

You *can* trust the ladder, but you shouldn't

1. Introduction

My claim in this paper is that, contra what I take to be the orthodoxy in the wider literature, we do trust inanimate objects—per the example in the title, there are cases where people really do trust a ladder (to hold their weight, for instance), and, perhaps most importantly, that this poses a challenge to that orthodoxy. My argument consists of 4 parts. In section 2 I introduce an alleged distinction between trust as mere reliance and trust as a rich, morally loaded notion. In the course of doing so, I briefly sketch some models of trust. This will help us to get a handle on how some of the mainstream models of trust preclude or ignore the trusting of inanimate objects. In section 3 I introduce cases designed to show that we have reason to think that we *do* trust inanimate objects. In section 4 I consider an objection, and, in section 5 I discuss some consequences of this finding, before concluding in section 6.

2. Kinds of trust

In the literature on trust it's common enough to see a two-fold distinction drawn between different kinds of trust. Following Faulkner (2015), I'll distinguish these kinds of trust as follows. First, we have trust as a two-place predicate, such that 'x trusts y'. Example: Tasha trusts Marvin. Second, an agent, x, may trust another agent, y, to perform some action, φ : x trusts y to φ . Call this 'contractual trust'. 'Contractual trust' is a three-place term, connecting two agents to some action. Example: Tasha trusts Marvin to lock the door. The third main notion associated with trust is general trustworthiness. 'Trustworthiness' is a one-place predicate. Example: Tasha is trustworthy. Of these notions of trust, this paper will be focused only on the notion of contractual trust: x trusts y to φ .

Within the literature that concerns itself with contractual trust, it's typical to see a further distinction drawn between trust as 'mere reliance' and a richer, morally loaded notion of trust. I cannot put the relevant points any better than Hawley (2014a: 2):

we often rely upon inanimate objects but we do not grant them the rich trust we sometimes grant one another; inanimate objects can be reliable but not genuinely trustworthy. Moreover our reactions to misplaced trust differ from our reactions to misplaced reliance. Suppose I trust you to look after a precious glass vase, yet you carelessly break it. I may feel betrayed and angry; recriminations will be in order; I may demand an apology. Suppose instead that I rely on a shelf to support the vase, yet the shelf collapses, breaking the vase. I will be disappointed, perhaps upset, but it would be inappropriate to feel betrayed by the shelf, or to demand an apology from it.

In what follows, I will speak of *mere reliance* and of *rich trust*, by way of short hand for these different notions. My argument in short is that whilst I agree with Hawley and others that (e.g.) it would be *inappropriate to* trust the shelf to hold the vase; to feel betrayed by the shelf if the vase falls; to demand an apology from it; to *blame* the shelf: nonetheless, I do not see that as any barrier to doing so. I am capable of doing the inappropriate (as those who know me can testify), as are many others. Demonstrating *that* this is (that we can and do trust inanimate objects) so and *that* this poses a challenge to the orthodoxy is my main aim here.

To get the ball rolling, we can start by noting that on Hawley's account we cannot have rich trust in inanimate objects.

Here is my account of trust and distrust:

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... (2014a: 11)

An inanimate object is not properly described as 'someone'. That being so, aside from the quote given above where Hawley *explicitly* states that we do not richly trust inanimate objects ('we often rely upon inanimate objects but we do not grant them the rich trust we sometimes grant one another'), her very definition of contractual trust prohibits rich (contractual) trust from being granted to inanimate objects. Elsewhere, Hawley (2014b: 2029-30—my emphasis) appears similarly clear:

There are rival philosophical accounts of *what trust is*. But most have a common underlying form: trusting someone to do something involves relying upon her to do it, plus some extra factor *which distinguishes genuine trust from the attitude of (mere) reliance we may take to inanimate objects like climbing ropes, cars or crash helmets*.

To unpack: the first sentence tells us that there are various analyses of what trust consists in—of what trust itself *is*. The first part of the second sentence tells us that these accounts of trust have a common form, and the second part tells us that these accounts distinguish cases of trusting someone to do something, from other cases, where we take an attitude of (mere) reliance, which is taken to be the attitude of trust we have towards inanimate objects. Thus, this seems to give us a fairly clear pronouncement to the effect that most philosophical accounts of trust distinguish genuine trust (that we have in people) from mere reliance (an attitude we may take to inanimate objects—though it is certainly true that we frequently merely rely on people, too).

Richard Holton (1994: 67—my emphasis) makes remarks that appear equally suggestive:

I think that the difference between trust and reliance is that trust involves something like a participant stance towards the person you are trusting. When you trust someone to do something, you rely on them to do it, and you regard that reliance in a certain way: you have a readiness to feel betrayal should it be disappointed, and gratitude should it be upheld. In short, you take a stance of trust towards the person on whom you rely. It is the stance that make the difference between reliance and trust. *When the car breaks down we might be angry; but when a friend lets us down we feel betrayed.*

Karen Jones (2006: 6—emphasis mine) takes a similar line:

[Three-place t]rust *is* accepted vulnerability to another person's power over something one cares about, where (1) the truster foregoes searching (at the time) for ways to reduce such vulnerability, and (2) the truster maintains normative expectations of the one-trusted that they not use that power to harm what is entrusted.

Jones here explicitly defines trust *as* a relation that one stands in to another *person's* power.

Similarly, Pettit (2004: 111, see especially Fig 1) categorises reliance as coming in two forms: on persons and on non-persons. Trust is then categorised as a particular kind of reliance on persons. That being so, given this is a categorical scheme, it follows that we *do not* trust non-persons, such as inanimate objects, to perform particular actions.

Baier (1986: 259) describes trust as follows: '[t]rust, I have claimed, is reliance on others' competence and willingness to look after, rather than harm, things one cares about which are entrusted to their care.'

Ladders, chairs, shelves: these objects, all inanimate, are neither willing nor unwilling. As such, we cannot (in current terms, *richly*) trust them. Given Baier's entirely justified influence in the literature, I think that this is significant.

Finally, at the outset of her SEP entry, McLeod (2015—emphasis mine) also appears to rule out the idea that trust extends to inanimate objects:

Trust *is* an attitude that we have *towards people* whom we hope will be trustworthy, where trustworthiness is a property, not an attitude.

This is an identity claim and one that identifies trust as a relation that we (as people) have towards other people.

Purely from the point of view of getting a feel for the field, I take it that McLeod's remarks are perhaps the most salient here. An SEP article is a gateway to a topic, written by an expert in the field, who is seeking to present that field to the interested reader. And, as she defines it right from the outset, McLeod presents trust *as being an attitude that we have towards people*.

Others *appear to* pursue a slightly different approach. For example, Hieronymi (2008: 213) is explicit that she is restricting her focus to cases of contractual trust *between people*. One natural way to see my paper in response to those who restrict in this way is that they are providing us only with a partial account of contractual trust. I don't think that we can say that same about those cited a little earlier, such as Hawley, Jones or McLeod who *seem* to be offering us very generally claims about what trust is, without any *stated* restrictions in place (I'll come on to the possibility of tacit restrictions in a moment).

So, here is my guiding thought: it seems apparent that the orthodoxy identifies trust as a relation that people stand in to other people. The analyses of trust are, *I think*, most reasonably interpreted as ruling trust inanimate objects impossible.

There is scope for disagreement. Here is a line of reply that one might offer to the thought that I'm correct and that there is an orthodoxy that rules out our trusting inanimate objects. Hawley (2014a: 11) says that, "[t]o 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." Hawley is, however, providing a definition of what it is to trust or distrust *someone*; that doesn't prevent us from having a further account of what it is to trust inanimate objects, and even though Hawley precedes her definition by describing the account as her 'account of trust and distrust', not merely of trust or distrust *in someone*, we shouldn't read anything of significance into that framing.

Similarly, Baier is concerned with the ethics of trust. That being so, it might be reasonable to view Baier as, albeit tacitly, thinking *only* about trust as it involves persons. And, finally, Pettit is concerned with the rationality of trust. That being so, it's plausible to simply view his paper as offering us an account of trust that is restricted in its scope, to tell us only about cases involving the trust of people. So, treating these cases as representative, rather than having an orthodoxy that's hostile to or inconsistent with the possibility of trust in inanimate objects, might it not be the case that these philosophers have, to this point, simply been tacitly restricting their target, and ignoring cases involving inanimate objects?

For what it's worth, I find that position somewhat tricky to sustain given what's said. When Hawley (2014b: 2029-30) surveys the field she says that "[t]here are rival philosophical accounts of *what trust is*. But

most have a common underlying form: trusting someone to do something involves relying upon her to do it, plus some extra factor'; this seems to be a clear move from seeking an analysis *of what trust is*, to then treating that very thing—that very account *of what trust is*—as something that obtains between people. And, as she goes on in that quotation, trust is to be differentiated from the mere reliance that we have in inanimate objects.

Another example of an author proceeding when in this way when taking a view of the field can be drawn from Simpson's recent work. Simpson (2012: 550) begins by outlining what he takes to be the dominant methodology within the philosophical literature on trust. I quote extensively to avoid the charge that I am cherry picking at all.

The task of this article is to elucidate the concept of trust. The paradigm way that analytic philosophers have addressed tasks of this nature is through conceptual analysis, assuming that the rules of correct use are determinate and therefore statable, perhaps because there is some particular phenomenon to which the term refers. At its zenith, the project issues in a definition stated in necessary and sufficient conditions, using only terms understood more clearly than the definiendum. The influential philosophical treatments of trust have generally assumed something akin to this: 'trust is this'.

Simpson (2012: 551) does not in fact believe that trust *is* amenable to conceptual analysis. At the outset of section 2 of the paper, he states the following:

Why should trust not be amenable to conceptual analysis? Against some, I do not think that conceptual analysis is always a hopeless project. Nor do I have an argument to show that trust must be too heterogeneous to be defined, so I have no knock-down reply to the question. But even if no conclusive reasons are available, I suggest that some forceful ones can be given for not persevering with attempts to analyse trust, and for trying a different way.

An inductive reason is the ease with which counter-examples can be produced to existing definitions.

At this point, then, Simpson has noted a general philosophical project: providing a conceptual analysis of trust. He has not looked to provide caveats, such that the project of providing a conceptual analysis of trust is to be understood as the project of providing a conceptual analysis of *trust in people*, and, at least so far as I can see, there is nothing in the framing that suggests that Simpson would intend the project to be shaped in that way. He then goes on to note that: '[t]hree major ways of thinking about trust have emerged in the literature' (Simpson, 2012: 551). The three major ways he identifies are those positions taken by Baier, Jones and Holton. And, as we saw above, *if* we treat those authors as providing a conceptual analysis of *trust itself* (as opposed to *trust in people*), then these analyses preclude our placing

trust in inanimate objects. Simpson thus appears to endorse my view of the literature. As does Hawley in her summary of the field. The prevailing view thus seems to be that trust—what trust *is*—is something that can only be given to objects that are not inanimate.

But, if we really can explain all of this away by appeal to tacit restriction, supposing that each philosopher is happy to simply *ignore* cases of trust in inanimate objects, then I think that we have a different challenge to the orthodoxy. If we can show that there are cases where we do trust inanimate objects, then we *should not* rule out cases of trust in inanimate objects from our discussions precisely because discussion of these further cases will, very likely, tell us something interesting about the nature of trust (itself). At the very least, cases of trust in inanimate objects will provide a plausibly fertile ground for developing counterexamples to accounts of trust. Indeed, and to come full circle and return to Hawley's (2014a) work on trust, a key part of her (I think crucial) methodological insight in that paper is that our analysis of distrust can tell us something of value about trust, and that as such we should not ignore distrust in our analyses. Similarly, then: if we *can* have trust inanimate objects, we *should not* ignore it when trying to understand what trust is, even if we are ultimately more interested in cases of interpersonal trust.

To see why, note the following. Suppose that we have trusting-people-to as one kind of trusting-to, and trusting-inanimate-objects as another kind of trusting-to. That would leave a further question in need of an answer: if these are both instances of trusting-to, then in virtue of what is this so? That is, if trusting-people-to is one kind of trusting-to, and trusting-inanimate-objects-to is another kind of trusting to, then what is trusting-to *simpliciter*? And, if there is no such thing, then in virtue of what do these two kinds of trusting-to warrant the name? Certainly, it would be most natural to suppose that there *is* an umbrella form of trusting-to such that these are both instances of it. For more on the question of why to think that the possibility of trusting inanimate objects may impact how we think about trust, see discussion in section 5.

I'm not going to dwell on the details of the many competing accounts of trust discussed here (nor will I later try to offer a corrective to these different accounts—though see section 5 for some limited discussion). All that matters is that each of these accounts of trust either precludes rich trust in inanimate objects, or else simply offers a restricted analysis of trust and thereby ignores cases that might involve inanimate objects.

I claim that debates about the nature of three-place trust should ensure that the analyses provided are compatible with instances of trust in inanimate objects and that they include such cases in their considerations. To show this, in section 3 I introduce some cases in which it appears correct to say that we can trust inanimate objects. In section 4 I'll explore the objection that I've not properly identified an orthodoxy and in section 5 some of the consequences of this result.

3. Cases

With those distinctions drawn, and accounts of trust given, I now want to describe the cases (more are mentioned in the footnotes—see fn3, fn5 and fn8) in which I think it's right and proper to agree that an agent richly trusts an inanimate object.¹ The general strategy in each case will be as follows. I will describe a case in which I think it's reasonable to say that the putative trustor really does trust some inanimate object. I will *suggest* (though this is speculative on my part) that the reason that the putative trustors engage in that trust is because they have, to some extent, likely anthropomorphised the object in question. I do not say that they have strongly committed to attributing to it beliefs, desires, and so on. Rather, I suggest that it looks reasonable in each case to suppose that the putative trustor is treating the inanimate object *roughly as if* it's some kind of agent. Here is the first case.

Wiley is 4; old enough to read a little, but young enough to still be credulous. Wiley's most treasured possession is his blackboard and chalk set. He loves writing and drawing on it. One day, Wiley comes down for breakfast to see a prediction about the weather stated on the blackboard: 'today it will rain'. Lo and behold, it does! The next day, a similar thing happens. This time the prediction is: 'today the sun will shine'. And, it does! The same pattern repeats itself each day, with the prediction on the blackboard coming true. Wiley forms the view that the blackboard is magical. Somehow, he thinks, the blackboard is communicating with him about the weather. Now, as it happens, what's going on here is that Wiley's sister, Jerlene, is reviewing the weather forecast and writing it onto his blackboard. She doesn't tell him that she's doing it and, overhearing some of the things that Wiley is saying to the blackboard, over time also starts to include responses to what Wiley says in her messages. She thinks it a wonderful prank. Over time, Wiley comes to trust the 'magic' blackboard to tell him what the weather will be (among other things); at least, he says that he does. (For reasons that may already be apparent, and that are associated with her tendency to be somewhat unkind to Wiley, Wiley does not trust Jerlene.)

Now suppose that Jerlene starts to use the blackboard to manipulate Wiley into wearing unseasonable clothing (t-shirt and shorts in the middle of winter, for instance), and suffering as a consequence. Wiley blames his magic blackboard for misleading him. To borrow language from the literature, Wiley will have many of the 'reactive attitudes' that are associated with a breach of trust (resentment, betrayal, disappointment, etc.).² Wiley is being tricked into forming these attitudes towards the blackboard; it is not appropriate for him to form these attitudes: nonetheless, form them he does. In this case, Wiley will self-report as trusting the blackboard and seem to be engaging with the board as a moral agent; blaming it, resenting it, being disappointed by it, etc. This seems to be good evidence that Wiley trusts the blackboard to tell him particular things about the weather. Certainly, there is nothing conceptually incoherent about Wiley coming to trust, and not merely relying upon, the blackboard to tell him things,

¹ Even if, and we must not lose sight of this, they *should not have done*. My claim that it is true to say that agents trust inanimate objects is not to be mistaken for the claims that it is appropriate for them to do so.

² For the general claim that these kinds of reactive attitude are signifiers of trust, *inter alia*, Holton (1994); Jones (1996).

and experiencing these reactive attitudes if that trust is broken. Of course, Wiley *shouldn't* trust the blackboard. He *shouldn't* blame the blackboard if it (/Jerlene) misleads him. Wiley just shouldn't have these reactive attitudes towards the blackboard. After all, it's a blackboard; it is an inanimate object. But, again, that is not the point. The point is that he *can* come to richly trust the blackboard, and that is all that we require—that an agent *can* trust an inanimate object to carry out particular actions.

A further concern: isn't it the case that Wiley trusts, because he (falsely) ascribes the blackboard an agency it does not have? Wiley ascribes agency and thus potentially even intentionality to an object. Wiley, we may say, takes a participant stance towards the blackboard (c.f. Holton (1994)). We don't ordinarily do that. It is plausible in this case to think that Wiley trusts the blackboard only because we are imagining the blackboard as being perceived by Wiley as magical; that is, as having agency. Otherwise this would be mere reliance.³

I'm not sure how this is supposed to help my opponent. I agree with this imagined opponent that rich trust is not appropriate.⁴ But, *if* Wiley is ascribing that agency, then that error is precisely what is doing the work for me in this case: trust occurs *because* of the error in agency ascription. This will be the key point to which we will return: contractual trust is entered into by agents *because* they make an error of some kind.

The second kind of case of richly trusting inanimate objects is one that trades upon superstition. Consider, for instance, a case where a member of a police force comes to 'trust' some particular item of body-armour that has saved their life on previous occasions. Our police-officer, Mel, refuses new armour, despite the fact that those around Mel insist that it would be for the best to have it replaced. In fact, Mel comes to form such a close bond with the armour that she refuses to allow the armour to be maintained by those who would normally do so, resorting to subterfuge to ensure that no-one else lays a hand on her armour. Mel believes: 'I trust this body-armour to protect me; no-one is allowed to touch it, but me'. Aside from being a little delusional, how should we understand Mel's state of mind here? Well, at least as I am describing the case, having been in many scrapes together, Mel has (irrationally) come to the view that her body-armour is somehow special; that without it she may come to harm, and, though she does

³ A similar kind of case: a child might trust a teddy or similar toy to do something for them and might well feel let down if the action is not performed. For instance, my daughter once self-described as trusting her bear to look after her in the dark, became scared in the dark, and then resented the bear for not having done a better job of looking after her. Of course, she clearly shouldn't have trusted the bear to look after her, but (at least so far as *I* can tell—and I was there) she really did trust the bear to look after her. Relatedly, a somewhat unkind piece of satire at the expense of President Bush (Jr), describes a fictional case where he felt betrayed by a magic 8 ball. <http://frowl.org/stuff/8ball.html> Google searches will reveal to the reader that a number of internet users will certainly describe themselves as trusting magic 8-balls, as well. The point, again, is simply that people can have the requisite, trust-related, reactive attitudes towards inanimate objects, whilst relying upon them, and that this is good evidence that they can also trust them to perform particular tasks.

⁴ Actually, that's a little incautious. Given Wiley's position, I think that from some points of view it's actually perfectly appropriate for him to trust the blackboard. It is inappropriate for him to do so only in the sense that his trust is based on a lie.

not go so far as to impute beliefs or commitments to the armour, nonetheless takes herself to have formed a tight bond with it and has an attitude (and emotional feeling of) trust towards the body-armour.⁵

Equally, it is not at all unreasonable to suppose that, in the event of the armour failing, Mel will experience many of the ‘reactive attitudes’ that are standardly taken to indicate if x’s trusting y to perform some action is betrayed (cf. Faulkner 2014: 335). To illustrate: if the armour fails her (perhaps a soft-spot has developed that fails to deflect a blow from a criminal she is pursuing), Mel might respond with feelings of resentment; feelings of betrayal; a feeling that the armour has let her down. These reactive attitudes may be thought to function a *little* like markers of trust. If they are present in such a case (and, again, it seems reasonable to suppose that they may be), then that gives us reason to think that trust is present. (Of course, a determined proponent of one or other of the models of trust might simply insist that this is not an instance of trust. But what I cannot see is why one would take that approach. Simply: let us begin by looking at a plausible case. Since this looks like a case of trust—with the right kinds of reactive attitude present, but in which there is no attribution of agency, for these attitudes do not obviously require any such attribution—so we should treat it as such.)

Similar cases of superstition-leading-to-trust could be generated out of any form of religion that advocated the worship (and rich trust) of an object that is in fact inanimate, but to which agency was nonetheless ascribed. If, for instance, we agree that the sun lacks agency, then the early Egyptian religion of Atenism, that advocated the worship of the sun (not a figure, such as Amun, who was supposed merely to reside within the sun) might present an interesting case. Those who followed Atenism could, quite reasonably I shall assume, have described themselves as trusting the Sun to deliver particular protections and bring particular goods. They would have been wrong to do so (or so I shall assume), for the sun is not in fact animate (with apologies to panpsychists). Nonetheless, they could very well have trusted the sun to protect them/to deliver goods unto them, because they wrongly ascribed agency to it.

In such cases, it seems to me that agents do in fact trust the object in question and, if it failed to protect them/deliver the desired goods, would not necessarily blame those *in fact* responsible for protecting them or delivering goods unto them, but might well simply feel betrayed by the inanimate object itself. In short: if the harvest failed then the Atenists would likely blame the sun, not the farmers. Again, I would not dispute that such an agent is doing something *wrong*; something that they ought not to do. Certainly, Mel is likely irrational to trust the body-armour in this way, and certainly quite wrong to blame it for any injury

⁵ At the risk of engaging in autobiography I have what I take to be a similar attitude towards my bike. We’ve been in lots of close scrapes together. I think I trust it. If the bike was to fail, I would likely (perhaps briefly) feel let down by it. I would experience the same kinds of reactive attitudes to it as I do to *someone* who lets me down. (Indeed, on one memorable occasion the details of which are beyond the scope of this paper, I did exactly that). I don’t think it has beliefs, commitments, or what have you. Nonetheless, the term ‘trust’ seems to me to apply, quite reasonably, and not just in the sense that I can rely on it. The reader is entitled to make what they will of my irrationality on this score.

she suffers; similar thoughts *may* apply to the superstitious sun-worshipper. But, to repeat a now familiar refrain, that an agent should not behave in some way is no guarantee that they will not do so.

Of course, what is common to the cases that I've described is that it *seems* as if the agent that I am alleging is doing the trusting is treating the inanimate object it trusts, *as an agent*. Wiley is treating the blackboard *as* an agent that is writing messages to him. Mel is treating her body armour *as* an agent that has a duty of protection towards her. Atenists, we are supposing, imputed agency to the sun. And from many points of view, I think that should be unsurprising. We know that humans anthropomorphise, typically as a quick and dirty way of forming predictive inferences about the world around them, or else in order to provide psychologically appealing narrative explanations of various occurrences. So much is not news.⁶

Abstracting away from the cases, the general challenge to those in the literature whose accounts (if I have interpreted them correctly) imply that we *cannot* trust inanimate objects may then be put thus: why should we suppose that it's *impossible* for an agent to trust an inanimate object? (Note, I don't think any of the aforementioned philosophers explicitly say *that* we cannot trust inanimate objects. My claim is that this result is simply a product of the way in which they analyse rich, contractual trust.) This seems a bold and unreasonable claim to make in the absence of both strong argument and clear empirical data. Since neither have been provided, I think that we ought *at the very least* to suspend judgement about whether it is possible to trust an inanimate object; better, I think, would be to concede that it *is* possible, after all. And, as a secondary point, my suspicion is that what tempts us into trusting in the cases described is the simple fact that we anthropomorphise. However, though the cases appear to leave that option open—and I'll say something a little more about it later on—it is speculation on my part. As above, the core aim of the paper is to show that we can trust the inanimate and to explore the consequences for the orthodoxy. I take it that these cases demonstrate that we can.

4. Objection: there is no orthodoxy

Few philosophical claims are uncontentious. Those made here are no exception. I am sure that the reader can think of many more, but in this section I want to consider what I take to be the most pressing

I've suggested that the accounts of contractual trust given by Baier, Pettit and Hawley (and others) are not compatible with our being able to trust inanimate objects, or else that they simply ignore cases involving inanimate objects. But, so goes the thought, there are other accounts. These other accounts are (or may be, for all that I've said) compatible with our trusting inanimate objects. I stand charged of over-reaching.

I don't have a definitive view on when a group of positions might be regarded as capturing orthodox sentiment. I'm happy enough to concede ground to my opponent here. If the objection I'm considering

⁶ See, e.g., Epley, Waytz and Cacioppo (2007) for an overview of some of the psychology here.

here is just to my use of the term 'orthodox', then so be it: I shall retract. So, for instance, Owens' (2017) offers an account of trust according to which, 'for X to trust Y is for X to engage with Y in a way that will realize the value distinctiveness of Y-things (provided this thing is a good Y)'. In discussion of promising, Owens moves from this to contractual trust (or in his terms 'X trusts that Y will A'). 'One trusts *that* a promise will be kept whenever one extracts some value or other from the promise by relying on its being kept'. Were one persuaded by what Owens' has to say here, then perhaps this could be generalised into an account of contractual trust that permits the trusting of inanimate objects.

For what it's worth, that strikes me as correct. But recall: my claim is not that *everyone* in the literature is prone to my concerns here. My target is but a part of that literature. If it turns out that Owens can give us an account of contractual trust that renders it compatible with our trusting inanimate objects—well then; all the better for Owens' account.⁷ Similarly, according to Holton, 'the difference between trust and reliance is that trust involves something like a participant stance towards the person you are trusting....trusting someone is one way of treating them as a person.' (1994: 4)

For reasons that may be obvious, I'm very sympathetic. In the cases of trust given to inanimate objects, my claim is that there is plausible anthropomorphism involved—at least to some degree. This aligns nicely with Holton's requirement that trusting someone, or something, is one way of treating them as a person. So, my position is not (and should not be taken to be) that there is no account of contractual trust that correctly allows that we trust inanimate objects.

Nonetheless, I think that this is all grist to my mill. Suppose that certain accounts of trust (those offered by the likes of Holton, for instance) carry over very naturally to dealing with cases where we trust inanimate objects and that others (perhaps those offered by Hawley and Jones—see section 5 for a little more discussion) don't. That would seem to speak squarely in favour of those accounts of trust where we can port our analysis to encompass cases where we trust inanimate objects, and against those where we cannot.

5. Consequences

What if I'm right? I think that at least three consequences follow. First, and this is obvious, it turns out that the orthodoxy is false or overly restricted: we *can* trust inanimate objects (even if we should not) and this requires our attention. That's a not insignificant point and one that needs to be recognised. Second, if we take their analyses at face value *as* analyses of what trust is (without tacit restriction), then those who define rich trust in such a way as to preclude the trusting of inanimate objects (that is, if I have the right

⁷ Nonetheless, it isn't clear to me that we can generalise from Owens' account. Abstracting away from his discussion of promising would seem to give us the following account of contractual trust: x trusts y to φ when x extracts some value from relying on y φ -ing. And that generalisation does not look to be right. If I trust my wife to pick up our daughter from School, and she fails to do so, then I extract no obvious value from relying on her to do so.

reading of the likes of Baier, Hawley, Jones and Pettit), would have to revise their accounts. If it is possible to richly trust an inanimate object then this is nothing more than an obvious consequence of what I say here.

More, I don't think that it would be entirely straightforward to make the necessary modifications—though I could be wrong about this. Since there are many such accounts of what trust is, with very different structures, I won't seek here to try to provide a fix. Simply, it strikes me as unlikely that there is any *single* fix that would apply to all such accounts of trust. Nonetheless, it's easy to see that some parts of this analysis will have to change. To illustrate, let us return briefly to Hawley's definition of trust (assuming we read it as an analysis of trust *generally*, and taking Hawley literally when she says that this *is* intended as an account of trust, and not just as an account of trusting people):

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. (2014: 11)

We will have to shift from 'someone' to 'something', at the very least.

I confess I remain uncertain as to how to think about the issue of commitment here. In the cases that I've described, it isn't clear to me that the trustors *need* to believe that the trusted entity *has a commitment* to doing something. For instance, consider once again the case of Mel and the body armour. Does Mel in this situation *need* to think that the body armour in fact *has a commitment* to protecting her? Hawley's analysis requires that Mel does need to think that.⁸ I'm honestly not at all sure how we decide that issue. Mel makes an error in trusting the armour; does Mel also make the error of imputing a *commitment* to the armour? It's far from obvious to me that this *need* be so. To be clear: I don't say that this poses an insuperable problem for Hawley. My more mundane point is simply that once we move to allowing that trustors can trust inanimate objects, I become uncertain as to the extent that we can sensibly ascribe to trustors the belief that the inanimate object *has a commitment* to doing something. *Perhaps* we can ascribe to trustors the belief that the inanimate object has a commitment in the right kind of way, but this seems to me an avenue for further research (see, also, discussion below). And that, in itself, is an interesting consequence to note and something that requires further exploration.

Of course, not all accounts of contractual trust are like Hawley's; not all trade on an explicit notion of commitment. Indeed, it is for this reason that I remarked just a moment ago that I would not be looking to provide a *general* fix. My aim is not to try and re-write competing accounts of trust, but to point out that some of them may be in need of such a re-writing. For instance, the account of trust given by Karen Jones (1996: 1) treats trust as:

⁸ See, also, fn 5.

[a]n attitude of optimism that the goodwill and competence of another will extend to cover the domain of our interaction with her, together with the expectation that the one trusted will be directly and more favorably moved by the thought that we are counting on her

Once again, and just to pick on one case again: in the body armour case there is simply no reason to think that Mel takes the body armour to be *favourably moved by the thought that Mel is counting on it*! Nonetheless, it does look as if Mel can (albeit inappropriately) trust the armour. Similarly, in the case of Wiley and the blackboard, though we *might* be prepared to commit to the idea that Wiley thinks that the blackboard has a commitment to telling him the truth (though I do find that notion *something* of a stretch), and so think this case compatible with Hawley's account, it seems far harder to swallow the idea that Wiley takes the blackboard to be favourably moved by the thought that he is counting on it. Even a credulous child should surely not be *required* to have that attitude in the situation in question.

There is a resurfacing of a potential moral, here, that I noted at the end of section 4: suppose Hawley's account fares a little better here than does Jones', with being able to account for trust in the case involving an inanimate object. If that's right, then it points nicely to the fact that cases involving inanimate objects may well be of bearing on how we try to clarify our overarching concept of trust, for certain cases involving inanimate objects might be taken to favour one account of trust over another. This, in turn, reinforces my point in section 2—that the possibility of being able to trust inanimate objects has ramifications for how we think about trust generally.

Third, there are *plausibly* some practical implications—implications that are broad in scope and worth pointing out here, and again these also bring front and centre the import of being careful to recognise that we can trust inanimate objects. As is obvious, we regularly interact with inanimate objects. Some of these inanimate objects are sold to us. Presumably, *some* inanimate objects that are trusted to carry out a particular task will sell better than those that are not. This is not speculation on my part. Potential customers already report distrust of driverless cars, for instance, and this is described in the media as a barrier to the widespread implementation of the technology.⁹ If am right about contractual trust, and if I am right about the role played by anthropomorphisation, there is therefore potentially some motivation for those designing and marketing this sub-class of products (those that will sell well/better if they are trusted to carry out some particular task) to do so in such a way as to deceive us into assigning agency to

⁹ See, e.g., <https://arstechnica.co.uk/cars/2015/07/self-driving-cars-are-coming-but-most-drivers-dont-want-them-yet/> For a very gentle introduction to research on the topic, see <http://theconversation.com/why-driverless-cars-still-need-driving-tests-62721>

Note, also, that if one reviews the stated objections from drivers it doesn't seem to be to the *reliability* of driverless cars that is at issue; the concern seems to centre on trust. This case is also a promising avenue for showing that we can richly trust inanimate objects. If we can, then, again, it seems that our account of trust ought to be modified to ensure that we can accommodate this possibility.

them, and thus to trust them to carry out the tasks that we *should* merely rely upon them to complete. If we couldn't trust inanimate objects, this would not be the case. If we couldn't trust inanimate objects then the only trust-related goal for the designers and marketers should be one of persuading the buying public that the product is *reliable*. (Of course, there will be other tasks for the marketers, too: persuading the public that the company can be richly trusted, for instance, may well be of commercial value.) Or, if those in the orthodoxy are simply not considering such cases, then it would be helpful in the extreme were they to do so.

One obvious (and current) domain where this point is especially pressing is the case of the driverless car (alluded to above and in fn 8). A recent MIT study found that of those who claimed that they would not ever purchase a driverless car, some 26% opted for the reason that 'they did not trust the technology' (Abraham, *et al.* 2017: 9). Note also that we have *some* reason to think that this notion of trust is the rich, morally loaded sense, since *a further 46%* responded that they would never purchase the technology, either because it wasn't safe or because it wouldn't work perfectly. In other words, this 46% *appear* to have been concerned that the product would not be *reliable*. Thus, we have at least some reason, albeit a defeasible reason, to think that respondents are going some way to distinguishing between (rich) trust and mere reliance. Notable, too, is that further research on this issue (e.g. Beggiato and Krems (2013)) explicitly distinguishes a notion of rich trust, contrasts it with mere reliance, and finds that one barrier to the acceptance of automated vehicles *is* a lack of trust. That being so, designers and manufacturers will want us to come to trust these vehicles in order to help potential customers move past this concern and consider making a purchase.

As an aside, it's worth noting that the account of trust that is given by Beggiato and Krems, though clearly gesturing in the direction of a morally loaded conception of trust, isn't satisfactory for our current purposes. As Beggiato and Krems (2013: 48) have it, borrowing from Lee and See (2004: 51), trust is: '...the attitude that an agent will help achieve an individual's goals in a situation characterised by uncertainty and vulnerability'. First, then, note that the account of trust just given is that of an attitude: it is of the form *x trusts y*. This is not contractual trust: it is not of the form *x trusts y to φ* . Second even as an analysis of attitudinal trust, this seems to me a flawed account. To see why, consider a case discussed by Hawley (2014a: 9). Imagine a scenario in which Alice brings with her an excess of food every day, leaving whatever she does not eat on a table in the staff cafeteria. She does so with sufficient reliability that, unbeknownst to her, I stop bringing lunch with me and stop bringing any money, too. I simply consume Alice's leftovers. I am, in short, relying on Alice's flawed judgements about quantity. Indeed, in such a situation I will likely have an attitude toward an individual, Alice, such that Alice will help me achieve my goals (getting fed) in a situation characterised by uncertainty and vulnerability (I do not *know* that she will bring an excess of food and if she does not then I do not know how to get access to other food and am now in a sense 'vulnerable' to being very hungry for the rest of the day). Nonetheless, it

seems quite wrong to say of such a case that I *trust* Alice.¹⁰ That being so, the account of (attitudinal) trust offered by Beggiato and Krems seems in need of revision.)

Even if one doesn't find the driverless care case especially persuasive there is also a rich literature that looks at the connection between automation on the one hand, and trust and reliance on the other. As Lee and See note in their excellent (2004: 76—my italics) overview:

Trust influences reliance on automation; however, it does not determine reliance... Because automation and computer technology are growing increasingly complex, the importance of affect and trust is likely to grow. Computer technology allows people to develop relationships and collaborate with little or no direct contact, but these new opportunities produce complex situations that require appropriate trust. As computer technology grows more pervasive, trust is also likely to become a critical factor in consumer products, such as home automation, personal robots, and automotive automation. Designing *trustable technology* may be a critical factor in the success of the next generation of automation and computer technology.

Therein, as is apparent, Lee and See differentiate trust from reliance, and note the importance (to consumer products) of developing technology—automated technology—that can be trusted (not just relied upon). Philosophers, too, must pay attention to cases of trust in automation, and look to provide an account of what that consists in.

In the face of all of this, it seems hard to maintain either that agents do not contractually trust inanimate objects or that this result is of little significance. The question of whether we can (e.g.) be brought to trust driverless cars for instance, and if we can, *how* we can, is one that will have a significant bearing on our financial futures, as well as the ways in which we organise the social and domestic elements of our lives.

6. Conclusions.

So, *can* you (richly) trust a ladder? Yes, but *likely* only if you treat it as an agent, and even then you *probably* shouldn't. Should you (richly) trust a driverless car? Again, probably not. But if you rely upon it and you anthropomorphise it with some sense of agency, then it seems to me you probably will.¹¹

References:

Abraham, H., Reimer, B., Seppelt, B., Fitzgerald, C., Mehler, B., and Coughlin, J.F. 2017. 'Consumer Interest in Automation: Preliminary Observations Exploring a Year's Change', White paper (2017-2), <http://agelab.mit.edu/sites/default/files/MIT%20-%20NEMPA%20White%20Paper%20FINAL.pdf>

¹⁰ See Hawley (2014: 9-10) for a defence of this claim.

¹¹ I'm very grateful to two referees for two rounds of comments on previous drafts that led to a significant re-shaping of the paper.

- Baier, A. 1986. 'Trust and Anti-Trust', *Ethics*, 96, 231-60
- Beggiato, M., and Krems, J.K, 2013. 'The evolution of mental model, trust and acceptance of adaptive cruise control in relation to initial information', *Transportation Research Part F*, 18, 47-57
- Epley, N., Waytz, A., and Cacioppo, J.T. 2007. 'On Seeing Human: a three-factor theory of anthropomorphism', *Psychology Review*, 114, 864-86
- Faulkner, P. 2014. 'The Moral Obligations of Trust', *Philosophical Explorations*, 17, 332-45
- Faulkner, P. 2015. 'The Attitude of Trust is Basic', *Analysis*, 424-9
- Hawley, K. 2014a. 'Trust, Distrust and Commitment', *Nous*, 48, 1-20
- Hawley, K. 2014b. 'Partiality and Prejudice in Trusting', *Synthese* 191, 2029-2045
- Hieronymi, P. 2008. 'The reasons of trust', *Australasian Journal of Philosophy*, 86, 213-36
- 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.
- Lee, J. D., and See, K. A. 2004. 'Trust in automation: Designing for appropriate reliance', *Human Factors*, 46, 50-80
- McLeod, C. 2015. 'Trust', *The Stanford Encyclopedia of Philosophy* (Fall 2015 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/fall2015/entries/trust/>>
- Owens, D 'Trusting a promise and other things', in P. Faulkner and T. Simpson. Eds. *The Philosophy of Trust*, Oxford: OUP, 214-229
- Pettit, P. 2004. 'Trust, Reliance and the Internet', *Analyse & Kritik*, 26, 108-21
- Simpson, T. 2012. 'What is trust?' *Pacific Philosophical Quarterly*, 93, 550-69