion that these aims are distinct is not so clear given that the view *also* holds that building trust is among the normal social functions that is played by (successful) ordinary trust. But if *that’s* right, then isn’t it the case that standard trust constitutively aims not *merely* at the trustee’s taking care of things as entrusted, but also, at this fact playing the social function of strengthening trust? If so, then it looks like the claimed difference between the constitutive aims of standard trust and overriding therapeutic trust collapses.

Reply: The fact that the constitutive aim of ordinary trust—viz., that the trustee take care of things as entrusted—is such that when this aim is met, it’s doing so has a characteristic social function, X, does not imply that its actually playing that function, X, is thereby included *as part of the constitutive aim*. The aim would still be met even if that social function characteristic of attaining that aim were *not* played.³⁶ (Compare: the aim of archery—hitting the target—is attained even if your hitting the target does not play any of the roles that attaining this aim would characteristically play, e.g., to build confidence, solidify social standing with peers, signal competence, etc.). Likewise, if you trust a reliable colleague to deliver an envelope to your boss without reading the message inside, and the colleague successfully does so without taking a peek, there is a clear sense in which your trust placed in your colleague on this occasion has attained *its* aim—no matter what further social functions your trust plays or does not play, including social functions you might reasonably *expect* it to play.

(Objection 2).

The competences involved in ordinary (non-therapeutic) trust and overriding therapeutic trust are claimed to be *different* competences. But is this really so? Here is a reason to think the answer is ‘no’. Adroit overriding therapeutic trust issues from a competence to reliably enough build trust through

³⁶ There is a precedent for this kind of thinking about aims and defective functioning found in Burge (2003, 509). According to Burge, evidence that something is or is not operating defectively offers us insight into what its aim (or, for Burge, function) is (or is not). For example, if we did not regard the heart as defective if it failed to pump blood, then this would cast doubt on the idea that the heart is normatively constrained by the aim of pumping blood. By parity of reasoning: my suggestion is that—in both archery and in ordinary trust—we would not regard a shot as defective if it hit the target but did not inspire confidence (a normal social function which, suppose, hitting a target plays) nor (ordinary) trust as defective if the trustee took care of things as entrusted despite this fact not going on to build further trust. This—with reference to the kind of reasoning we find in Burge—counts against the aim of archery being ‘hitting the target *and inspiring confidence*’, which is surely the right result, and likewise against the aim of ordinary trust as being ‘that the trustee take care of things as entrusted *in a manner than builds trust*’.

successful trust when one attempts to do so while appropriately situated. But then—being reliable at *this* was said to require a capacity for the kind of risk assessment that’s needed to prevent one from too often ‘over’ or ‘under’ trusting (as in the CEO and cactus cases, respectively). But even if that’s right—and here’s the worry—doesn’t being reliable at attaining the aim of *ordinary* trust *also* require a capacity for risk assessment? That is: a competence to attain the aim of ordinary trust reliably enough (when appropriately situated) surely requires a capacity to evaluate risks of *betrayal*, including risks of betrayal generated by, e.g., incentives the trustee has to betray, the difficulty of the task relative to the trustee’s perceived abilities, etc. But once these points are granted, the distinction between the substance of the competences relevant to (i) ordinary trust versus (ii) overriding therapeutic becomes blurred.

Reply: In short, the answer is the kind of competence that matters for overriding therapeutic trust *asymmetrically entails* the kind of competence that matters for ordinary (non-therapeutic trust). While risk assessment is undeniably important to both kinds of competences, and thus to both adroit overriding therapeutic trust as well as adroit ordinary trust, the kind of risk assessment that competent overriding therapeutic demands is more sophisticated, and accordingly more demanding, than the kind of risk assessment that competent ordinary trust demands. Given that the constitutive aim of ordinary therapeutic trust (that the trustee take care of things as entrusted) is a *component* of the constitutive aim of overriding therapeutic trust (that trust is built through successful trust—viz., through the trustee taking care of things as entrusted), reliably attaining the latter will require the very same kind of risk assessment needed to reliably secure the former, *plus* the capacity to assess *additional* risks—risks specifically to the non-obtaining of *trust built through successful trust*. Ordinary trust competence doesn’t demand one have the capacity to assess these further risks.

The above illuminates an interesting wider point about the difference between ordinary and overriding therapeutic trust, which is that the latter is—in short—*more difficult* to do well. Being competent at overriding therapeutic trust requires all the skills required to be competent at ordinary trust, plus others which the latter doesn’t require. A corollary is that the achievement of *apt* overriding therapeutic trust is more substantial, arguably more valuable³⁷, than the achievement of *apt* ordinary trust, in that the former issues from a comparatively more sophisticated and demanding kind of competence to acquire and exercise.

(Objection 3).

Certain kinds of ‘forced-choice’ cases seem like they would work as counterexamples to the proposed account of overriding therapeutic trust. Consider the following:

FORCED CHOICE: You’ve just moved to an apartment building in a new city, where the only person you know is teenager who lives in the flat below you. You need to leave town for the weekend—suppose your job depends on it—and need someone to feed, water and walk your dog (you’ve tried kennels, etc., and all are fully booked). Your hand forced, you trust the teenager who lives in the flat below you with this task—someone whom, had you had a better range of options—you wouldn’t have chosen, as they’ve not established any track-record yet of responsibility with you, and your dog’s welfare is important to you.

³⁷ For some notable arguments that difficulty adds value to achievement, see, e.g., Bradford (2013), (2015), and Pritchard (2014).

Two things seem, *prima facie*, to be true in FORCED CHOICE. First, (i) it looks like a case of therapeutic trust of a philosophically interesting sort (you are, after all, placing trust in a teenager to whom you wouldn't ordinarily trust a task like this); but, second, (ii) it doesn't get ruled in on the account proposed. This is because in FORCED CHOICE, it's *not* the case that the aim of ordinary trust is *not* sufficient to lead the truster to risk trusting. That aim *is* sufficient, *ex hypothesi*.

Reply: My response to FORCED CHOICE is to accept (ii) and press back against (i). It is a mistake to think that all cases in which one trusts a non-ideally suited trustee (e.g., by selecting someone regarded as being less reliable than would be preferred) are, in virtue of this, 'therapeutic' in some interesting sense. On the view I've proposed, therapeutic trust of a philosophically interesting sort misses its mark—viz., is defective—even if the trust is successful, provided the trust fails to build or strengthen through this successful trust. FORCED CHOICE, however, is a case where the trust placed in the teenager *succeeds perfectly well* so long as the teenager takes care of the dog as entrusted. This is so in a way that is not interestingly different than were the teenager perceived to have been much more reliable than they are actually perceived to be. The situation is, however, very different if we suppose that the basic aim of successful trust *weren't* enough (as it is in FORCED CHOICE) to motivate you to risk trusting the teenager—viz., as would be the case when you trust the teenager, e.g., rather than someone you think has a better track record, with the aim of using successful trust to building trust. The above diagnosis not only explains why we would be wrong to, in short, lump all high-risk cases together—but it also helps to highlight the important sense in which therapeutic trust of the philosophically interesting kind *uses* trust in a way that ordinary trust does not.³⁸

³⁸ This paper was written as part of the Leverhulme-funded 'A Virtue Epistemology of Trust' (#RPG-2019-302) project, which is hosted by the University of Glasgow's COGITO Epistemology Research Centre, and I'm grateful to the Leverhulme Trust for supporting this research.

REFERENCES

- Adler, Jonathan E. 1994. "Testimony, Trust, Knowing." *The Journal of Philosophy* 91 (5): 264–75.
- Alfano, Mark. 2016. "The Topology of Communities of Trust." *Russian Sociological Review* 15 (4): 30–56.
- Baier, Annette. 1986. "Trust and Antitrust." *Ethics* 96 (2): 231–60.
- Baker, Judith. 1987. "Trust and Rationality." *Pacific Philosophical Quarterly* 68 (1): 1–13.
- Bradford, Gwen. 2013. "The Value of Achievements." *Pacific Philosophical Quarterly* 94 (2): 204–24.
- . 2015. *Achievement*. Oxford University Press, USA.
- Bratman, Michael E. 1992. "Practical Reasoning and Acceptance in a Context." *Mind* 101 (401): 1–15.
- Buckareff, Andrei. 2010. "Acceptance Does Not Entail Belief." *International Journal of Philosophical Studies* 18 (2): 255–61.
- Carter, J. Adam. 2020. "On Behalf of a Bi-Level Account of Trust." *Philosophical Studies* 177: 2299–2322.
- . Forthcoming. "Trust as Performance." *Philosophical Issues*.
- Carter, J. Adam. 2021. "De Minimis Normativism: a New Theory of Full Aptness." *The Philosophical Quarterly*, 71(1): 16-36.
- Carter, J. Adam and Emma C. Gordon. 2014. *American Philosophical Quarterly*, 51(1), 1-13.
- Chrisman, Matthew. 2012. "The Normative Evaluation of Belief and the Aspectual Classification of Belief and Knowledge Attributions." *Journal of Philosophy* 109 (10): 588–612.
- Cohen, L. Jonathan. 1989. "Belief and Acceptance." *Mind* 98 (391): 367–89.
- Domenicucci, Jacopo, and Richard Holton. 2017. "Trust as a Two-Place Relation." *The Philosophy of Trust*, 149–60.
- Emmons, Robert A, and Michael E McCullough. 2004. *The Psychology of Gratitude*. Oxford University Press.
- Faulkner, Paul. 2007. "On Telling and Trusting." *Mind* 116 (464): 875–902.
- . 2011. *Knowledge on Trust*. Oxford University Press.
- Fernández Vargas, Miguel Ángel. 2016. *Performance Epistemology: Foundations and Applications*. Oxford University Press.
- Frost-Arnold, Karen. 2014. "The Cognitive Attitude of Rational Trust." *Synthese* 191 (9): 1957–74.

- Geach, Peter T. 1956. "Good and Evil." *Analysis* 17 (2): 33–42.
- Goodman, Nelson. 1954. *Fact, Fiction, and Forecast*. Harvard University Press.
- Greco, John. 2010. *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge University Press.
- Hall, Mark A. 2005. "The Importance of Trust for Ethics, Law, and Public Policy." *Cambridge Quarterly of Healthcare Ethics* 14 (2): 156–67.
- Hawley, Katherine. 2014. "Trust, Distrust and Commitment." *Noûs* 48 (1): 1–20.
- . 2019. *How to Be Trustworthy*. Oxford University Press, USA.
- Hieronymi, Pamela. 2008. "The Reasons of Trust." *Australasian Journal of Philosophy* 86 (2): 213–36.
- Holton, Richard. 1994. "Deciding to Trust, Coming to Believe." *Australasian Journal of Philosophy* 72 (1): 63–76.
- Horsburgh, Howard John Neate. 1960. "The Ethics of Trust." *The Philosophical Quarterly (1950-)* 10 (41): 343–54.
- Jones, Karen. 1996. "Trust as an Affective Attitude." *Ethics* 107 (1): 4–25.
- . 2004. "Trust and Terror." In *Moral Psychology: Feminist Ethics and Social Theory*, edited by Peggy DesAutels and Margaret Urban Walker, 3–18. Rowman & Littlefield.
- Keren, Arnon. 2014. "Trust and Belief: A Preemptive Reasons Account." *Synthese* 191 (12): 2593–2615.
- . 2019. "Trust, Preemption, and Knowledge." *Trust in Epistemology*.
- McGeer, Victoria. 2008. "Trust, Hope and Empowerment." *Australasian Journal of Philosophy* 86 (2): 237–54.
- McHugh, Conor. 2012. "The Truth Norm of Belief." *Pacific Philosophical Quarterly* 93 (1): 8–30.
- McKittrick, Jennifer. 2018. *Dispositional Pluralism*. Oxford University Press.
- McLeod, Carolyn. 2002. *Self-Trust and Reproductive Autonomy*. MIT Press.
- . 2015. "Trust." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Fall 2015. <https://plato.stanford.edu/archives/fall2015/entries/trust/>; Metaphysics Research Lab, Stanford University.
- McMyler, Benjamin. 2011. *Testimony, Trust, and Authority*. Oxford University Press USA.
- Pritchard, Duncan. 2009. "Knowledge, Understanding and Epistemic Value." *Royal Institute of Philosophy Supplement* 64: 19–43.
- . 2012. "Anti-Luck Virtue Epistemology." *Journal of Philosophy* 109 (3): 247–79.
- Quine, Willard Van Orman. 1960. *Word and Object*. MIT press.

- Ryle, Gilbert. 1949. "The Concept of Mind." *London: Hutchinson*.
- Simion, Mona. 2019. "Knowledge-First Functionalism." *Philosophical Issues* 29 (1): 254–67.
- Simion, Mona, Christoph Kelp, and Harmen Ghijsen. 2016. "Norms of Belief." *Philosophical Issues* 26 (1): 374–92.
- Solomon, Robert C., and Fernando Flores. 2003. *Building Trust: In Business, Politics, Relationships, and Life*. Oxford University Press USA.
- Sosa, Ernest. 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume 1*. Oxford University Press.
- . 2015. *Judgment and Agency*. Oxford: Oxford University Press.
- Stocker, Michael. 1979. "Desiring the Bad: An Essay in Moral Psychology." *The Journal of Philosophy* 76 (12): 738–53.
- Turri, John. 2016. "Knowledge as Achievement, More or Less." *The Present and Future of Virtue Epistemology*.
- Wanderer, Jeremy, and Leo Townsend. 2013. "Is It Rational to Trust?" *Philosophy Compass* 8 (1): 1–14.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Zagzebski, Linda Trinkaus. 1996. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge University Press.