

# Trust as an Unquestioning Attitude

C. Thi Nguyen

(This is a pre-print draft. Forthcoming in *Oxford Studies in Epistemology*)

*Most theories of trust presume that trust is a conscious attitude that can be directed at only other agents. I sketch a different form of trust: the unquestioning attitude. What it is to trust, in this sense, is not simply to rely on something, but to rely on it unquestioningly. It is to rely on a resource while suspending deliberation over its reliability. To trust, then, is to set up open pipelines between yourself and parts of the external world — to permit external resources to have a similar relationship to one as one’s internal cognitive faculties. This creates efficiency, but at the price of exquisite vulnerability. We must trust in this way because we are cognitively limited beings in a cognitively overwhelming world. Crucially, we can hold the unquestioning attitude towards objects. When I trust my climbing rope, I climb while putting questions of its reliability out of mind. Many people now trust, in this sense, complex technologies such as search algorithms and online calendars. But, one might worry, how could one ever hold such a normatively loaded attitude as trust towards mere objects? How could it ever make sense to feel betrayed by an object? Such betrayal is grounded, not in considerations of inter-agential cooperation, but in considerations of functional integration. Trust is our engine for expanding and outsourcing our agency — for binding external processes into our practical selves. Thus, we can be betrayed by our smartphones in the same way that we can be betrayed by our memory. When we trust, we try to make something a part of our agency, and we are betrayed when our part lets us down. This suggests a new form of gullibility: agential gullibility, which occurs when agents too hastily and carelessly integrate external resources into their own agency.*

## 1. Introduction

Most accounts of trust take it to be a conscious attitude directed towards other agents. According to some accounts, to trust somebody is to think that they have goodwill towards you, or that they will be responsive to your needs.

According to other accounts, trust is an attitude we adopt for various social reasons, which encourages us to rely on others and believe what they say. All these accounts share two central features. First, trust is supposed to be a clear and present rational force — an active participant in an ongoing deliberative process. Second, trust can only ever be directed at other agents.

In this paper, I would like to explore a very different alternative: that there is a form of trust which involves suspending the deliberative process. To trust something, in this sense, is to put it outside the space of evaluation and deliberation — to rely on it without pausing to think about whether it will actually come through for you. To trust an informational source is to accept what that source delivers without pausing to worry or evaluate that source's trustworthiness. To trust, in short, is to adopt an unquestioning attitude. Which is not to say that one can't question this sort of trust, or reason about whether one ought to trust something. Trust can certainly arise out of deliberation and it can certainly be called into question. But it is to say that when one has come to trust, one has adopted an unquestioning attitude. And limited beings like us must often take up such unquestioning attitudes as part of a reasonable strategy for coping with the overwhelming cognitive onslaught of the world.

Crucially, we can hold this unquestioning attitude towards non-agents: simple objects, parts of my body, and features of the natural world. I can trust my legs and I can trust the ground. To understand why this is a significant departure, we need to look at the history of the philosophical work on trust. That literature springs from a couple of inquiries. First, philosophers have been interested in the morality of trust and how it plays out in various efforts of cooperation and social relationships (Baier 1986; 1992; Baker 1987; Holton 1994; Jones 1996; 2012; McLeod 2002; O'Neill 2002a; 2002b). This conversation tends to focus on how trust works in distinctively moral, social, and political settings. Second, philosophers have been concerned with the epistemology of trust and how we might acquire knowledge through testimony (Hardwig, 1991; Hinchman 2005; Faulkner 2007a; 2011; Hieronymi 2008; Lackey 2008; Nickel 2012; Keren 2014). Both these conversations share a central presumption: that trust is agent-directed. That is, trust is taken to be an attitude of one agent *directed toward some other agent*.

We can find this presumption clearly articulated in the opening moments of the modern conversation on trust. Annette Baier's pioneering work explicitly sets the focus on agent-directed attitudes. Importantly, she says, our colloquial use of "trust" blurs together two very distinct concepts. She proposes

that we adopt a terminological refinement. First, there is the attitude of *mere reliance*, in which we simply depend on something. Second, there is the attitude of *trust*, in which we depend in some more normatively loaded manner. Suppose I notice that you pass by my door every day at noon and I start using your passage to time my lunch-break. In this case, I have merely come to rely on you. If you suddenly stopped passing by at noon every day, I might be disappointed that I had lost this convenient signal, but I could make no reasonable criticism of you. But if you failed to knock after having promised to do so, I would feel, not only disappointed, but *betrayed*. I had trusted you and you let me down. Our relationship towards objects, says Baier, can be one of, at most, reliance. It is only other people that we might come to trust. And the possibility of betrayal is a telling sign of the presence of full-blooded trust — and we do not feel betrayed by objects (Baier 1986).

The ensuing conversation has largely followed Baier's basic framework. Philosophers have accepted her claim that trust is essentially agent-directed. And they have followed Baier in treating the possibility of betrayal as the sign that trust is present. Thus, they have studied betrayal in order to understanding the content of trust. How might the reaction of betrayal be appropriate? In what might it be normatively grounded?

Baier's own account is that trust involves ascribing goodwill to the trusted, and that our sense of betrayal comes from the discovery that there is no such goodwill after all. Baier's account has seen some notable counterexamples — such as Onora O'Neill's observation that you may trust a doctor simply for their professionalism, with no expectation of goodwill whatsoever (O'Neill 2002, 14). Most theorists have since abandoned Baier's particular emphasis on goodwill, but many of the proposed modifications still retain the general shape of Baier's proposal. Many of these newer theories replace Baier's focus on goodwill with a focus on *responsiveness*. According to responsiveness theories of trust, to trust somebody is to think that they will respond to your trust positively — to think that the fact that you trust them will give them a reason to fulfill that trust. As Karen Jones puts it, a trustworthy person "takes the fact that they are counted on to be a reason for acting as counted on" (Jones 2012, 66). Similarly, Paul Faulkner suggests that when one person trusts another, the truster knowingly depends on the trusted to do something, and expects the trusted's knowledge of this dependence to motivate them to do it (Faulkner 2007b, 313). Betrayal, then, is grounded in the betrayer's failure to be properly

responsive. Katherine Hawley, on the other hand, rejects the details of the responsiveness account, but still analyzes trust in agent-directed terms. For Hawley, to trust somebody is to take them to have made a commitment to do something and to rely on them to fulfill that commitment (Hawley 2014). Hawley's account grounds the sense of betrayal in the trusted person's failure to live up to their commitments.

Note that Baier's account, Hawley's account, and the responsiveness account all share the presumption of agent-directedness. To put it more precisely, they all presume that the truster must *ascribe some complete agential state* to the trusted — be it a belief, motivation, disposition, or commitment. It follows, then, that we can only trust the stand-alone bearers of agential states: individuals, certainly, and also group agents, like nations and corporations (Hawley 2017). Perhaps we can also trust certain complex technological artifacts, like Google Search, precisely when we can attribute some form of agency to them. But we cannot trust or distrust dumb objects with no agencies of their own.

My proposed account rejects the presumption of agent-directedness. I will describe a form of trust that need not ascribe any complete agential states to its target. We can take the unquestioning attitude towards a wide variety of objects and artifacts. To trust, in this sense, is to have stepped away from the deliberative process. It is a way of settling one's mind about something. To trust is to lower the barrier of monitoring, challenging, checking, and questioning — to let something inside, and permit it to play an immediate role in one's cognition and activity. It is, in a sense, to give an external resource a direct line into one's reasoning and agency. Such trust is our mechanism for attempting to *integrate* other people and objects into our own functioning. This form of trust is still deeply bound up with agency, but it need not be directed only towards complete, external agents. We can be betrayed by objects, then, not because some distinct external agent has failed us, but because they have failed us in our attempt to integrate them into our own agency. Our response of betrayal towards those objects, then, is a close cousin of the betrayal we feel towards our own recalcitrant, failing parts.

This form of trust functions as an engine for agential outsourcing. When we trust things, we are granting them a degree of cognitive intimacy which approaches that of our own internal faculties — with all the access privileges and unquestioned acceptance that status accords. When I wholeheartedly

trust my online calendaring system, I usually act on its notifications unquestioningly — just as I usually do with my own biological memory. This allows for enormous efficiency, but at the price of exquisite vulnerability. This also suggests that there is a distinctive form of gullibility: agential gullibility, in which we too readily bolt external processes onto our own agency. Many of our relationships to emerging technologies — search algorithms, smartphones, social media networks — are marked by such agential gullibility.

I do not suggest that this account of the unquestioning attitude be taken to replace, or subsume, the traditional agential accounts of trust. Rather, I will suggest that the unquestioning attitude is one form of trust, and that the more familiar agent-directed accounts of trust explore another form of trust. And these different forms of trust can interact. For example, I might take an unquestioning attitude towards somebody precisely because I take them to have the right sort of goodwill or responsiveness. But these forms of trust also come apart. I trust the ground in the unquestioning attitude sense, but not any agent-directed sense. Finally, I will suggest that there is a reason that we group these various attitudes together under the umbrella of “trust”: they are all ways we have of expanding our agency by integrating in bits of the external world. Responsive cooperation and the unquestioning attitude are two tools for trying to integrate external resources into one’s agency.

## **2. Trusting the ground**

The trust literature shares a common founding presumption: that trust is a relationship we could only have towards other independent, complete agents. Talk of trust towards non-agential objects has been easily dismissed. After all, the colloquial language here is fuzzy. Everyday talk of trust in objects can simply be interpreted, in our newly technical language, as concerning mere reliance. After all, how could you ever be betrayed by an object? And isn’t the response of betrayal only appropriate when directed towards agents?

But if we look beyond the philosophical discussion of trust and morality — if we look to literature and to life — it’s easy to find talk of trust in objects. I will, for the rest of this paper, use “trust” to refer to the full-blooded, normatively loaded sense, and use “objects” to refer to non-agential objects. And I

will take onboard Baier's diagnostic. A sign that we aren't merely relying on objects, but actually trusting them, is the presence of that characteristic, negative, normatively loaded reaction to trust's breach. We know that we are in the presence of trust when we are willing to speak of betrayal.

So: do we actually trust objects, in the full-blooded and normatively loaded sense? Climbers speak of trusting the rope; they react with something far sharper than mere disappointment when a rope goes bad. And this sort of talk isn't just limited to human artifacts. We will say that we feel betrayed when the ladder gives way beneath us, but also by the collapse of that solid-seeming tree which we were climbing. We speak of trusting the ground — and of being betrayed by it when good footing turns unexpectedly bad. And we speak of the shock of discovering the untrustworthiness our own faculties and parts — of being betrayed by the shakiness of our hands or by our faltering memory. Superficially, these sorts of examples seem to weigh against the insistence that trust always be directed at agents.

I will begin in the familiar mode of conceptual analysis, but that is only a starting point. My aim here is to key in on a real-world phenomenon, using our language and concepts as a pointer. And I have a larger purpose in this investigation. I think there is a distinctive form of relationship we can have with objects, which goes beyond mere reliance, and which is best described as a form of trust. Contemporary life is significantly marked by trust in technological artifacts and technologically-mediated social environments: Google's search algorithms, smartphones, the ranking algorithm behind Facebook and Twitter, the emergent networks of interconnection on social media<sup>1</sup>. Our relationships with these objects, I suggest, is far more potent than mere reliance. One might respond that this is not really trust in objects, but trust in the designers behind those objects. Sanford Goldberg, for example, suggests that we can have normatively loaded relationships with designed artifacts, since we are willing to hold those designers to account when their artifacts fail us (Goldberg 2017). According to Goldberg, then, the sense of betrayal I feel when my iPhone fails me is properly directed at the corporation and manufacturers behind it.

---

<sup>1</sup> For starters, see (Pariser 2011; Miller and Record, 2013; Frost-Arnold 2014; Frost-Arnold 2016; Rini 2017; Nguyen 2018b). Many of these conceptualize trust as only directed towards other agents, but mediated by these technologies. I suspect, however, that that is partly due to the present lack of available theoretical resources for making sense of trust in objects.

But I think that there is a distinctive sense in which we can trust objects themselves — including non-designed objects, like the ground. And even with designed objects, I think our trust often cannot be wholly cashed out in terms of trust in the people and institutions who designed those objects. First, in many cases, the artifacts we trust have run beyond their creators' abilities to understand or control. One of the pressing issues in the ethics of technology involves thinking about machine learning algorithms, which have been built using evolutionary techniques, and whose innards and proceedings aren't understood by those who have built them (Resch and Kaminski 2019; Carabantes 2019). Second, key features of the network architecture of online social structures, such as social media networks, have evolved beyond the intentional control of the institutions that have made them. Third, in many cases the question of our trust in a particular object is something distinct from our trust in its manufacturers. Climbers, for example, need to decide whether to trust an old rope. What matters is not the manufacturer's goodwill or intent in manufacturing ropes. The question is whether this particular rope should still be trusted after its own particular history of abuse.

So let's start with trust in some obviously non-agential and undesigned parts of the world. When I walk, I usually trust the ground. This means more than simply relying on the ground. When I trust the ground, I walk without bothering to consider whether it will be steady beneath my feet. I don't evaluate the ground or ponder its supportiveness. I simply walk on it while thinking of other things. When I distrust the ground, I am constantly questioning its reliability. I worry about it; I test it. If I am walking across a muddy field, riddled with gopher holes, I'll examine the ground carefully before each step. And even when I do force myself to rely on that muddy ground — when I commit my weight to it — I still don't trust it. My reliance is tentative and demands constant reassurance. When, on the other hand, I come to fully trust the ground, I stop worrying about it. And the difference between mere reliance and unquestioning trust tracks our different negative reactions, too. When I hesitantly rely on the ground, I am only glumly disappointed when it gives way. But when it is when I walk without thinking about it — when the ground has become automatically and unthinkingly integrated into my background physical processes — that I react with shock and betrayal when it gives way.

I suggest that trust, here, involves taking on an unquestioning attitude. To trust something in this way is to rely on it while putting its reliability out of

mind. When we don't trust, we question. Sometimes the answers to our questions might be positive; sometimes they might be negative. Sometimes we may decide to rely on something after we've questioned it thoroughly. But the lack of trust is shown in the very process of active investigation itself. It is only when we have settled our mind and stopped actively questioning something that we truly trust it, in this sense. This does not mean that when you trust something, you never question it at all. To trust something is to have a general disposition not to question it. That disposition can be disrupted or overwhelmed for the moment, but we are still trusting something so long as we are generally disposed to not question it. We only lose trust when we lose the disposition itself.

I have found that philosophers who work on trust and testimony often find this use of "trust" is bizarre and unintuitive — especially locutions like "trusting the ground" and feeling "betrayed by the ground." But it seems to me that, in fact, these expressions are entirely natural and comprehensible, and it is only excess immersion in modern, narrowed philosophical theories of trust that renders them odd to the ear.

We can find talk of trust in and betrayal by the ground throughout ordinary speech. Consider this advice from a manual on trail-running.

So pay attention... Don't trust wooden structures. Stiles, bridges, fences, tiger traps, path edges: no matter how inviting they look, unless you have thoroughly tested them before, DON'T TRUST THEM.... Very few running mishaps result in such painful or long-lasting injuries as overconfident approaches to wooden structures. Just slow right down for a few strides and, if possible, find something to hold on to as you go... Oh yes, and don't trust the ground on either side of wooden structures either — in case you were thinking of leaping over one... The ground on either side will be much trodden and thus probably churned up, slippery, and generally untrustworthy. Just relax, take that extra second and speed up again when you're on the other side. (Askwith 2015, 150)

Notice that the runner here is not being told to avoid relying on the ground. Sometimes they must rely, because there is no other place to step and nothing else to hold on to. The runner is being asked to suspend their unthinkingness, to pay attention, to be careful. They are not being told to avoid any form of reliance; they are being told to suspend their trust. They are being told to rely on the ground, but in a mode of interrogating, suspicious awareness.

The presence of trust and betrayal are quite clear in certain traumatic experiences of hostile environments. From a sociological investigation into the



experience of war:

The veteran also suffers from a problem of trust, a building block on which all of social life is erected. The everyday, taken-for-granted reality of civilian life ignores much; civility assumes the nonlethal intentions of others. In war, however, all such assumptions evaporate: one cannot trust the ground one walks on, the air one breathes, nor can one expect with full assuredness that tomorrow will come again. (Kearl 1989, 353)

The best explanation here is not that soldiers in war have suspended their reliance — after all, one cannot but rely on the air and the ground. What changes is their attitude towards that reliance. They become suspicious, unable to rest easy on the assurance that the ground and air will continue to support them.

Tellingly, the language of trust and betrayal often crops up in stories about the emotional aftermath of earthquakes. Douglas Kahn writes:

I will never forget being in an earthquake near Seattle in which the ground itself became acoustic, with swelling waves traveling down through the road making houses I knew well bob up and down like ships on the sea. “A moment destroys the illusion of a whole life,” writes Alexander von Humboldt in *Cosmos*. “Our deceptive faith run the repose of nature vanishes, and we feel transported as it were into a realm of unknown destructive forces. Every sound — the faintest motion in the air — arrests our attention, and we no longer trust the ground on which we stand.” (Kahn 2013, 133)

And here is Betty Berzon’s earthquake story:

The house rocked and rolled, the glassware fell out of the cabinets, the pictures slid off the walls, the furniture skidded across the floor, and light fixtures came crashing down from the ceiling... I was frozen with fright and sure the house would topple over and end up in the street below. I was certainly going to die... The 6.6 earthquake and the aftershocks continued into the next day, but the house didn’t fall down. There is something about being betrayed by the ground underneath you that feels like the ultimate treachery. It took weeks to regain my equilibrium. (Berzon 166, 2002)

These samplings make clear that we can trust objects in the more substantive sense. And the loss of trust can hit us in a similarly sharp register, whether it be in other people, the ground, or the air. These narratives make clear that this loss of trust must be something beyond the loss of mere reliance. For after an earthquake, we must still rely on the ground. After war, we must still rely

on the air. But suspicion intrudes upon us, and we can no longer take their reliability for granted. Our mind is profoundly unsettled. (The fact that many philosophers find it odd to speak of being betrayed by their environment is perhaps best explained by the fact that most philosophers have lead, by and large, pretty cushy lives.)

Of course, one might continue to insist that these uses of “trust” and “betrayal” are merely metaphorical. They do not sound so to my ear. Saying that one felt betrayed by the ground after an earthquake, or by one’s failing memory, strike me as paradigmatic invocations of the concept. But I do not think that we can settle the matter here just by comparing the intuitive rings of various locutions in our various ears. More importantly, even if this use is merely metaphorical, there is a reason why we reach for this particular metaphor — a reason why we reach for the terms “trust” and “betrayal” when we find ourselves profoundly perturbed by an earthquake. What’s most important here is to understand the nature of the heightened form of relationship we sometimes hold towards objects, which goes beyond mere reliance, and whose breach could reasonably ground a sharply negative, normatively loaded response — a response which is entirely natural to describe as betrayal.

### **3. Trust in the background**

Before we attempt to give a detailed accounting of the unquestioning attitude is, let’s take a bit of time to get clearer on when and where it occurs. Crucially, I think it is easy to miss the unquestioning attitude precisely because it is so pervasive. The unquestioning attitude blends into the background of our cognition. And this is exactly what we should expect, since the function of the unquestioning attitude is precisely to put certain worries out of mind. To settle your mind about something is, in part, to shove it into the cognitive background.

First, we can take the unquestioning attitude towards objects, but we can also take it towards other agents. I trust my doctor about medical advice insofar as I take their medical suggestions as immediate reasons to act, without pausing to worry about what ulterior motives they might have for selling me this drug. I trust the newspaper when I simply accept its pronouncements

without worrying about whether its staff might be financially biased or lazy. A soldier trusts their squad mates when they plunge ahead, believing without question that their squad mates have their back.

I was once involved in a car accident; another driver lost control of their car and swerved across a narrow country highway, hitting me head-on. Afterwards, I lost my trust in other drivers. I hadn't realized, in fact, how much I had been trusting other drivers until that trust suddenly evaporated. What had changed? It wasn't my attributions of goodwill or responsiveness to the other drivers. If you had asked me before and after the accident, I would have made the same estimate of how much goodwill and responsiveness I could reasonably expect from other drivers. What's more, I had relied on other drivers before the accident, and I still relied on other drivers after the accident. What changed in the accident was my ability to sink into the unquestioning state. I was stuck in an endlessly suspicious mood. What had changed — what had evaporated — was my easy, settled, unquestioning state of mind.

When we say we trust agents, then, often much of that trust is actually best cashed out in terms of the unquestioning attitude sense of trust, instead of strictly in terms of agent-directed accounts. But this is often misunderstood, because the conversation about trust has sometimes focused on the wrong sorts of cases. We often focus our analysis on those cases where trust comes to mind. But that focus is misleading. So much of our trust occurs in the cognitive background. Often, it is only when our trust is threatened that we suddenly realize how much we have been trusting all along — as with my car accident. As Baier puts it, we inhabit trust like we inhabit the air, and we only notice it when it has departed (Baier 1986, 99; Jones 2004). Similarly, Thomas Simpson suggests that all our varied talk of trust descends from a simple primitive phenomenon: our need to cooperate. But this ur-trust is such a background feature of our lives that we barely think or talk about it. Trust only comes to mind once it has been threatened. Thus, the fact that we are actively thinking and talking about our trust actually indicates that we are likely at the peripheries of the core phenomena. In fact, talk of trust has many apparently different faces, which reflect all the different ways ur-trust can be threatened (Simpson 2012, 560-1). The implications of this analysis are quite striking: if we only analyze those incidents where issues of trust have come to the forefront of our consciousness and conversation, we will actually miss out on the heart of the matter.

In this light, let's reconsider some of the standard examples that have

fueled the literature on trust. Take Richard Holton's central case, from which he builds much of his account of trust: the trust fall, an exercise beloved of acting classes and management training courses, where we make ourselves fall into the arms of others. When we take a trust fall, says Holton, we decide to trust. We will ourselves to trust. Cases like this suggest to Holton that trust can be voluntary, and that it can outrun the evidence. We do not know if people will catch us, but we decide to trust them in order to find out (Holton 1994). If we focused on these sorts of cases, we might end up thinking that trust is not so unthinking, after all. The process of questioning and weighing considerations seem quite prominent with the trust fall. The novice climber, too, typically engages in a tentative and deliberate process as they learn to trust the rope. If we took trust falls to be paradigmatic instances of trust, then it would be serious mark against my account — especially when we think of trust falls onto the climbing rope. After all, here is a moment of trust in an object full of consciousness, indecisiveness, and questioning.

But notice that these staged trust falls are actually at the periphery of trust. Trust falls are done between people that do not trust, as an exercise in learning how to trust.<sup>2</sup> Similarly, the novice climber who nervously talks themselves into taking practice fall after practice fall onto the rope is not yet fully trusting; they are at an early stage on the path to trust. The paradigm of trust in being caught actually looks quite different. Consider the experienced rock climber's attitude toward their rope and their gear. A novice rock climber tests the rope gingerly, occasionally weighting it, telling themselves over and over again to trust it. This is part of the process of coming to trust — but it is only the beginning of the process. While they are engaged in this process of self-negotiation and self-reassurance, we would say that they do not yet fully trust the rope. It is the experienced rock climber who truly trusts their rope. Their trust is reflected in the fact that concerns about the rope's reliability occupy no mental space for them at all. And that trust lets them focus all their mental efforts on the climb itself (Ilgner 2006).

---

<sup>2</sup> Pamela Hieronymi offers a similar explanation of Holton's discussion: that what we are doing here is not full-blooded trust, but merely entrusting — acting *as if* we trusted, as part of the process of building trust (Hieronymi 2008).

#### 4. Trust and resolve

I now owe an account of how one might reasonably trust objects, and reasonably expect something of those objects, in a way that might justify a sharply negative, normatively loaded response to the breach of trust. But, at the same time, that account should not require that trust involve attributing any complete agential states to the trusted. I will now attempt to give such an account.

I take inspiration here from a very different sector of Holton's philosophical work: his analysis of weakness and strength of will. Let's examine his account in some depth. To exercise willpower, says Holton, is to close yourself to a certain kind of reconsideration. It is to settle your mind in a certain direction. Holton is building here on Michael Bratman's account of intentions. Suppose I form an intention. Later on, I can act directly from that previously formed intention. I don't re-deliberate, treating my past self's decision as merely one input among many. In other words, I don't treat my remembered intention as the mere issuances of some distinct past self. In ordinary circumstances, I simply act on my past intention. What it is to form an intention is to make up one's mind in a way that extends to one's future self. It is to have decided for one's future self.

In order to perform that role, intentions must have a certain stability. They must exhibit *cognitive inertia*. This doesn't mean that intentions can't ever be re-considered. Rather, the standards for re-consideration are much higher. For example: I might take all sorts of reasons into account when deciding where to go to dinner tonight — like what my current cravings are, or the exact state of my bank account. But once I form the intention to go to Roscoe's House of Chicken and Waffles, I close myself to reconsidering that intention from minor fluctuations in the usual run of considerations. I will only re-open deliberation in the face of some drastic change in the circumstances, like Roscoe's catching fire, or a violent case of the stomach flu.

What's more, says Holton, we sometimes need something stronger than an intention: we need a *resolution*. We make resolutions when we need to steel our ourselves against future temptations. A resolution, says Holton, is a pair of intentions: it is an intention to do something, and then a second-order intention not to let that first-order intention be deflected (Holton 2009, 11). In other words, willpower includes the power to *actively refuse to reconsider* intentions. It is, we might say, willful inertia. And we breach a resolution when we open it up to the possibility of revision. That refusal to re-consider is required

for resolutions to play their particular role. As Holton puts it:

...Much of the point of a resolution, as with any other intention, is that it is a fixed point around which other actions — one's own and those of others — can be coordinated. To reconsider an intention is exactly to remove that status from it. (121-2)

Why do we need to close ourselves off in this way — both acquiring cognitive inertia itself, and then sometimes doubling down with active and willful support for that inertia? First, says Bratman, we need to fix intentions in order to make plans, both with ourselves and with others.<sup>3</sup> And resolutions firm up our plans, solidifying them in the face of temptation. But behind this thought lies deeper considerations about cognitive finitude. I only have so many cognitive resources to go around, and this is a way to conserve them. The cognitive inertia of intentions is central to their role in cognitive resource management, and we sometimes need to cement that cognitive inertia through an active effort of will. Limited beings need to settle their minds about about some investigations in order to free up their cognitive resources for other investigations, and they need to work to protect that settled state of mind.

I suggest that the unquestioning attitude of trust plays a roughly analogous role to that of resolutions in settling the mind. I do not mean to claim here that trust is some sort of resolution. I mean, instead, to indicate a functional similarity between trusting and making resolutions. Trust is a strategy to cope with our cognitive finitude and manage our limited cognitive resources — to steel ourselves across time to the cognitive disruption of dealing with new evidence by (defeasibly) closing our minds against re-consideration. Trust is a way of establishing fixed points in our deliberation. And trust is distinct from resolutions, because it is a way of establishing *external* fixed points — resources that we will always accept without question, resources that we will rely on without thought.

Here is the *unquestioning attitude* account of trust:

To trust X to P is to have an attitude of not questioning that X will P.

We can also offer a specific instantiation of this unquestioning trust, for trusting informational sources:

---

<sup>3</sup> (Holton 2009, 2-4; Bratman 1987). My terminology and framing here draws on Holton's presentation of Bratman.

To trust X as an informational source in domain Z is to have an attitude of not questioning X's deliverances that Z.<sup>4</sup>

Let's take a look under the hood. I intend the "attitude of not questioning" to have a similar two-tiered structure as Holton's account of resolutions. To trust X to P is to have a first-order disposition to immediately accept that X will P, and a second-order disposition to deflect questioning about the first-order disposition.

First, note that having an unquestioning attitude that X will P does not involve a disposition to come to a particular conclusion from deliberation about whether X will P, or to discount certain forms of evidence while deliberating about whether X will P. It is a disposition against deliberating, in the first place, whether X will P.<sup>5</sup> Second, an unquestioning attitude is defeasible, but the reasons needed for defeating the unquestioning attitude that X will P need to be significantly stronger than merely being reasons that bear against belief that X will P. The unquestioning attitude towards X's doing P is thus *resistant*: it maintains itself against some classes of considerations that would normally weigh against my believing that X will P. Third, the account is intended to indicate a spectrum concept. One can trust with varying degrees of unreservedness, since one can hold the dispositions with varying degrees of force. Finally, in almost all cases, the scope of trust in X will be restricted to some particular functions of X. However, in colloquial usage, the scope of trust is often implicit and understood from context — usually in light of some highlighted role or function. When I say I trust my doctor, I can be understood to mean that I trust my doctor to perform their medical duties, and not that I trust them to do modal logic or play jazz.

---

<sup>4</sup> "Deliverances" here is meant to be a general term for transmitting propositional content. The deliverances of other people are usually what we call "testimony". But other cases of trusting informational sources' deliverances include: trusting my watch to tell the time; trusting my eyes to deliver accurate visual information; trusting my calendaring system to auto-synch between my phone, laptop, and tablet and report to me the events that I have entered into it; and trusting Google Search to deliver search results organized by relevance

<sup>5</sup> My account here shares certain thematic similarities to Lara Buchak's account of faith as steadfastness in the face of counter-evidence (Buchak 2017), but the key difference is here. Buchak's analysis concerns cases in which it is rational for me to commit myself in a way so as to ignore counter-evidence during deliberation; my analysis concerns when we suspend deliberation altogether.

Next, notice that the account does not say that to trust X to P is to not question X in any way. It says, rather, that to trust X is *not to question that X will P*. That is, when I trust X to P, I don't question X's efficacy in particular instances of doing P. By the account I've given, it is possible to trust that X will P while asking questions about X in general — so long as we accept X's particular deliverances and affordances in regards to P. It's possible to trust something and, at the same time, to ask general questions about that thing's reliable functioning, *so long as we don't question particular instances of that functioning*. For example: suppose Esi is a memory researcher. Her research focuses on the fallibility of memory; she frequently asks questions about human memory in general, and is willing to extend those theoretical worries to her own memory. So long as she doesn't question her memories of what she had for breakfast or what time her doctor's appointment is, then she still can be said to trust her memory to deliver particular contents. Of course, questioning her memory in general may lead to questioning particular contents presented by her memory — but the two levels of questioning are conceptually distinct. Similarly, academic philosophers can ask as many questions as they like about the justifiability of accepting the deliverances of their senses, but they can still be said to trust their senses to deliver accurate presentations of the world so long as they unquestioningly accept particular sensory presentations.

Importantly, trust is an unquestioning *attitude* — understood as a two-tiered set of dispositions — and not a total cessation of questioning. Those dispositions can be defeated in particular instances and yet still remain as dispositions. I may trust my friend about all movie trivia, and then come across good reason to think they have probably made a particular mistake about the casting history of Ozu's *Late Spring*. That doesn't destroy my trust in my friend's encyclopedic movie knowledge, because it doesn't budge my overall disposition to accept their claims about cinema unquestioningly. The inertia of trust can survive the occasional disturbance. I only lose trust when I lose that inertia — when I lose the dispositions against questioning and let any sort of considerations trigger redeliberation about their claims.

Furthermore, the account specifies that trust is an unquestioning attitude, and not that it has gone unquestioned. The account is entirely compatible with my having questioned, sought justifications for, and deciding to trust, prior to my actually trusting. Don't confuse the issue of what it is to trust with the issue of the basis on which one has come to trust. I can decide to trust this rope to hold my weight because it has held it so many times in the past. But what it is



to decide to trust the rope is to decide, henceforth, to stop questioning it. Notice, though, that the unquestioning attitude account is also compatible with forms of trust which have never been questioned. Naive trust in authority and in the physical environment often has such a character. As Baier says, any adequate account of trust has to take into account the trust of children for their caregivers (Baier 1986, 240-6). And I take it that when my toddler eats the food I give him, he may have no reasons for his unquestioning acceptance of what I hand him. He has always trusted me and he plunges the food into his mouth without a moment's hesitation. A point in favor of the unquestioning attitude account is how well it models such naive, unconsidered trust. My toddler's trust in me is one of the paradigmatic instances of trust — but it is not best explained in terms of his attributing some commitment or benevolence to me. It is something more unthinking than that. His trust, I suggest, is constituted by his unquestioning acceptance of my food offerings.

The unquestioning attitude account also explains why trust is often recalcitrant. I often find it hard to trust, even if all my reasons indicate that I should.<sup>6</sup> I may have every reason to think my belayer and climbing rope trustworthy — but, still, I might find myself unable to trust them. I have come to trust only when I actually have made the transition to the unquestioning attitude. Consider a well-documented exercise for learning to trust: Arno Ilgner's technique for training climbers to trust their gear and their belayer. A beginning climber has likely done their research and learned that modern ropes simply do not break in standard circumstances, and that modern climbing gear is at least as trustworthy as, say, a car. They have likely chosen a trustworthy belayer. But many beginning climbers find that they cannot banish worries about the rope and the belayer from their mind, which greatly limits their ability to climb fearlessly and efficiently. Ilgner's solution to this mental difficulty is simply practice and repetition with falls. The climber must climb a little bit above their last anchor point, and jump off so many times that they simply become bored. Then they must climb a little higher, and jump off, over and over again. And, over time, the evidence they have that the rope will not break becomes something else: a confidence so complete that it recedes into the background, and

---

<sup>6</sup> The observation of recalcitrance comes from (Baker 1987). Karen Jones has offered a different account of the recalcitrance of trust. She suggests that trust is an affective attitude of optimism about the trusted's goodwill and competence (Jones 1996) My reasons for thinking trust is an attitude, rather than a set of reasons, borrows from Jones' analysis. Obviously, we cannot generalize her account's particular references to goodwill and competence to understand trust in objects.

then the climber can focus entirely on the climb itself, without having to worry about their gear or having to rehearse to themselves all the reasons they have to think it trustworthy. The transition from mere reliance to trust here is exactly the transition between having the reasons to trust and having the further attitude of unquestioning acceptance (Ilgner 2006).

## **5. The integrative stance**

Does the unquestioning acceptance account meet our desiderata for an account of trust? Let's start with Hawley's demand for a tripartite account of trust. Hawley notes that reliance has only two states — we either rely on something, or we don't. But with trust, there are three distinctive states. We can actively trust or actively distrust. But we can also occupy a third, neutral state — what Hawley calls non-trust. Actively distrusting somebody is a very different state from merely not trusting them. Any theory of trust needs to account for all three of these possible registers.

The unquestioning acceptance account meets Hawley's demand quite tidily. To trust is to have an unquestioning attitude. To distrust is to have an actively questioning attitude. And to non-trust is to have a neutral attitude, which is entirely open and unresistant to both questioning and non-questioning, as the situation suggests. I trust the ground when I don't think about it, and when that unthinkingness has been adopted as a disposition with some weight and resistance behind it. One uneven bit of sidewalk doesn't, by itself, disrupt my trust in the ground. I come to distrust the ground only when I acquire a disposition to actively worry about and question each step. And I non-trust the ground when I maintain no disposition in particular, but let whatever considerations arise push me in either direction. Most of the time, my trust settings are something like this: I trust the sidewalk and the highway; I distrust swampy and icy ground; and I non-trust natural grassy plains and the average backcountry hiking paths.

Now for the main event: we need to explain how this form of trust is something above and beyond mere reliance. And we need to do so in a way that could help explain why we might feel betrayal. On a first pass, our language is full of talk of betrayal by non-agents: of being betrayed by our body when it fails to obey our will, of being betrayed by our memory when it starts to go; of

rage and anger at our malfunctioning computers and recalcitrant devices. But is this merely a sloppy or metaphorical use of “betrayal”? To claim that it is a full-throated use of “betrayal”, we would need to explain what might ground and make appropriate the reaction of betrayal. But how could it ever be reasonable to have normatively charged expectations of objects? How could objects ever be the appropriate subjects of criticism?

Thinking about how we can be betrayed by our parts will shed some light on the matter. Betrayal by our body and our mental faculties seems to be a paradigmatic case of non-agential betrayal. We would be profoundly shocked if one of our limbs were to suddenly resist our control, refusing to move in accordance with our intent, or lunging about of its own accord. We rage and blame when our memory starts to go. What could justify the sharpness of that reaction? It can't be that my faculty of memory has made some commitment or that it bears goodwill towards me. It can do no such thing. My memory isn't responsive to my trust as such, either. My faculty of memory is too cognitively simple to recognize or be motivated by my trust in it. My memory is no separate agent at all. Rather, I feel betrayed because my memory had been tightly integrated into my basic functioning — until it started to let me down.

The external objects that evoke the strongest sense of betrayal are those whose functions are most tightly integrated into our own thinking and functioning: our musical instruments, our wheelchairs, our smartphones, our social media networks, our walking sticks, our cars. Even the more distant examples — like the ground — are part of our background system of affordances.<sup>7</sup> The ground's stability is a part of how I walk, and especially how I walk with ease. It is not exactly a part of myself, but it is tightly integrated into my background functionality. It is the loss of that effortless integration — the suggestion that the earth might have, so to speak, a mind of its own — that makes earthquakes so disturbing. Let me suggest, then, that the normativity here arises, not from there being any moral commitments in play, but from *teleological integration*. It is the normativity of integrated functionality, of parts

---

<sup>7</sup> Consider, for example, the well-known cases where an instrument seems to become an integrated part of one's perceptual system, like a walking stick. Classical discussions are in (Merleau-Ponty 1962) and (Gibson [1979] 2014). The theme has been taken up by the extended mind literature, especially (Clark 2008, 30-43). Notice that the argument I give here doesn't turn on any robust version of the extended mind thesis, that such external objects can become literally part of one's mind. My argument only depends on the weaker commonplace, that affordances can become phenomenally integrated into one's practical functioning and cognition.

knitted together into a functional whole. The negative reaction here towards the failure of one's trusted memory, I think, is one of alienation — a type (or at least close neighbor) of betrayal.<sup>8</sup> One feels alienated towards a part when one discovers that what one thought was a perfectly integrated piece of one's self is, in fact, failing to function as a smooth part of one's agency. That part is failing to be a good participant in one's functional whole. It seems perfectly appropriate for me to feel betrayed by my parts — my memory, my hands — for failing in their tasks. This reaction is, then, not entirely unrelated to agency. But it is not a reaction necessarily orientated towards independent, self-sufficient agents. It is a reaction directed at *parts of agents* by the whole agent — or by other parts of that agent — for failing the rest. And those parts need not be independent agencies in and of themselves in order to merit our sense of alienation when they come apart from us. Reproach at one's parts for failing is appropriate, not on moral grounds, but on grounds of functional unity.

Some might think that it is odd to think of betrayal as a response to the failure of one's integrated parts. Betrayal, in most philosophical accounts, turns out to be a specifically moral notion, directed at other people. This is why various analyses of trust have tried to ground betrayal in obviously morally-involved phenomena, such as responsiveness and commitments. But betrayal, it seems to me, is more deeply connected with notions of integration than it is with notions of commitment or responsiveness. After all, there are plenty of ways in which somebody can fail to be responsive to my needs or fail to live up to their commitments, but where I don't feel betrayed. I depend on the administrative assistants in my university payroll department to be motivated by my needs. Their failures might leave me frustrated and angry, but not *betrayed*. Betrayal is a more intimate notion; we are betrayed by those that are close to us, with whom we work in intimate concert. We are betrayed when something we were trying to make into a part of ourselves shears away from us, or when we are let down by somebody with whom we were trying to form a collective unit. It is far more natural to speak of being betrayed by one's memory than of being betrayed by some distant bureaucrat on whose cooperation one's visa application depends. The primary axis around which betrayal revolves, I suggest, is that of agential integration. Moral criticism often comes into the picture in those cases where we make use of various moral apparatus — like commitments — to help enable the integration.

---

<sup>8</sup> I don't mean alienation here in the very particular modern notion, such as the Marxist usage (Jaeggi 2014). I mean to be drawing on a more colloquial use of the term.

Holton suggests that the responses of trust are part of what he calls the *participant stance*. This is the characteristic stance that one agent takes towards another agent, which involves taking part in a network of agent-directed attitudes and actions: praise, blame, ascribing responsibility, feeling betrayal. And interpersonal trust of this sort plausibly occurs against the background of the participant stance. But thinking about object-oriented trust reveals another kind of stance we can take up, which we might call the *integrative stance*. This is our attitude towards things that we take to be part of us, and towards things with which we are supposed to be integrating to form some larger whole. I take the integrative stance to be my stance towards my own parts — like my hands and my memory — but also towards my fellow parts — like my fellow team-members, fellow employees, or fellow citizens. And I think the integrative stance begins to explain the fact that we feel betrayed by the failure of some objects, but not others. I may rely on my shelf to hold up my books, but I do not feel betrayed if it collapses — only deeply annoyed. But I have a much sharper reaction if the steering on my car suddenly breaks down, or if my computer mouse begins to respond erratically to input, or if my smartphone begins to scroll at random, or when the files on my computer desktop suddenly rearrange themselves, entirely unbidden. My car, my mouse, my laptop — these objects have come to be functionally integrated with me to various extents, and the breakdown of that integration is a violation from the point of view of the integrative stance. When I integrate other objects into my agency, I, in a sense, extend my agency into them, and so invest them with such status so as to be the appropriate objects of a very particular sort of reproach.

This suggests an account of the functional importance of the unquestioning attitude. For most of my sub-parts, holding a questioning attitude towards them would impede my efficient functioning. I need to trust my parts — and what this means is not simply that I rely on them, but that I can take my reliance for granted. When I truly make something a well-integrated part of my functional system, I drop the barriers. When I trust my memory, I let my various cognitive processes use the deliverances of my memory without a moment's hesitation. The same is true, I suggest, when we begin to incorporate external resources into our functioning. The unquestioning attitude lets me give external resources a direct pipeline into my cognitive and practical functioning. When one member of an elite and tightly knit unit of soldiers shouts, "Duck!" the other members simply duck. I trust the calendaring function on

my phone since I treat its alerts as immediate directives for where I am supposed to go, and its silence as an unquestioned indication that I have no immediate obligations. To trust something is to *let it in*, to let it muck about directly with one's practical and cognitive innards. To trust something is to attempt to bring it to the *inside* of one's practical functioning. Again: such trust is not indestructible. The unquestioning attitude can be defeated if the right sorts of evidence and considerations arise. Rather, to trust is to be strongly disposed to take the unquestioning attitude — to take the unquestioning attitude as a moderately static default stance

So here is the answer to our question about how we might ground negative reactions towards objects — how we might explain their air of normative bite. Trust is an unquestioning attitude. The primary use of the unquestioning attitude is as cognitive grease for functional integration. This explains the sharply negative reaction that arises from failures of trust. Such reactions are ones of alienation, arising from the integrative stance, towards things we thought were well-integrated parts of our functioning, when they fail to be well-behaved parts of our integrated, functional whole.

Let me offer a sketch of an even larger thought. This sketch will take us quickly over some very heady philosophical terrain, but I think it might be useful to scout out where this line of thinking might take us. Many philosophers have suggested that there is a deep relationship between our intentions and the unity of our agency over time. As Edward Hinchman puts it, when I make up my mind to do something at one time, I don't usually re-deliberate about that decision later on. It is self-trust, says Hinchman, that binds me together as an agent over time. Let's say that, this morning, I decided to make a chicken stir-fry for dinner. This evening, I buy bok choy and chicken breast because of that earlier decision. When I'm in the grocery store, I don't re-deliberate, treating my past self's decision as, say, testimony from an alien source which provides an input into my present deliberation. What it is to trust my past self is to simply buy the bok choy because my past self had settled on it. As long as things are going as usual, and my intention hasn't been defeated, my past deliberation has closed the matter (Hinchman 2003).

But it is not just trust in my past self; it is also trust in my present self's faculties to maintain the connection to the past self. I trust my memory. As Tyler Burge puts, memory doesn't just supply propositions about past events — "...it *preserves* [propositions], together with their judgmental force" (Burge

1993, 462). That is, if I perceived something in the past such that my perception was conclusive to my past self, and my memory conveys that conclusion to my present self, I don't relate to that remembered conclusion as one bit of evidence among many; I don't treat it as a mere input to some further deliberation. The conclusiveness itself transmits. For this to work, my memory must transmit the force of my past self's conclusiveness even if I don't remember the details of the process of reasoning that lead to that remembered conclusion. So long as I trust my memory, it functions as a direct pipeline from my past self into my present self. And we, as cognitively finite beings, all need to trust in this way. We must often engage in chains of reasoning that are longer than our consciousness can grasp in any single moment. We must, then, have the capacity to use our memory to integrate past conclusions into present deliberations, even when those past conclusions are presented to us shorn of their accompanying evidence and reasoning. That is the only way we can manage to pass long chains of reasoning through the limited pinhole of our consciousness.

I simply don't have the cognitive resources to constantly question my parts. The unquestioning attitude is needed for the seamless, efficient functioning of my integrated parts. When we trust others, I suggest, we are bringing them into a relationship roughly analogous to what we have with our own faculties. Self-trust and other-trust, then, turn out to be very much of a kind. When one trusts external informational sources, one grants them a similar cognitive status as one's own memory and other internal cognitive resources. When my spouse, who I trust entirely, shouts to me that the child has gotten his hands on a knife, I just start sprinting towards him. Her testimony is simply entered, instantly, into my set of accepted beliefs, just as would be a belief presented to me through my own memory. When I trust Google Search, I let its ordering of the search results direct my attention almost as if they were part of my own cognitive processes.

One use of the unquestioning attitude, then, would be to let one agent integrate other bits of the world into its system of cognition and action — to plug them in directly. Another use would be for individual agents to integrate each other, along with some non-agential resources, into a smoothly functioning collective agency — to approach, in their relationships with one another, the sort of the direct-pipeline relationship that we find between an individual's

internal cognitive faculties.<sup>9</sup>

My relationship to evidence, when I acquire a belief through trust in another's testimony, would then turn out to be something like my relationship to the remembered conclusions of my past self. Often, I don't possess the evidence and epistemic reasons for those past conclusions at the present moment. Rather, the conclusive force of the reasoning I performed in the past is transmitted to my present self — though stripped of awareness of the actual evidence and reasons that my past self reasoned with. Thus, I can be in a position where I am following the best norms of practical rationality, but yet still my present belief outruns the evidence that I presently grasp. Self-trust opens the door to acquiring beliefs for which we don't have immediate access to the supporting evidence. This may seem terrifying, but it is, in fact, the only way for cognitively limited beings to proceed. We must trust our past selves to have reasoned properly according to the relevant norms of deliberation. When our past selves fail to do well by those norms, we may be in a position where our self-trust brings us to believe what there is no good reason to believe, precisely because self-trust involves accepting the deliverances of our past self without re-checking our past self's reasoning.

My suggestion is that trust in external resources puts us in the same exquisitely vulnerable position towards them, and for similar reasons. Trust transmits the conclusive force of their reasoning; it transmits the conclusion to me, shorn of its support.<sup>10</sup> By trusting something else as an informational source, I can enter it into my cognitive network and take up a relationship towards it similar to the relationship I have to my own cognitive sub-faculties. It is not that, when I accept testimony through trust without deliberation, that I have necessarily failed to go through a proper practical deliberation. I am *deferring to deliberation that was run elsewhere*.<sup>11</sup> Again, this makes us terrifyingly vulnerable and makes our deliberative procedures vastly open-ended.

---

<sup>9</sup> For a compelling account of something like this, see (Hutchins 1996) for a classic study of how submarine crews act as a single mind.

<sup>10</sup> Some have worried that this sort of extended-mind approach to knowledge leads to a kind of epistemic bloat, in which we "know" far too many things. (Carter and Kallestrup 2019) offer a useful response to this worry by distinguishing between what we have in principle access to via extended faculties and what we have actually called forth into our awareness.

<sup>11</sup> Benjamin McMyler offers somewhat similar view. According to McMyler, when we accept a belief through testimony, we defer the justification of that belief to the testifiers. However, McMyler situations the deference in a voluntary taking of responsibility by the testifier (McMyler 2017). McMyler here is offering what has been called an assurance



But, as has been often observed, trust in others is the only way to proceed in the modern era, where human knowledge has vastly outgrown the reach of a single mind (or even of a single institution or discipline). We are no longer capable of individual intellectual autonomy; at best, we can autonomously manage our participation in a vast and distributed community of inquiry (Hardwig 1985; 1991; Millgram 2015; Nguyen 2018a).

My cognitive system typically runs with open pipelines, internally. What one part of me accepts, the other parts of me use without question.<sup>12</sup> The unquestioning attitude is the internal grease that lets me function quickly and efficiently. The unquestioning attitude also lets me weld open pipelines from outside resources into my cognitive and practical systems. And this also goes a long way to explaining the sharply negative reactive attitude when what's at the other ends of those pipelines lets us down. When we not only rely on a resource, but give it a direct pipeline into our thought and action, we are more profoundly alienated and disturbed when it goes awry.

Here's a real-life story — and an interpersonal echo of Hinchman's individualist story. My spouse and I keep a shopping list, in the form of a shared, cloud-based Google Doc that we each access from our smartphones. Each of us will add items to the shared list, usually without consulting the other. Then, when one of us is in the store, they simply buy everything that's on the list. When I am at the grocery store, I don't question the list. I don't try to remember which items I entered and which ones she did. I don't worry about whether or not she might have made some miscalculation or forgotten to update the list properly. I trust the list — which includes trusting my spouse, but also the hardware and software involved, as well as the procedures that my spouse and

---

view of trust — that what it is to trust somebody is to accept their assurances, in which they voluntarily take on responsibility for what follows from another's acceptance of their assurances. Such assurance theories make trust an essentially second-personal relationship — it is one where I trust you because you gave me your assurances about that trust. My view doesn't depend on any such action on the part of the testifier, or on any second-personal relationship. I can decide to trust somebody who has no idea who I am, and no relationship towards me, by observing their actions and following them without question. Imagine, for instance, that I am following somebody else through treacherous terrain. I can trust them by following, unhesitatingly, and stepping where they step. They need not offer me second-personal assurance for me to trust; in fact, they may not know I am there at all. My trust in them is entirely a matter of my own attitude towards their actions.

<sup>12</sup> The discussion of cognitive integration with external sources has been deeply inspired by Bryce Huebner's discussion of distributed cognition, and the kinds of integration required to count as distributed cognition (Huebner 2014).

I have instituted to maintain that list. And since I trust the list, I simply let its contents direct my actions without question, under normal circumstances. I trust the list in the same way that I trust my own memory about my past decisions. And trusting that shared list gives my spouse the power to input certain things directly into my practical reasoning and activity.

The unquestioning attitude account also helps to explain the divide between the sorts of objects with which we seem to engage in relationships of robust trust and the sorts of objects with which we don't. I have claimed that we can be betrayed by ropes, phones, computers, and the ground. On the other hand, I have seen far less talk of trust in and betrayal by the weather and by natural ecosystems. Farmers may rely on the weather, and when it fails them, they may be profoundly disappointed — but there seems to be no sense of profound betrayal. I may depend on the flowers in my garden to bloom, but if they do not, I am disappointed, but not betrayed. My account suggests a reason. The weather and my flowers are not immediately integrated into my system of practical affordances; I do not try to make them a part of my agential system. Likely, I don't try to integrate them because it is abundantly clear that they have some significant degree of independence. The ground, on the other hand, is mute, simple and seemingly easy to integrate. A smartphone is more complex, but it seems designed to be pliable and to conform itself to my will. These are the things that I try to integrate into my practical and cognitive self, and so they become things by which I can be betrayed.

The unquestioning attitude account also helps to explain the characteristic ways in which trust can go terribly wrong. Once we have welded together some cognitive pieces together with trust, errors can propagate easily. Cognitive elements that have been joined together with the unquestioning attitude are more efficient and more capable of seamless cooperation. But they are also more susceptible to infection as a whole.<sup>13</sup> According to the picture I've suggested, this is not a mere byproduct of trust. The upsides and downsides both arise directly from the fact that trust welds open pipelines directly into our functioning. The unquestioning attitude permits both collective efficiency and collective fragility.<sup>14</sup>

---

<sup>13</sup> For a relevant case study, see my discussion of echo chambers as trust manipulators (Nguyen 2018b).

<sup>14</sup> I am inspired here by Charles Perrow's discussion of natural disasters. According to Perrow, some organizational systems have "loose linkages", where each functional unit questions and interprets what's passed to it. Systems where a person has active interpretational agency at each juncture are such systems. Other systems have "tight linkages",

The unquestioning attitude account, then, could be taken as a first step towards a more radically non-individualist epistemology. The literature on trust has largely presumed that the basic unit of analysis is the individual. This presumption is shared across a broad swath of philosophy. Understanding our social and moral lives is about understanding relationships between autonomous individuals. Understanding our epistemic lives — even our social epistemic lives — is about, first and foremost, understanding how information is processed by individuals. But the unquestioning attitude account suggests a different take. The basic units of analysis could be larger collectives, and a form of trust could be the glue that holds them together helps assemble them into a collective. And betrayal could be the response, not of one individual to another, but of a part of a collective towards a recalcitrant part.<sup>15</sup>

However, this more radically non-individualistic line of thinking have only been intended as exploratory proposals, to feel out what a possible fuller account might be like. What matters most, for the present purpose, is to see that we have some need for cognitive and practical integration and that the unquestioning attitude has a clear role to play in such integration. We take the unquestioning attitude towards our own parts, but also use it to integrate other resources into our functioning. And the stance of integration brings it with certain loaded expectations, the failure of which leads to a sharply negative reaction. That is enough to see how failures of such integration can ground sharply negative attitudes of betrayal, or something very close to it.

And this helps us to bring reunite our discussion of trust with concepts of intimacy — a relationship that has become rather distant, of late, in the philosophical conversation on trust. Baier, in her originating discussion, made note of the deep association between trust and intimacy (Baier 1986, 247, 252). But

---

where each system simply takes what it's been given without interpretation and operates on it directly. Computer subsystem that simply takes a variable from another computer system and plugs that number directly into its calculations and operations — that is a tight linkage. Tight linkages, says Perrow, are very efficient, but they don't fail well. He attributes many kinds of systems failures — like the Three Mile Island nuclear meltdown — to cascading unpredictable failures in large, complex, tightly linked systems (Perrow 1999).

<sup>15</sup> I have been influenced here by recent literature in group agency (List and Pettit 2011; Gilbert 2015), especially Carol Rovane's discussion of the metaphysics of groups (Rovane 2019). Matt Strohl and I have offered a discussion of how some norms of behavior towards groups might be grounded in considerations of group intimacy — and about how acts of group intimacy might help constitute certain group agents (Nguyen and Strohl 2019). See also my discussion of how there might be emotional group agents (Nguyen 2019).

that connection has largely been lost, perhaps because philosophers seem to understand intimacy poorly and have usually avoided invoking the concept. But these thoughts about integration help us understand why trust and intimacy seem closely associated. Trust is aimed at achieving agential integration — about letting something inside, about uniting with it. And closeness and unification are some of the key markers of intimacy (Inness 1996; Nguyen and Strohl 2019).

## **6. Gullibility and Agential Outsourcing**

I have made a linguistic claim: that our natural use of “trust” includes the unquestioning attitude, and our natural use of “betrayal” includes disappointment from resources which we have taken the unquestioning attitude towards. But I don’t want these linguistic claims to get in the way of the more substantive proposal. What is most important here is the description of the phenomena itself. What I really care about is the unquestioning attitude itself and how it functions in our cognitive and practical lives.

And I think that it is vital that we get a handle on the unquestioning attitude, especially when it concerns our relationship to new and emerging technologies. Many of us, I think, have come to take the unquestioning attitude towards our smartphones, Google Search, and our social media networks. And this means that we have integrated complex processes and structures into our agency — often without adequate reflection how deep a change we are effecting. Each of these technologies structures our activities and cognitive processes in substantive ways. Google Search guides our attention. Social media networks filter what information gets to us, and what we pay attention to (Pariser 2012; Miller and Record 2013; Heersmink and Sutton 2018; Gillet and Heersmink 2019). Many of seem to have integrated our portable music players into our systems of emotional self-regulation (Krueger 2013; Colombetti and Krueger 2015). Infrastructural features of technologies can suggest conceptual schemes — like the menu bar on a news site suggesting a basic division of the important categories of news (Alfano, Carter, and Cheong 2019). And technologies can even suggest goals and structure our motivations. Gamified technologies can change our goals with respect to an activity. A fitness tracker, such as FitBit, highlights certain measures and, by giving the user daily scores

and rankings based on those measures, invites the user to change their reasons and motivations for physical activity. Such technologies can engender what I have called “value capture” — where clear quantified presentations of value take over in an agent’s self-conception, practical reasoning, and value capture (Nguyen 2020, 189-215).

This suggests an enlarged notion of gullibility. First, let’s start with what we already understand: the nature of gullibility in agent-directed forms of trust. The details will depend on which particular theory you accept. Perhaps we should only trust people if they actually have goodwill; or if they will actually be responsive to our needs; or will actually be likely to live up to their commitments. Gullibility, then turns out to be trusting somebody more than their trustworthiness warrants.<sup>16</sup>

What, then, is the analogous mistake with the unquestioning attitude? What would gullibility look like for this form of trust? Gullibility here would involve being too ready to set up those open pipelines — of being too quick to weld external objects into various aspects of our cognition and practical agency. The results are familiar when we take the unquestioning attitude towards informational sources. As with traditional gullibility, the problem there involves being too willing to accept the testimony of others. But the unquestioning attitude opens up the possibility of a new form of gullibility — what we might call *agential gullibility*. We can take the unquestioning attitude, not just towards informational sources, but towards processes that we incorporate into our agency. We can take an unquestioning attitude towards the agential infrastructure of the world. When we take an unquestioning attitude towards a news-site, we integrate its conceptual schemas and ways or organizing the world into our thinking. When we take an unquestioning attitude towards, say, a streaming musical service and use its algorithmically generated playlists to help regulate our emotions, we are integrating its emotional content — and its algorithmic selection process — into our system of emotional self-regulation. When we take an unquestioning attitude towards our FitBit, we let its embedded goals and metrics guide our valuing and decision-making. When we are agentially gullible, we are too willing to hastily integrate external resources into our own agency.

---

<sup>16</sup> One caveat: voluntarists like Holton think that our trust can exceed the trustworthiness of its target, when we have a reason to so exceed — like inspiring somebody to live up to our trust. Gullibility, in this case, would be trusting beyond what the trustworthiness of the target, combined with our good aspirational reasons, allow.

I am not here urging categorical resistance to the unquestioning attitude. Rather, I am suggesting that it is a powerful resource which also carries with it enormous risks. And its powers are inseparable from the vulnerabilities it creates. Those vulnerabilities are part and parcel of the basic functioning of the unquestioning attitude: to create efficiency by removing checks. Taking the unquestioning attitude is something like one country deciding to have open borders with another country, with all the efficiency, freedom, and vulnerability that entails. We should, then, try to deploy it appropriately and with great care.

Let's take a step back. Trust, in all its forms, runs far beyond our ability to manage or control. This is true even of mere reliance on testimony. Each person I rely on as an informational source has relied on others, who, in turn, rely on others. When I rely on my doctor's testimony, I am also thereby relying on whoever my doctor relies on: all their teachers, the medical researchers who provided the results which my doctor uses, the peer review process. And that reliance iterates, since I am also relying on whomever those latter people rely on: the statisticians on whose methods the medical research relies, the engineers who made their research instruments, and on and on. Reliance on testimony is *fractally iterated*. And because of that, we usually have no idea about how far our reliance extends, and on whom we are relying.

The danger is compounded with the unquestioning attitude, especially since the unquestioning attitude can be taken towards processes and agencies. This is already true for simple environmental features — when I trust a particular path or a ladder, my movements and decisions are significantly conditioned by those features. But the consequences for my agency are particularly sharp when I take the unquestioning attitude towards complex technologies. When I take the unquestioning attitude towards Google Search or my social media network, I am permitting complex processes to play a crucial role in my cognition and practical activity. Google Search is actively ranking and filtering search results. My social media network is actively amplifying some forms of discourse and suppressing others, as ranking algorithms intrude into what each node-member sees. In social media networks, dynamic network architecture encourages certain forms of expression to enter explosive viral feedback loops, while burying other forms of expressions out of sight (Tufekci 2017; 2018). Importantly, the unquestioning attitude doesn't simply add discrete, self-contained functions to our own agency. It *outsources* our agency — and that outsourcing can be iterated. When I trust Google Search, I let it guide my

attention, thus outsourcing a part of my agency. But I actually have very little idea what or who I'm outsourcing to — especially since Google Search itself outsources much of its operation to external resources. Google Search is built from modules collected from thousands of different researchers and technological institutions. And, what's more, those modules aren't all stable and finalized. Contemporary computing technologies often outsource their processing on-the-fly, each integrated resource itself open to constant revision and change.

This system is, of course, vastly powerful and efficient. (Try asking anybody who's lived with Google Search to give it up.) But agential outsourcing is particularly open to subterranean tinkering — and the more complex the trusted resource is, the more forms of invisible tinkering the truster becomes open to.<sup>17</sup> The basic functionality of Google Search might change without our knowing — and so a key part of our outsourced agency changes without our knowing. The agentially gullible person, then, is a certain kind of early adopter. Their mistake is being willing to outsource their agency too readily — to let any old thing in.

It might be useful, then, to update our paradigm of gullibility. The traditionally gullible person is the person who believes anything that anybody tells them. In our age, there is a new form of gullibility — a very particular type of agential gullibility. The technologically gullible person is the one that quickly and carelessly welds new form of technology into their agency. They unreflectively integrate smartwatches that highlight and make salient particular ways of valuing, in the form of metrics about exercise and sleeping, which can condition their values and motivations. They unreflectively integrate social media networks that transform their experience of discourse, argument, and interaction. They take up the unquestioning attitude without a moment's pause, towards any new technology, without considering the vulnerabilities and changes they're allowing into their agency.

## **7. Different forms of trust**

I have suggested that there is a form of trust that involves taking up an unquestioning attitude, and that this form of trust has an important place

---

<sup>17</sup> An interesting parallel to this discussion arises in Josh DiPaolo's analysis of the epistemology of fanaticism (DiPaolo forthcoming).

alongside the agent-directed forms of trust. But why do we call these two very different attitudes “trust”? And why are our negative response to both grouped together under the notion of “betrayal”? Let me end by suggesting that these various attitudes and reactions are grouped together because of their relationship to our attempts to expand our agency.

It will help here to focus, for the moment, on one particular account of agent-directed trust. Recall Jones’ responsiveness account of trust. According to her, to trust somebody is to depend on them because you take them to be trustworthy. And to be trustworthy is to be motivated to act to fulfill others’ dependence on you. Trust and trustworthiness go hand in hand, says Jones; they let us coordinate our actions by permitting us to actively depend on others. When we trust somebody, we know they will be responsive to our needs, and so we can take their expected responsiveness into account when deciding what to do.

Jones suggests that we can get clearer on the particular value and normativity of trust by imagining a world without any trust in it at all. Imagine, she says, a world full of all the other norms that we use, such as those of morality. Imagine that the people in this world follow these norms, and that they’re quite predictable. Imagine that everybody is fully rational and perfectly transparent to one another. But imagine that they simply do not trust, and that nobody is trustworthy, in her sense. This world, she says, would be perfectly safe to live in, but there would be something very important missing. Because nobody would act out of the awareness that they were being depended on — and nobody would depend on others to so act — “agents would lack the capacity to directly enlist the agency of another in the service of their ends” (Jones 2017, 100). What trust, in Jones’ sense, enables us to do is to “extend our agency” — to be able to recruit the agency of others into our own (101-2).

I think this is quite right — and plausibly right of any agent-directed theory of trust. And it points the way to a broader account of trust that encompasses both agent-directed trust and unquestioning trust. Both forms of trust are methods by which we seek to extend our agency, to integrate the functionality of parts of the external world into our own efforts. We have at least two tools for this integration: we can take into account others’ responsiveness to our needs, and we can turn off the questioning process. That is: we can cooperate with others and we can plug things directly into our agency. We can use these tools separately. But we can, and often do, deploy these tools together — as, for example, I usually do with other drivers on the road. And betrayal is the



characteristic response we have to failures of either form of integration.

Simpson suggests what unites all the various talk of trust is something very basic. The phenomenon at the center — the ur-trust — is, he says, simply the relationship we have towards other people we need to cooperate with. I suggest we borrow the structure of Simpson’s account, but swap in a new phenomenon for the center. The basic form of ur-trust, I’m suggesting, is agential integration. Trust — all trust — involves the attempt to bring other people and things into one’s agency, or of joining with other people and things into collective agencies. The various tools of interpersonal cooperation — commitments, goodwill, responsiveness — are part of one approach to agential integration, but there at least one other. There is also the adoption of the unquestioning attitude. And we are betrayed when we are let down by something with which we had tried to agentially unite. Trust, in the broad sense, then turns out to be a response to our essential cognitive and practical finitude. We need help, and we need to make the sources of that help things that we can rely on unquestioningly. Trust of both sorts involves various attempts to integration of other beings and objects into our practical functioning, of bringing them more into, or at least knit with, the boundaries of our selves.<sup>18</sup>

C. Thi Nguyen  
Utah Valley University

---

<sup>18</sup> I’d like to thank, for all their commentary, wisdom, and aid: Endre Begby, Julia Bursten, Anthony Cross, Sandy Goldberg, Kevin Lande, Neil Levy, Michaela McSweeney, Elijah Millgram, Geoff Pynn, Tim Sundell, Greta Thurnbull, Dennis Whitcomb, and Stephen White. I’d also like to thank the University of Kentucky Philosophy Department for a particularly exciting conversation that generated the entire discussion of agential gullibility and agential outsourcing.

## Bibliography

- Alfano, Mark, Joseph Adam Carter, and Marc Cheong. 2018. "Technological Seduction and Self-Radicalization." *Journal of the American Philosophical Association* 4 (3): 298–322.
- Askwith, Richard. 2015. *Running Free: A Runner's Journey Back to Nature*. London: Yellow Jersey Press.
- Baier, Annette. 1986. "Trust and Antitrust." *Ethics* 96 (2): 231–260.
- Baker, Judith. 1987. "Trust and Rationality." *Pacific Philosophical Quarterly* 68 (1): 1.
- Bratman, Michael. 1987. *Intention, Plans, and Practical Reason*. Center for the Study of Language and Information.
- Buchak, Lara. 2017. "Faith and Steadfastness in the Face of Counter-Evidence." *International Journal for Philosophy of Religion* 81 (1–2): 113–133.
- Burge, Tyler. 1993. "Content Preservation." *Philosophical Review* 102 (4): 457.
- Carabantes, Manuel. 2019. "Black-Box Artificial Intelligence: An Epistemological and Critical Analysis." *AI and Society*, 1–9.
- Carter, J. Adam, and Jesper Kallestrup. 2019. "Varieties of Cognitive Integration." *Noûs*. <https://doi.org/10.1111/nous.12288>.
- Clark, Andy. 2008. *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford: Oxford University Press.
- Colombetti, Giovanna, and Joel Krueger. 2015. "Scaffoldings of the Affective Mind." *Philosophical Psychology* 28 (8): 1157–1176.
- DiPaolo, Josh. Forthcoming. "The Fragile Epistemology of Fanaticism." In Michael Klenk (ed.), *Higher Order Evidence and Moral Epistemology*. Abingdon: Routledge.
- Faulkner, Paul. 2007a. "On Telling and Trusting." *Mind* 116 (464): 875–902.
- . 2007b. "A Genealogy of Trust." *Episteme* 4 (3): 305–321.
- Frost-Arnold, Karen. 2014. "Trustworthiness and Truth: The Epistemic Pitfalls of Internet Accountability." *Episteme* 11 (1): 63–81.
- . 2016. "Social Media, Trust, and the Epistemology of Prejudice." *Social Epistemology* 30 (5–6): 513–531.
- Gilbert, Margaret. 2015. *Joint Commitment: How We Make the Social World*. Reprint edition. New York: Oxford University Press.
- Gillett, Alexander, and Richard Heersmink. 2019. "How Navigation Systems Transform Epistemic Virtues: Knowledge, Issues and Solutions." *Cognitive Systems Research* 56 (56): 36–49.
- Goldberg, Sanford C. 2017. "Epistemically Engineered Environments." *Synthese*, 1–20. <https://doi.org/10.1007/s11229-017-1413-0>.
- Hardwig, John. 1985. "Epistemic Dependence." *Journal of Philosophy* 82 (7):

- 335–349.
- . 1991. “The Role of Trust in Knowledge.” *Journal of Philosophy* 88 (12): 693–708.
- Hawley, Katherine. 2014. “Trust, Distrust and Commitment.” *Noûs* 48 (1): 1–20.
- . 2017. “Trustworthy Groups and Organisations.” In *The Philosophy of Trust*, edited by P. Faulkner and T. Simpson. Oxford: Oxford University Press.
- Heersmink, Richard, and John Sutton. 2018. “Cognition and the Web: Extended, Transactive, or Scaffolded?” *Erkenntnis*, 1–26. <https://doi.org/10.1007/s10670-018-0022-8>.
- Hieronymi, Pamela. 2008. “The Reasons of Trust.” *Australasian Journal of Philosophy* 86 (2): 213–236.
- Hinchman, Edward S. 2003. “Trust and Diachronic Agency.” *Noûs* 37 (1): 25–51.
- . 2005. “Telling as Inviting to Trust.” *Philosophy and Phenomenological Research* 70 (3): 562–587.
- Holton, Richard. 1994. “Deciding to Trust, Coming to Believe.” *Australasian Journal of Philosophy* 72 (1): 63–76.
- . 2009. *Willing, Wanting, Waiting*. Oxford: Oxford University Press.
- Huebner, Bryce. 2014. *Macrocognition: A Theory of Distributed Minds and Collective Intentionality*. New York: Oxford University Press.
- Hutchins, Edwin. 1996. *Cognition in the Wild*. Revised ed. Cambridge, Mass: A Bradford Book.
- Ilgner, Arno. 2006. *The Rock Warrior’s Way: Mental Training for Climbers*. Desiderata Institute.
- Inness, Julie C. 1996. *Privacy, Intimacy, and Isolation*. New York: Oxford University Press.
- Jaeggi, Rahel. 2014. *Alienation*. Translated by Frederick Neuhouser and Alan E. Smith. New York: Columbia University Press.
- 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. Lanham: Rowman & Littlefield.
- . 2012. “Trustworthiness.” *Ethics* 123 (1): 61–85.
- . 2017. “But I Was Counting On You!” In *The Philosophy of Trust*, edited by Paul Faulkner and Thomas W. Simpson, 161–176. Oxford: Oxford University Press.
- Kahn, Douglas. 2013. *Earth Sound Earth Signal: Energies and Earth Magnitude in the Arts*. Oakland: University of California Press.

- Kahneman, Daniel. 2013. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.
- Kearl, Michael C. 1989. *Endings: A Sociology of Death and Dying*. New York: Oxford University Press.
- Keren, Arnon. 2014. "Trust and Belief: A Preemptive Reasons Account." *Synthese* 191 (12): 2593–2615.
- Krueger, Joel. 2013. "Affordances and the Musically Extended Mind." *Frontiers in Psychology* 4: 1–12.
- Lackey, Jennifer. 2008. *Learning From Words: Testimony as a Source of Knowledge*. Oxford: Oxford University Press.
- List, Christian, and Philip Pettit. 2011. *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press.
- McLeod, Carolyn. 2002. *Self-Trust and Reproductive Autonomy*. Cambridge, Mass.: MIT Press.
- McMyler, Benjamin. 2017. "Deciding to Trust." In *The Philosophy of Trust*, edited by Paul Faulkner and Thomas W. Simpson, 161–176. Oxford: Oxford University Press.
- Merleau-Ponty, Maurice. 1962. *Phenomenology of Perception*. New York: Routledge.
- Miller, Boaz, and Isaac Record. 2013. "Justified Belief in a Digital Age: On the Epistemic Implications of Secret Internet Technologies." *Episteme* 10 (2): 117–134.
- Millgram, Elijah. 2015. *The Great Endarkenment: Philosophy for an Age of Hyperspecialization*. Oxford: Oxford University Press.
- Nguyen, C. Thi. 2018a. "Expertise and the Fragmentation of Intellectual Autonomy." *Philosophical Inquiries* 6 (2): 107–124.
- . 2018b. "Echo Chambers and Epistemic Bubbles." *Episteme*, 1–21. <https://doi.org/10.1017/epi.2018.32>.
- . 2019. "Monuments as Commitments: How Art Speaks to Groups and How Groups Think in Art." *Pacific Philosophical Quarterly*. <https://doi.org/10.1111/papq.12279>
- . 2020. *Games: Agency as Art*. New York: Oxford University Press.
- Nguyen, C. Thi, and Matthew Strohl. 2019. "Cultural Appropriation and the Intimacy of Groups." *Philosophical Studies* 176 (4): 981–1002.
- Nickel, Philip J. 2012. "Trust and Testimony." *Pacific Philosophical Quarterly* 93 (3): 301–316.
- O'Neill, Onora. 2002a. *Autonomy and Trust in Bioethics*. Cambridge: Cambridge University Press.
- Pariser, Eli. 2011. *The Filter Bubble: How the New Personalized Web Is Chang-*

- ing What We Read and How We Think*. Reprint edition. New York: Penguin Books.
- Perrow, Charles. 1999. *Normal Accidents: Living with High-Risk Technologies*. Revised ed. Princeton: Princeton University Press.
- Resch, Michael, and Andreas Kaminski. 2019. "The Epistemic Importance of Technology in Computer Simulation and Machine Learning." *Minds and Machines* 29 (1): 9–17.
- Rini, Regina. 2017. "Fake News and Partisan Epistemology." *Kennedy Institute of Ethics Journal* 27 (S2): 43–64.
- Rovane, Carol. 2019. *The Bounds of Agency: An Essay in Revisionary Metaphysics*. Princeton: Princeton University Press.
- Simpson, Thomas W. 2012. "What Is Trust?" *Pacific Philosophical Quarterly* 93 (4): 550–569.
- Tufekci, Zeynep. 2017. *Twitter and Tear Gas: The Power and Fragility of Networked Protest*. New Haven: Yale University Press.