

This is an author-produced, post-review, pre-copyedit typescript of a chapter to appear in The Moral Psychology of Trust, eds. David Collins, Iris Vidmar Jovanović, and Mark Alfano (Lexington Books, 2023). Please cite the version of record.

OK, Google, Can I Trust You?

An Anti-Trust Argument for Antitrust

Trystan S. Goetze, Harvard University
tgoetze@fas.harvard.edu

Abstract

In this paper, I argue that it is impossible to trust the Big Tech companies, in an ethically important sense of trust. The argument is not that these companies are untrustworthy. Rather, I argue that the power to hold the trustee accountable is a necessary component of this sense of trust, and, because these companies are so powerful, they are immune to our attempts, as individuals or nation-states, to hold them to account. It is, therefore, literally impossible to trust Big Tech. After introducing the accounts of trust and power that I deploy, I argue that Big Tech companies have four kinds of power that render them unaccountable: fiscal power, political power, data power, and cognitive power. I conclude by reflecting on recent calls to break up the Big Tech firms, suggesting a new antitrust test in the light of my arguments.

1. Introduction: Trust and the Techlash

This leaves us with the question that I think we are still dealing with today: “What do you do when the most powerful institutions in society have become the least accountable to society?” And I think that’s the question that our generation exists to answer.

—Edward Snowden, speaking via video call to the Web Summit conference (2019).

Public confidence in technology firms has fallen precipitously since 2015, in the wake of data breaches, political scandals, biased AI systems, and suggestions that social media is a driving force in the rise of online misinformation anti-democratic authoritarianism.¹ In the media, this shift in

¹ In 2010, Pew Research found that U.S. Americans placed the tech sector with small businesses, religious organizations, and colleges and universities, as the only institutions of which they had a net positive impression (Rosentiel 2010). From 2015 to 2019, however, the proportion of respondents who had a favourable view of the tech sector fell from 71 per cent to 50 per cent; this drop was consistent across the Democratic–Republican partisan divide (Auxier, Anderson, and Kumar 2019).

public opinion has widely been called the *techlash*. In response, technology firms have stepped up their efforts to enact ethical best practices – or at least, to give the *appearance* of having enacted them (Floridi 2019) – in order to regain public trust. This is much as we might expect: in business as in personal relationships, the received wisdom is that when trust is violated, the once-trusted party must take steps not just to make amends to those whose trust was directly violated, but also to assure all who have trusted them, or who may need to trust them in the future, that they are worthy of being trusted again. And it is surely to the good if technology firms aspire to be worthy of our trust and take steps in this direction.

However, when it comes to the largest players at the centre of the techlash – namely, the *Big Tech* firms – it is possible that there is *nothing* they could do to earn our trust.² But this is not because their actions and inactions have been so egregious that they are forever untrustworthy. Rather, I will argue that these firms are untrustable. The main thrust of my argument is that, because of their power, there is no way for consumers, or even their political representatives, to hold these firms accountable. But, it is a central requirement for trust that the truster have the power to hold the trustee to account for violations of trust. Thus, no matter how the Big Tech firms may change for the better in response to the techlash, *we literally cannot trust them* as long as they are worth calling “big” tech.

Here is the plan. In §2, I outline a theory of trust, and explain why, on this view, it is a necessary condition of trust for the truster to be capable of holding the trustee accountable for violations of trust. In §3, I outline an account of power, and argue that the ability to hold a wrongdoer accountable is a form of power. In §4, I make the case that the power of Big Tech firms makes it impossible for us to hold them accountable, thereby rendering them untrustable. In §5, I address two objections. The first is that this view of trust and its relationship to power mean that children cannot trust their parents. The second is a worry that more needs to be said about group agency in order for group accountability or trust in groups to make sense. In §6, I conclude by reflecting on the implications of my argument for calls to break up the Big Tech firms.

² For the purposes of this paper, I take ‘Big Tech’ to refer to the so-called ‘Big Five’ – Alphabet (the holding company that owns Google, FitBit, DeepMind, and Waymo, among others), Apple, Amazon, Meta (the holding company that owns Facebook, Instagram, WhatsApp, and Oculus, among others), and Microsoft. However, there is good reason to expand this list to include Alibaba, IBM, Tencent, and Baidu (Webb 2019).

2. Trust

In moral philosophy, it has become customary to distinguish between *trust* and mere *reliance*. While we may trust *or* rely upon people, organizations, or states, we can *only* rely on things, animals, and autonomous systems. For example, I can rely on my word processing software to work without crashing unexpectedly, but I may trust the software developers to have designed a stable application. I can rely on a geyser to erupt at a regular interval, but I may trust the tour guide who assures me and my fiancée that Ol' Faithful will go off when we visit for a wedding photoshoot. I can rely on a bee colony to pollinate my crops, but I may trust the beekeeper whom I hired to bring the bees to my farm.

Both trust and reliance involve some element of expectation – one trusts or relies upon another to do something – and some element of vulnerability – if the agent one trusts or relies upon fails to do as one expects, one will be worse off. The nuances of the trust/reliance distinction are subject to much debate.³ For instance, Annette Baier argues that trust involves an expectation that the trustee will act with goodwill towards the truster (Baier 1986). By contrast, Russell Hardin argues that the trustee must act out of an interest in maintaining the trusting relationship, which involves “encapsulating” the truster’s interests within their own (Hardin 2002). For the purposes of this paper, I will deploy the theory developed by Margaret Urban Walker (2006) – she, in turn, draws on Richard Holton’s work (Holton 1994).⁴

I use Walker’s view over the alternatives for several reasons. First, as we shall see shortly, her account draws attention to the role of moral accountability in trust, which highlights features that are important to my critique of the Big Tech companies. Second, because she does not base trust in an assumption that the trustee has a particular motive or mental state (*viz.*, goodwill or encapsulated interests), it allows us to sidestep some questions about the moral psychology of group agents, such as large corporations. Third, Walker develops her account of trust as a component of her broader theory of *moral repair*, the process whereby we restore relationships after betrayals of trust and other forms of moral wrongdoing. As such, a critique of Big Tech from this angle can be seen as a step towards moral repair between Big Tech and the rest of society. Fourth, Walker’s view covers both the case of trusting someone *simpliciter*, and trusting someone

³ For a more thorough summary of the literature, see McCleod (2021).

⁴ These views focus on *moral* trust, but there are important differences when we think about *epistemic* trust. I am concerned only with the former in this paper.

to do something specific, whereas other views usually require a three-place relation: in Baier's terms, "A trusts B with valued thing C" (Baier 1986, 236). The broader scope of Walker's account is useful for my purposes, because, in my view, it isn't entirely clear *what* we might take ourselves to trust Big Tech to do (to provide their services? not to misinform us? not to misuse our data?), though it is clear that we now rely on them in many ways.

On Walker's view, to *rely* on someone or something is merely to expect that they or it will behave in a particular way, and to plan one's own life in anticipation of that outcome. Reliance is a purely descriptive concept; nothing about one's reliance on something implies that it *should* behave in the expected way, in any sense of "should." *Trust*, on the other hand, adds a normative dimension to reliance, which can only be satisfied by a responsible agent:

[T]rust links reliance to responsibility. In trusting one has normative expectations of others, expectations of others that they will do what they should and hence that we are entitled to hold them to it, if only in the form of rebuking and demanding feelings. (Walker, 2006, p. 80, italics hers)

In other words, when we trust someone, our expectations are not just about how they are *likely* to behave; rather, we expect that they will *act as they should*. And since this expectation is grounded in normative demands, we are entitled to hold the trustee accountable for failing to do as they ought.

The scope of these normative expectations depends on the context and on the relationship between the truster and the trustee. The truster may have general normative expectations that the trustee will act ethically, such as the implicit trust that we show our fellows that they will not attack us in the street. Or the truster may have expectations that the trustee will act as they should given their social role, such as the trust we have towards journalists that they will not mislead us or exploit the power that their social role gives them. Or the truster may expect that the trustee will act in accordance with the norms of a specific interpersonal relationship that obtains between truster and trustee, such as the expectation of faithfulness in romantic partnerships. Or the truster may have an expectation that the trustee will perform a specific task at a specific moment without exploiting or misleading the truster, such as the trust one shows to a stranger on the station platform when one asks something like, "Pardon me; is this the Chattanooga Choo-Choo?"

Why add the further detail about accountability?⁵ Walker defends this move by turning to Peter Strawson's much-discussed essay on moral responsibility (Strawson 1962), where he defends a distinction between two different stances we may take when responding to the harmful acts of another. When we take the *participant attitude* towards someone, we treat them as a fellow member of the moral community, someone who is a responsible agent, responsive to moral reasons and capable of goodwill towards others. When someone to whom we take the participant attitude causes harm or otherwise transgresses a moral norm, we may respond to their wrongdoing with communicative and punitive responses, such as feelings of resentment, the purpose of which is to make them appreciate the moral reasons that they have ignored or flouted, feel bad for the harms that their actions have caused, and take steps to make amends – as well as, potentially, our own private catharsis. The software developer, tour guide, and beekeeper are agents toward whom one might take the participant attitude – people we might blame if they fail to act as they should given the normative expectations we have of them. For Walker, accountability is crucial to her theory of moral repair. For, one component of making amends to one another after betrayals of trust is to respond to being held accountable for the betrayal by demonstrating one's remorse and commitment to doing better in the future.

By contrast, those to whom we take the *objective attitude* are treated instead as non-responsible entities who must be trained, incentivized, corrected, managed, restricted, avoided, or even terminated in order to protect ourselves from their harmful behaviour. We do not engage with them as responsible agents who are responsive to moral reasons, opting instead for strategies of behaviour modification that truncate or bypass any rational faculties that they may or may not have. In cases where changing the entity's behaviour is not possible, all we can do is regret what happened or be disappointed that something didn't happen, and move on. As Strawson notes, we sometimes must take the objective attitude towards other human beings, when attempting to reason with them would be futile or simply not worth the effort. On the other hand, the word processor, geyser, and bee colony are the sorts of things to which we might take the objective attitude, but *not* the participant attitude. We might be disappointed if they fail to act as we are relying on them to

⁵ In this paper, I understand the notion of *accountability* to encompass the activities we also refer to as *holding someone responsible* for wrongdoing. Holding someone accountable includes blaming, resenting, punishing, and other methods of attempting to make the wrongdoer recognize that they acted wrongfully and to do better in the future. See Shoemaker (2011) for discussion of this sense of "accountability" and how it is distinct from other senses of responsibility in the literature.

do, but it would make little sense to *blame* them. Therein lies the difference between things on which we can only rely, and people whom we can trust: to be trustable, one must be a fellow member of the moral community, the sort of entity that one could, in principle, hold accountable for betraying one's trust. Hence, reliance *plus* responsibility constitutes trust.

3. Power

The next piece of conceptual machinery I want to add to this account of trust is the notion of *social power*. The relative power of wrongdoer and wronged often goes missing in theorizing on moral responsibility, but it has profound effects on how we hold one another to account, and whether such efforts are safe and effective.⁶

The account of power that I deploy here is that presented by Miranda Fricker as a prelude to her discussion of epistemic injustice.⁷ On her view, social power is

a practically socially situated capacity to control others' actions, where this capacity may be exercised (actively or passively) by particular social agents, or alternatively, it may operate purely structurally. (Fricker, 2007, p. 13, italics removed)

Let's unpack this definition. By *practically socially situated*, Fricker means that power requires a functioning social world with shared institutions, meanings, expectations, and so on, because any particular operation of social power requires the coordination of various individuals. In addition, the notion of social situation, which stems from feminist accounts of how differences in social identity are relevant epistemologically and otherwise, highlights that aspects of one's social identity, such as one's gender or race, directly affect the amount of power that one has in various contexts.

Fricker also notes that power comes in two main types. *Agential power*, on the one hand, is a capacity that an individual agent can use to control the actions of others. It can be exercised actively or passively, meaning that those who have it can intentionally choose to use their power to exert control over others, or they may simply allow the shared understanding that they *have* such power to influence others' behaviour without taking direct action themselves. An example of agential power is a traffic cop's capacity to issue tickets for illegal parking. The cop can exercise this power actively, modifying people's behaviour by issuing tickets, or they can allow the shared

⁶ Baier (1986) includes a discussion of power in her account of trust, which I will return to in §5 when discussing an objection that Walker's account suggests that children cannot trust their parents.

⁷ Fricker, in turn, is drawing on Michel Foucault (1980, 1982, 2003) and Thomas Wartenberg (1992).

understanding that people may be ticketed for parking violations to do the work, without actively issuing any tickets.

Structural power, on the other hand, operates without the need for a particular agent who *has* power to exercise. One way structural power operates is by establishing, in our shared conceptions of the social world, that people with certain characteristics or identities tend to think or behave in certain ways, which then produces that behaviour in actuality.⁸ For example, a stereotype that men are better at mathematics than women may contribute to the under-representation of women in technical fields such as computer science. Another form of structural power is when opportunities available to some groups are not available to others. For example, educational inequality arising from inequitable distributions of government resources may mean some people have access to different educational opportunities. Alongside this, an image of some groups rather than others being “the sort of people” who seek and take up particular educational opportunities (e.g. university vs. trade school) may shape people’s decisions about which forms of education to pursue even if others are available to them.

Finally, while she acknowledges that theorists tend to invoke conceptions of power only when it is causing harm to someone or some group, Fricker notes that her account is agnostic as to whether particular operations of power are good, bad, or neutral in their effects. For example, in some neighbourhoods, such as those in a downtown core, it may be beneficial to residents if the traffic cop is aggressive in their exercise of the power to issue tickets for parking violations. But in others, such as a residential neighbourhood, aggressive ticketing may be a form of harassment.

Fricker notes elsewhere that the act of blaming – one of our principal ways of holding people morally accountable – can be thought of as “a moral species of social power” (Fricker, 2016, pp. 181-2).⁹ On her view, in blaming the wrongdoer (or holding them responsible by other means), the blamer is engaged in an attempt to change the wrongdoer’s behaviour. Namely, the blamer aims to spur the wrongdoer to recognize that they have done wrong, to acknowledge the moral reasons they ignored or flouted, to correct their moral understandings and/or future

⁸ Cf. Ian Hacking’s concept of the *looping effect* of human kinds (Hacking 1995), or Foucault’s notion of *subjectification*.

⁹ It is perhaps worth noting that not all forms of blame are exercises of social power. Private feelings of blame, because they are internal to the subject, are not exercises of power, for example. However, private blame is not a way of holding the blamee accountable.

behaviour, and perhaps to engage in moral reparative work. In context, Fricker's observation is intended to flag that the power of blame can be used for good – e.g. holding wrongdoers to account – or for ill – e.g. browbeating those with whom one disagrees into conformity. Hence, in this passage Fricker raises the concern that a blamer might have *too much* power over the blamee, or that the blamer may abuse that power.

But we can just as easily invert this connection between power and blame: *there are some instances where the blamer has too little power relative to the wrongdoer*. In these cases, the victim of wrongdoing lacks the capacity to influence the behaviour of the wrongdoer by holding them responsible for their actions. This situation can result from any of the forms of power presented above. The wrongdoer may be in a social position where they can actively use their agential power to silence or undercut their victim's attempts to hold them responsible. For example, in response to being accused of wrongdoing, the wrongdoer could take legal action to silence the victim, or terminate the victim's employment. Or the mere fact that the wrongdoer is in a position to take such actions may passively operate to silence the victim before any attempt to hold the wrongdoer accountable is made.

Structural power can also block the victim from holding the wrongdoer responsible. For example, the victim may be a member of a group that is socialized to be non-confrontational and submissive, discouraging them from speaking out about their mistreatment. Or, social institutions may fail to provide any practical mechanism by which the victim can hold the wrongdoer to account, for example, by pricing out people in the victim's socio-economic class from exercising their right to legal action, or by forcing the victim to use formal channels that are designed to frustrate attempts to bring about meaningful change, or by sustaining patterns of credibility deficit that result in testimonial injustices when the victim attempts to explain how they were wronged. Even if all of these different machinations do not deter the victim from confronting the wrongdoer with their blame, the wrongdoer's power may place them in a position from which they may safely dismiss or ignore the blamer's attempts to hold them to account without consequence, thereby undermining the very act of blaming.

For example, consider the well-known difficulties that accompany sexual harassment and sexual assault allegations made by students against professors. In such a situation, the student is at a disadvantage in their ability to hold their abuser to account because of their relative lack of social power. They may find the power that the professor holds in the university institution to be a

deterrent to reporting the incident in the first place. They may find that confronting their abuser leads to retaliation by way of lost opportunities or academic penalties. They may find it difficult for their allegations to be taken seriously by the relevant authorities. Their allegations might be dismissed as the words of a student against those of a colleague. If their allegations are taken seriously, they still might not be handled with care, or bureaucratic and legal processes may fail to deliver a meaningful response to the incident. Even if some resolution is reached, in many instances it will not have a lasting impact on the offender's career, undercutting the effectiveness of holding them to account. In each of these variations, the relatively powerful wrongdoer is able to escape some or all of the effects that we would expect to follow from the victim's attempt to hold them responsible, either by active or passive exercise of their social power, or by the operation of structural power. A powerful wrongdoer can overrule the power of the blamer to hold them responsible.

4. Big Tech

Let's put these pieces together now, and apply them to Big Tech.

When one is incapable of holding a person responsible, this precludes one's ability to trust them. This is because, as we saw, trust is reliance plus responsibility. On the account of trust that I presented, we can only trust entities of which we can have normative expectations. Built into this notion is that we must be able to take the participant attitude – to treat the trustee as a member of the moral community – which implies that we have the power to hold the trustee accountable when they transgress moral norms. An entity that one cannot hold morally accountable is, thus, not an entity toward whom one can take the participant attitude. And so, *if one lacks the power to hold someone or something accountable for their actions, that person or entity is not something that one can trust.*

As suggested by Snowden in my epigraph, there are few institutions in modern life more powerful than the Big Tech companies.¹⁰ First, they are among the wealthiest institutions in the world, granting them significant agential power. That is to say, because they are wealthy, they have the resources to actively exercise control over others in a variety of ways, and this capacity also

¹⁰ In context, Snowden's remarks are about not just Big Tech but also Big Government (Web Summit 2019). By the latter, I mean law enforcement, intelligence, and military organizations, and the governmental and judicial bodies that have consistently supported these institutions in expanding state and corporate surveillance.

operates passively to discourage those who may be harmed by their actions from pursuing remedies. For example, Alphabet's annual revenue in 2019 was about US \$162 billion (Wallach 2020), more than the 2020 GDP of Ukraine and over one hundred other sovereign nations (World Bank 2020). The sheer fiscal power of these firms, then, puts them quite out of reach of attempts by ordinary citizens, and even a good number of nations, to hold them accountable for their wrongdoing. While internal dispute mechanisms offered by these firms are available in some instances, their enforcement is notoriously capricious and changes frequently and without notice. While lawsuits are an option, it is safe to say that very few people possess the means to launch a legal case against a multinational with the resources of a medium-sized country. And this is not to mention that the terms of use for the services provided by Big Tech – binding contracts that we all agree to without reading – often explicitly require disputes to be resolved via arbitration, rather than civil suits. Furthermore, in the United States, case law has often favoured technology corporations. For example, since the 1996 decision in *Zeran v. America Online, Inc.*, companies controlling online communications platforms, such as social media, cannot be sued in the USA for harmful content posted by their users (Sheridan 1997). Taking these factors together, practically speaking, very few of us can hold Big Tech accountable.

What about our elected representatives? One reason to have a government, after all, is that it can pool resources and enforce regulations intended to protect us from powerful bad actors. But in part because of their tremendous economic footprint, Big Tech firms also command a great deal of *political* power, not just in back-room lobbying, but also in tactics that increasingly resemble nation-to-nation tit-for-tat in diplomatic disputes. Indeed, writing in *The Atlantic*, Adrienne LaFrance describes Facebook as a “hostile foreign power” and “the largest autocracy on Earth,” which is “engaged in a cold war with the United States and other democracies” (LaFrance 2021). For example, in early 2021, the Australian parliament introduced draft legislation that would require Big Tech companies to share their advertising revenue with news companies whose links are shared on social media and social news platforms. In response, Facebook blocked all news from being displayed to Australian users, and Google threatened to disable its Search tool for Australians – both in the middle of the biggest public health emergency in the last century (BBC News 2021). Google backed down from their threat after global backlash; Facebook ended their information blockade only after concessions from legislators were made (Cellan-Jones, 2021). This was undoubtedly a flexing of power intended to test just how far these companies can go to resist

efforts to regulate them. And, since the Australian legislation was intended in part to rectify a harm caused by companies such as Facebook and Google – namely, the decline of traditional journalism, and with it, journalistic standards in online media – their resistance to this regulation can be seen as resistance to being held accountable. Not every nation and not every government can stand firm against such resistance – and, as mentioned, Australia was forced to compromise. This form of power undermines even the efforts of powerful democratic governments to hold Big Tech to account.

A third kind of power possessed by Big Tech is tied to their control over our *data*. In exchange for the free services offered by these companies, we allow them to collect, process, and sell data on many facets of our lives. These data are collected not just from our interactions with Big Tech services, but from our online and technologically enabled activities generally. Carissa Véliz argues that in giving up our data so freely, we are abdicating a distinctive kind of power. As Francis Bacon once wrote, knowledge is power, an adage Véliz reads in a pragmatist spirit: the more you know about someone or something, the more effectively you can plan your actions when they concern the subject of that knowledge. Véliz connects this observation to the idea that by collecting data on a person, one is collecting knowledge about them, and that with knowledge comes power: “Through protecting our privacy, we prevent others from being empowered with knowledge about us that can be used against our interests” (Véliz 2020).

In other words, because Big Tech companies understand so much about us from our data – in many cases, better than we understand ourselves – they can tap into this knowledge to predict our thoughts and actions in response to various situations. As Simson Garfinkel wrote in 2000 (when Google was still just a search engine and the most popular social media platforms were LiveJournal and classmates.com),

next-generation agents will scan the world for personal information about an individual, then construct a predictive model for use by marketers and others...The profile could know every document you’ve ever read, every person you’ve ever known, every place you’ve ever been, and every word you’ve ever said that has been recorded. Your identity would no longer exist just inside of you, but in the model. (Garfinkel 2000, 252)

Garfinkel refers to this process as the *extraction of the self*: a model of one’s self is captured in a computational system and used to predict how one will behave. The level of data power that the Big Tech firms have today makes extracting the self possible. This is both a kind of agential power

that enables Big Tech to “nudge” our behaviour subtly in ways that may benefit them – including pushing us away from attempts to hold them to account – and a kind of structural power that makes life without their services difficult or nigh-inconceivable.¹¹

Indeed, if we were to try to hold Big Tech accountable by denying them access to our data, they would still have ways to coerce our compliance. While many of these firms put on airs of giving us data control, our power to take control of our data can actually be quite limited. Facebook and Google, for example, prompt the user on occasion to perform a “privacy check-up” to confirm their data protection settings. But while these controls allow *some* data collection and processing to be switched off, there is no option to eliminate it entirely or even to reduce it to strictly necessary purposes. In fact, these firms perform some data collection and aggregation on you even if you do not have an account registered with them, thanks to trackers connected to their advertising services across the web. We would not fare much better were we to attempt to boycott their services instead. It has become difficult to opt out in this way because Big Tech services underpin much of the structure of the modern internet. For example, Amazon Web Services controlled a third of the market for cloud computing in 2020, and hosts a significant number of other companies’ online services, including Netflix and Slack (Runkevicius 2020). Attempts at digital boycott are likely to be incredibly difficult and ineffectual, as reporter Kashmir Hill discovered when she struggled through an attempt to live without interacting with any of the Big Five for a week while carrying on a normal life as a digital citizen – in her words, “it was hell” (Hill 2019).

The fourth kind of power possessed by Big Tech is what we may call *cognitive* power. This is the power not just to influence our behaviour, but to shape our very thoughts and values, which, it goes almost without saying, are some of the most important sources of human behaviour. As James Williams argues, the Big Tech companies have cognitive power in virtue of how they command our *attention* (Williams 2018). Using a series of illuminating metaphors, Williams outlines three levels of attention that can be hijacked by the platforms controlled by Big Tech. First, there is what he calls the *spotlight*, our capacities to focus our immediate attention on a particular task or object. Second, there is the *starlight*, our broader capacities to direct our actions so that they will align with our values. Third, there is the *daylight*, our capacities that enable us to have and to reflectively

¹¹ On nudging and the subtle power of design to influence behaviour, see Thaler and Sunstein (2008).

revise our beliefs and values in the first place. Our spotlight can be misdirected by technological intrusions and nudges, such as notifications and alerts that pull our attention away from what we mean to be doing and back to our social media news feeds. We can be pulled off-course from our guiding starlight by coming to internalize the reward mechanisms that exist within the platforms controlled by Big Tech. For example, Williams reports feeling compelled to act in ways that would maximize the number of “likes,” “favourites,” “follows,” “friends,” “connections,” “shares,” “reblogs”, and so on, on various platforms. He attributes a growing attitude of pettiness within himself to this internalization of social media rewards as valuable objects to pursue, and goes on to connect this sort of pettiness to political polarization, increases in the incidence of narcissism, and deadly risk-taking behaviour by social media personalities. Finally, the daylight by which we discern what is true and good can be occluded by the online spread of misinformation and moral outrage, which are rewarded by algorithmic content feeds that prioritize a thin, behaviouristic metric of “engagement” over thicker, more meaningful measures of quality.¹²

Cognitive power is especially dangerous to our ability to hold Big Tech accountable. By leveraging this form of power, these companies can keep us distracted from their wrongdoing (misdirecting our spotlight). They can encourage us to act as if the things that their platforms value are what *we* value, keeping us under their spell (pointing us away from our starlight). And they can confuse us as to what is really true or false, and who is really good or bad, by virtue of the content their platforms serve up (eclipsing our daylight). Each of these forms of cognitive power can undercut our ability to hold Big Tech companies accountable by preventing us from doing so in the first place – or by preventing us from even seeing the need.

Taken together, the fiscal power, political power, data power, and cognitive power possessed by Big Tech firms renders them out of reach of attempts to hold them accountable. This goes not just for individuals, but also for many – potentially all – nation-states. But, it is a requirement of trust that the truster have the power to hold the trustee accountable for violating normative expectations. Therefore, it is impossible for us to trust the Big Tech companies. Furthermore, the fact that Big Tech can neither be trusted nor held accountable suggests something even more disturbing: *these companies may not be members of the moral community*. Since their power insulates them from the giving, asking for, and responsiveness to moral reasons; expectations of

¹² Cf. Nguyen’s discussion of how social media, and Twitter in particular, gamify communication (Nguyen 2021).

goodwill; and responsiveness to being held responsible that are part and parcel of participation in the moral community, taking the participant attitude towards them no longer makes sense. Big Tech companies are more akin to autocrats – or dangerous beasts – and should be treated as such.

5. Objections

Before we come to what we might do to counter the situation we find ourselves in, in this section, I address two objections to my arguments. The first concerns the potential consequences of the view of trust and accountability that I have presented, namely, that this view implies that children cannot trust their parents. The second concerns the nature of Big Tech firms as group agents: it is possible that accountability simply works differently when dealing with groups rather than persons.

5.1. Think of the Children!

A potentially troubling consequence of the account of trust that I have presented is that children might not be able to trust their parents or guardians.¹³ After all, compared to adults, children have much less social power, which may well mean that they lack the power to hold adults to account. Indeed, it is an important aspect of childrearing that parents retain a significant degree of power over their children; when exercised with care, parental power helps ensure the safety, healthy development, and moral education of children. But at the same time, we also think that it is important for children to be able to trust their parents; in fact, this trust might be taken to underpin the permissibility of parental power as minors gradually become autonomous agents. Yet, on the account I have presented, a child's ability to trust their parents may seem at odds with their relative lack of power. If children are relatively powerless compared to their parents, this might well preclude their ability to hold their parents accountable for wrongdoing, which would, according to the arguments I gave above, preclude children from being able to trust their parents.¹⁴

One way to respond to this objection would be to simply bite the bullet. Perhaps children cannot trust their parents after all. It may be that children can only rely upon their parents. This

¹³ For brevity, I will refer only to parents from here on, but it should be understood that the conception of *parent* that I have in mind is broad, including different- and same-gender, biological, surrogate, adoptive, foster, and polyamorous parents, as well as legally assigned guardians, grandparents or godparents assuming a primary parental role, other adults in an extended family, and so on. Who counts as a parent in one family or another is dependent on their role in that family, not biological relation.

¹⁴ Thanks to Carolyn McLeod, Letitia Meynell, Sue Sherwin, and others, for independently pushing me to respond to this objection.

conclusion, while reshaping some of the contours of how we understand trust in our social relations, does not strike me as especially strange. Given that children are, in many other ways, developing their capacities as moral agents, their participation in the moral community is already truncated. Why should their ability to trust be any different?

Another potential response to this objection is to note that in fact children *do* sometimes have the power to hold their parents accountable – though they may not always understand what they’re doing. For example, my mom is fond of telling an anecdote about my childhood, wherein I was upset at some demand or restriction of hers, and, in a fit of tiny rage, I exclaimed, “You’re a bad mummy!” Now, she knew full well that this was not true – the source of contention, though I don’t remember it now, was a perfectly reasonable request – yet this reproach still hurt her feelings and made her question herself for a moment. If a child can succeed in rebuking their mother out of misplaced frustration, surely they can do so when their parents have genuinely done something wrong.

Or, it may be that a different, but still ethically significant, sense of trust obtains in the case of children and parents. Baier describes “trust between infant and parent” as a “primitive and basic trust” (1986, 245). By this, she means that children trust their parents by default. In fact, children *must* uncritically trust their caregivers, both as a matter of biological survival and because they initially lack the moral reasoning capacities needed to trust in the more robust sense outlined by Walker. If this primitive form of trust is distinct from the sense I have employed, however, it is plainly not the sort of trust we should have in Big Tech: it would be an affront to human dignity for us to be reduced to a child-like state of helpless, uncritical vulnerability *vis-à-vis* these corporations.

A final line of response to this objection is possible by noting that the problem of children’s ability to trust their parents is just one form of a more general problem. Namely, when someone or some group of people is vulnerable and relatively powerless compared to someone else, how can we ensure that trust is still possible? This problem occurs with regard to the relationship between citizens and the state, between students and teachers, between patients and medical professionals, between civilians and the police, between account holders and banks, and generally whenever the vulnerability of the truster is heightened by the greater power of the trustee. In each of these cases, we set up formal structures to ensure accountability, and so secure the possibility of trust. Citizens can oust their leaders in elections. Teachers and medical professionals are beholden to the policies of their professional associations, and violating their codes of conduct leads to a bar on practising

in their field. The police have internal and external oversight. Banks are subject to various consumer protection regulations. And in all of these cases, laws and practices have been established to give vulnerable people the ability to hold those who abuse their power to account in courts and tribunals. These systems are far from perfect, but when they work, their existence is an essential part of what makes it possible for us to trust powerful people and institutions.

The same is true for children. Most democratic countries have strict child protection laws, as well as legal duties for those acting in a parental role. Child neglect and abuse are serious crimes. Special government agencies exist to protect children whose vulnerability has been exploited or ignored by parents. Other adults have legal obligations to report suspected wrongdoing by parents. And underpinning these legal mechanisms is a prior moral commitment common to all folk moralities that the vulnerability of children means that they are owed special protection and care, which adults are morally obligated to provide. Even if children are not often the ones who launch such legal and ethical accountability procedures, the fact that these systems exist enables children to maintain an attitude of trust towards their parents. Though they might not be able to hold their own parents accountable themselves, others can do it for them.

The corollary, of course, is that when such institutions of accountability are systemic failures, vulnerable people cannot trust the powerful after all. We have seen this most starkly with the police in recent years, as their established impunity in unprovoked killings of people of colour is subject to repeated challenge and outcry, to little avail. So far as the theory of trust goes, we should welcome this result, as it provides a compelling and urgent reason to ensure that these formal systems of accountability function properly. It explains, in part, why the police cannot be trusted, why many who grow up in foster care do not trust the social services responsible for that system, why many marginalized people are leery to trust medical professionals, and so on. Moreover, if my arguments above are right, our situation with regard to Big Tech is substantively similar to that of the child whose society has failed to provide working formal accountability processes. Our regulatory frameworks are largely toothless when it comes to Big Tech, and that is part of the problem.¹⁵

¹⁵ My account might also imply that we cannot trust all-powerful entities, such as God. What I would suggest in response is that we can make a distinction between trust and faith. This might then show that while we cannot trust the powerful when they are unaccountable to us, we can still have faith in them. Thanks to several audience members at presentations I gave of this paper for raising this concern.

5.2. Group Accountability

Another worry might be that I have assumed too much in using an account of trust that was designed for trust *between human persons* to analyze a case of trust *between human persons and large businesses* – that is to say, between individuals and collectives. If groups can be held accountable, we might think that it must be in a way different from Strawson’s conception of the participant attitude. One might contend that because a collective, as such, has no mental states, it cannot feel shame and cannot be subject to the power of moral blame to reshape its moral understandings. Or, if collectives *can* be subject to the power of moral blame, at the very least we are owed some account of group agency that explains how.¹⁶ While a full defence of an account of group agency and the responsibility of collectives is beyond the scope of this essay, I do have a few potential responses to this objection.

One approach would be to adopt an account of group knowledge on which we can ascribe moral understandings to a collective. The epistemology of collectives is largely focused on how to explain the fact that collectives have knowledge, rather than establishing that group knowledge is possible.¹⁷ And, if a group can have knowledge or beliefs, then, *a fortiori*, a group can have moral knowledge or beliefs. It stands to reason, then, that a collective could be influenced to change its moral beliefs in response to reproach or other forms of pressure.

However, even if one rejects the possibility that a collective, such as a corporation, could have moral beliefs or understanding, the possibility remains that a collective could be influenced by the efforts of individuals or other collectives. There are several accounts of group agency that pay particular attention to how we can hold corporations and similar collectives to account.¹⁸ Adjudicating between these positions on group beliefs and group accountability is not important here; what matters is that a cogent account could, in principle, explain how corporate agents can be held morally accountable.

Finally, regardless of how one conceives of collectives, the fact remains that they are at least partially constituted by individual agents. And even if the collective has no psychology to influence, its members do. By holding individual members of the collective to account, especially those in

¹⁶ Thanks to Ian Brooks for suggesting this objection.

¹⁷ See, among others, Bird (2014), Gilbert (1989; 1987; 2004), Tuomela (1992), Wray (2001).

¹⁸ See, among others, Cooper (1968), Feinberg (1968), French (1984), May (1992; 1987), Mellena (1997), Pettit (2007).

positions of power within the collective's formal decision-making procedures, the collective itself can be held responsible for its actions.¹⁹

Any remaining skeptics with regard to group accountability will have to contend with the fact that concerted efforts to influence morally salient corporate behaviour have sometimes worked. For example, from the 1930s to 1980s, a class of chemicals called chlorofluorocarbons (CFCs), which include Freon, were commonly used as refrigerants and aerosol propellants. In the late 1970s, scientists discovered that CFCs were destroying the Earth's ozone layer, which shields terrestrial life from dangerous ultraviolet solar radiation. By 1989, an international coalition of governments had successfully passed regulations to phase out and ban CFCs in favour of non-destructive chemicals. This top-down measure likely would not have succeeded without bottom-up pressure from ordinary citizens and activist groups, who were quite reasonably outraged that corporations were putting life on this planet at risk in pursuit of profit. And had it not succeeded, a NASA Earth Observatory model predicts that by 2020, the concentration of the ozone layer over North America would have declined by about half, and would have nearly completely been depleted by 2060 (Carlowicz, Lindsey, and Simmon 2009). Regardless of whether one thinks that manufacturing corporations had a change of *moral belief* regarding the permissibility of using CFCs in their products, the power exercised by concerned individuals, groups, and states succeeded in changing these corporations' behaviour for the better. And this is one of the goals of exercising the power of moral blame. So I see no reason not to think that blame can, in some situations, effect positive change in corporate behaviour.

The case of Big Tech, however, is disanalogous to that of CFC-using manufacturers. Big Tech is, well, bigger, in the ways I have described earlier. Their power is greater and further-reaching than that of companies who marketed aerosols and refrigerants in the 1970s. Attempts to regulate Big Tech thus run up against more difficult barriers than environmental protection legislation in the late twentieth century – which, of course, is not to say that such regulations were easy to pass in the first place. If my arguments in this paper are right, then even if there is a possibility that corporations' moral beliefs or morally salient actions can be controlled by accountability mechanisms, Big Tech firms remain an exception: their power renders them unaccountable, untrustable, and dangerous.

¹⁹ Thanks to Duncan MacIntosh and Jeff Behrends for suggesting this point.

6. Conclusion: Break Them Up

In this paper, I have argued that the Big Tech companies cannot be trusted. The issue is not that they are untrustworthy, but rather, that they are untrustable. The reason is that, on the view of trust that I presented, trust requires that the truster have the power to hold the trustee accountable when the trustee fails to act as they ought. And since the Big Tech companies possess a tremendous amount of power – fiscal, political, data, and cognitive – they are able to resist and undercut attempts by individuals and states to hold them to account. Thus, they quite literally cannot be trusted. In light of this conclusion, I want to briefly suggest some next steps.

When an entity causes harm, but is not one to whom we take the participant attitude, Strawson tells us that we may resort to cruder means of correcting its behaviour or protecting ourselves from it. We may even be permitted to do so with someone who is otherwise a responsible agent, should they fail to respond to our reasonable attempts to hold them accountable. By retreating to the objective attitude, new lines of action open up – potentially including *violence*. And one way to strike back violently against a corporation is to break up its holdings.²⁰

In the nineteenth century, society was faced with a similar concentration of power in the hands of a small number of companies. The oil barons – among them John D. Rockefeller, the owner of Standard Oil and one of the richest men in world history – controlled so much wealth and influence that the Canadian and U.S. governments were impelled to enact some of the earliest modern antitrust legislation. The breakup of the Standard Oil monopoly followed in 1911, and laid out the conditions for forcing the break-up of a large corporation or trust: break up the entity when it can raise prices without losing customers.

Today, it is common to compare the Big Tech companies to the oil barons with the slogan ‘Data is the new oil’ (Economist 2017). As Catherine D’Ignazio and Lauren Klein observe,

It’s a metaphor that resonates uncannily well... The idea of data as some sort of untapped natural resource clearly points to the potential of data for power and profit once they are processed and refined, but it also helps to highlight the exploitative dimensions of extracting data from their source – people – as well as their ecological cost... (D’Ignazio and Klein 2020)

²⁰ An alternative could be to somehow *legitimate* the power held by Big Tech. For some suggestions on this, see Greene and Gilbert (ms.).

But the old antitrust test fails with the corporate powers that control the new oil. Because many of their services are *free* – or rather, offered to us *gratis* in exchange for our data – there are no prices for them to raise. A new antitrust test might emphasize a company’s ability to escape accountability for its actions: *the more easily a company can evade accountability, the stronger the case for breaking it up*. The point of such a test is not to protect the *market*, but to protect the *people*. By breaking up the Big Tech companies, we may reduce their power by reducing their financial assets, their political influence, their data assets, and their avenues for misdirecting our attention. In so doing, we may restore these companies’ status as participants in the moral community, instead of as dangerous forces outside it.

Acknowledgements

I presented earlier forms of this paper several times in 2021: at Congress of the Canadian Philosophical Association, to the Dalhousie University philosophy graduate students’ writing group, as a job talk at the University of King’s College, to the Talk Shop at Harvard University, and to the Philosophy Colloquium at Dalhousie University. For feedback on earlier versions of this material, I’d like to specifically thank Carolyn McCleod, Jacquie Burkell, Char Harrison, Clarisse Schuetz, Lara Millman, Erik Nelson, Katrina Ingram, Duncan MacIntosh, Darren Abramson, Letitia Meynell, Sue Sherwin, Richmond Campbell, Ian Brooks, Jenna Donohue, William Cochran, Jeff Behrends, Selim Berker, and Marc-Kevin Daoust. This work was supported by a Banting Postdoctoral Fellowship from the Social Sciences and Humanities Research Council of Canada and an Honorary Research Fellowship from the University of Sheffield. The lands on which I wrote this essay are the traditional territories of the Mi’kmaq, Mississauga, and Massachusett peoples.

Works Cited

- Auxier, B., Anderson, M., & Kumar, M. (2019). 10 tech-related trends that shaped the decade [Survey]. Pew Research Center. <https://www.pewresearch.org/fact-tank/2019/12/20/10-tech-related-trends-that-shaped-the-decade/>
- Baier, A. (1986). Trust and Antitrust. *Ethics* 96(2): 231–60.
- BBC News. (2021, February 18). Australia news code: What’s this row with Facebook and Google all about? *BBC News*. <https://www.bbc.com/news/world-australia-56107028>

- Bird, A. (2014). When Is There a Group That Knows? In *Essays in Collective Epistemology*, ed. J. Lackey. Oxford: Oxford University Press. 42–63. DOI: 10.1093/acprof:oso/9780199665792.003.0003
- Carlowicz, Michael, Rebecca Lindsey, and Robert Simmon. (2009, May 13). The World We Avoided by Protecting the Ozone Layer. *Nasa Earth Observatory*. <https://earthobservatory.nasa.gov/features/WorldWithoutOzone/page1.php>
- Cellan-Jones, R. (2021, February 23). Facebook v Australia: Who blinked first? *BBC News*. <https://www.bbc.com/news/technology-56168843>
- Cooper, D. E. (1968). Collective Responsibility. *Philosophy* 43(165): 258–68. DOI: 10.1017/S0031819100009220
- D’Ignazio, C., and L. F. Klein. (2020). *Data Feminism*. Cambridge, MA: MIT Press. URL: <https://datafeminism.io/>
- Economist, the. (2017). The world's most valuable resource; Regulating the data economy. *The Economist* 423(9039). URL: <https://www.economist.com/weeklyedition/2017-05-06>
- Feinberg, J. (1968). Collective Responsibility. *The Journal of Philosophy* 65(21): 674–88. DOI: 10.2307/2024543
- Floridi, L. (2019). Translating principles into practices of digital ethics: five risks of being unethical. *Philosophy and Technology* 32(2): 185–93. DOI: 10.1007/s13347-019-00354-x
- Foucault, M. (1980). *Power/Knowledge: Selected Interviews and Other Writings 1972–1977*. C. Gordon, (ed.); C. Gordon, L. Marshall, J. Mepham, & K. Soper (trans.). Brighton: Harvester & Wheatsheaf.
- . (1982). How is Power Exercised? In *Michel Foucault: Beyond Structuralism and Hermeneutics*. H. Dreyfus & P. Rabinow (eds.), & L. Sawyer (trans.). (pp. 208–226). Brighton: Harvester Press.
- . (2003). The Subject and Power. In P. Rabinow & N. Rose (Eds.), *The Essential Foucault: Selections from Essential Works of Foucault: 1954–1984* (pp. 126–144). New York: The New Press.
- French, P. A. (1984). *Collective and Corporate Responsibility*. New York: Columbia University Press.
- Fricke, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford and New York: Oxford University Press.

- . (2016). What's the Point of Blame? A Paradigm-Based Explanation. *Nous* 50(1): 165–83.
DOI: 10.1111/nous.12067
- Garfinkel, S. (2000). *Database nation: The death of privacy in the 21st century*. Sebastopol, CA: O'Reilly Media.
- Gilbert, M. (1987). Modelling Collective Belief. *Synthese* 73(1): 185–204. DOI: 10.1007/BF00485446
- . (1989). *On Social Facts*. New York: Routledge.
- . (2004). Collective Epistemology. *Episteme* 1(2): 95–107. DOI: 10.3366/epi.2004.1.2.95
- Greene, A., and S. J. Gilbert. (ms.) More Data, More Power? Towards a Theory of Digital Legitimacy. Unpublished manuscript. Last revised: 6 July 2021.
SSRN: <https://ssrn.com/abstract=3773898>
- Hacking, I. (1995). The looping effects of human kinds. In D. Sperber, D. Premack, & A. J. Premack (Eds.), *Causal cognition: A multidisciplinary debate* (pp. 351–94). Clarendon Press/Oxford University Press.
- Hardin, R. (2002). *Trust and Trustworthiness*. New York, NY: Russell Sage Foundation.
- Hill, K. (2019). Week 6: Blocking them all. *Gizmodo*. <https://gizmodo.com/i-cut-the-big-five-tech-giants-from-my-life-it-was-hel-1831304194>
- Holton, R. (1994). Deciding to Trust, Coming to Believe. *Australasian Journal of Philosophy* 72(1): 63–76. DOI: 10.1080/00048409412345881.
- LaFrance, A. (2021, September 27). The largest autocracy on Earth: Facebook is acting like a hostile foreign power; it's time we treated it that way. *The Atlantic*.
<https://www.theatlantic.com/magazine/archive/2021/11/facebook-authoritarian-hostile-foreign-power/620168/>
- May, L. (1987). *The Morality of Groups*. Notre Dame: University of Notre Dame Press.
- . (1992). *Sharing Responsibility*. Chicago: University of Chicago Press.
- McCleod, C. (2021). Trust. *The Stanford Encyclopedia of Philosophy*. Fall 2021 Edition Edward N. Zalta (ed.). <https://plato.stanford.edu/archives/fall2021/entries/trust/>
- Mellena, G. (1997). *Collective Responsibility*. Amsterdam: Rodopi Press.
- Nguyen, C. T. (2021). How Twitter Gamifies Communication. In *Applied Epistemology*, ed. Jennifer Lackey. Oxford: Oxford University Press. pp. 410–36. DOI: 10.1093/oso/9780198833659.003.0017

- Pettit, P. (2007). Responsibility Incorporated. *Ethics* 117: 171-201.
- Rosentiel, T. (2010). United We Stand ... on Technology [Survey]. Pew Research Center.
<https://www.pewresearch.org/2010/05/05/united-we-stand-on-technology/>
- Runkevicius, D. (2020). How Amazon Quietly Powers The Internet. *Forbes*.
<https://www.forbes.com/sites/danrunkevicius/2020/09/03/how-amazon-quietly-powers-the-internet/?sh=751145733092>
- Sheridan, D. R. (1997). Zeran v. AOL and the Effect of Section 230 of the Communications Decency Act upon Liability for Defamation on the Internet. *Albany Law Review* 61(1), 147-179.
- Shoemaker, D. (2011). Attributability, Answerability, and Accountability: Toward a Wider Theory of Moral Responsibility. *Ethics* 121: 602-632.
- Strawson, P. F. (1962). Freedom and Resentment. *Proceedings of the British Academy* 48, 1-25.
- Thaler, R. H., and C. R. Sunstein. (2008). *Nudge: Improving Decisions About Health, Wealth, and Happiness*. New Haven, CT: Yale University Press.
- Tuomela, R. (1992). Group Beliefs. *Synthese* 91: 285-318.
- Véliz, C. (2021). *Privacy is Power: Why and How You Should Take Back Control of Your Data*. London: Bantam Press. Kobo e-book edition (no fixed page numbers).
- Walker, M. U. (2006). *Moral Repair: Reconstructing Moral Relations after Wrongdoing*. Cambridge, UK: Cambridge University Press.
- Wallach, O. (2020). How Big Tech Giants Make Their Billions. *Visual Capitalist*.
<https://www.visualcapitalist.com/how-big-tech-makes-their-billions-2020/>
- Wartenberg, T. E. (1992). "Situated Social Power." In *Rethinking Power* (pp. 79-101). Albany, NY: SUNY Press.
- Web Summit. (2019, November 4). Morality in the age of tech surveillance—Edward Snowden.
https://www.youtube.com/watch?v=X4_7A-SGLo8
- Webb, A. (2019). *The Big Nine: How the Tech Titans and Their Thinking Machines Could Warp Humanity*. New York: PublicAffairs.
- World Bank, The. (2020). GDP Ranking. *The World Bank Data Catalog*.
<https://datacatalog.worldbank.org/dataset/gdp-ranking>
- Williams, J. (2018). *Stand Out of Our Light: Freedom and Resistance in the Attention Economy*. Cambridge, UK: Cambridge University Press.

Wray, K. B. (2001). Collective Belief and Acceptance, *Synthese* 129(3): 319–333. DOI:
10.1023/A:1013148515033