

The Enterprise of Platform Governance Development

Transnational organizations and hegemonic donor states such as the United States have concerned themselves in recent years with building the state capacity of developing and transitional states. The United States made efforts to develop the military and police capacity of postwar Iraq (after tearing down its prior version of those institutions) (Gates 2010). The European Union has made efforts in Eastern Europe to build anti-corruption capacity (Börzel and van Hüllen 2014). There have been substantial global efforts to build state capacity in Somalia (Menkhaus 2014). At least to some extent, these efforts are motivated by the self-interest of those leading the capacity development efforts – the more functional Somalia’s state is, the more it can control the piracy that harms its neighbors. And it may be that a functional Somali state would be better at controlling piracy than external states trying to do so themselves, for example, because of its superior local knowledge or sociological legitimacy. In other words, the enterprise of capacity development is often predicated on the ideas that governments that are unable or unwilling to control their own citizens create *negative externalities* for the rest of the world, and that the rest of the world is unable to control those behaviors directly, without helping the state with direct access to the people causing the problem to build the capacity to intervene.

This chapter argues that, from the standpoint of government regulators, analogous incentives exist with respect to platforms. Governments have strong reasons to control much of the pathological behavior that appears on networked internet platforms, from commercial fraud on Amazon to fake news on Facebook (although the example of fake news also illustrates that government reasons are far from unequivocal and may be outweighed by other interests, such as individual freedom). And it is likely that governments may be ineffective at directly regulating some of this behavior, for example, because of the tendency of platforms and their users to span jurisdictional boundaries, because of the property rights of owners of those platforms, or simply because of the bureaucratic burden of regulation without direct access to the immediate site of behavior in the form of platform code and interactional surfaces. Such governments accordingly have reasons to promote the

development of a kind of analogy to state capacity in platforms in order that they may require or encourage those platforms to control behavior that appears via their infrastructure.

2.1 WHY IS CONDUCT ON PLATFORMS SO HARD TO REGULATE?

Some of the ways in which many commentators worry about externalities caused by platforms represent what we might think about as pure conflicts of interest between those platforms and society at large. Social media companies tend to profit by collecting large amounts of information about user activity and using that information for ad targeting. That is (arguably) bad for society as a whole, insofar as it allows advertisers to manipulate consumers and leads to the risk of dangerous third-party data leaks. But it's not bad for companies (the data leaks are – but the manipulation is precisely the service they offer to advertisers). Similarly, “sharing economy” transactional platforms tend to make money by evading regulations governing services like taxicabs and hotels as well as workplace law in general. Again, that's a problem for the rest of us. Under such circumstances, the best we can do as a society is use the legal system to simply coerce companies to act in ways that are less self-serving and more public-serving by, for example, requiring them to count their drivers and delivery people as employees rather than contractors or prohibiting them from sharing user data between different corporate functions.

By contrast, other kinds of external harms – the sorts within the domain of interest alignment described in the Introduction – likely also harm platform company interests, at least under certain interpretations of those interests. The clearest examples involve various kinds of fraudulent behavior which can harm companies' long-run interests in running platforms that their users can trust for reliable information. For example, viral disinformation on social media undermines the implicit value, at least over the long run, that those users can attribute to information found on those platforms. The same goes for counterfeit products and fraudulent ratings on transactional platforms.¹

With respect to such harms, it is a genuine puzzle (which this book aims to address) as to why the companies cannot simply stop them. This puzzle is highlighted by the

¹ To the reader who is skeptical about the claim that company and social interests are aligned in these cases: Wait a few paragraphs. I absolutely acknowledge that sometimes, company interests are more ambiguous or conflicted than may appear on the surface, and Chapter 4 addresses such cases at length. It may be that the harmful behavior under examination, for example, promotes short-term profit (by generating social media platform engagement or transactional platform sales) at the expense of long-term profit; that different constituencies within a company's employees and/or ownership differ as to their evaluation of the relative benefits and burdens of such activity; or that companies operate under external pressure (such as by government officials who benefit from social media disinformation) which effectively changes platform interests by giving them an incentive to permit the harmful behavior. But sometimes, as in the Myanmar case which opened this book, company interests are not conflicted, yet they still cannot stop the behavior.

fact that there are important senses in which technology platforms have vast governance advantages over physical world governments.

First, unlike physical world governments, company personnel can observe anything (though not everything at once) that happens under their control (with the sole exception of end-to-end encrypted platforms like WhatsApp – but even there, the encryption is itself a platform design choice under company control). They need not install surveillance cameras and peek through windows, they simply need to query a database. Such observation is not costless – it still requires paying personnel to actually examine some piece of on-platform behavior (a product for sale, a tweet), but the costs are substantially lower per observation than for states to send police to follow people around, get search warrants, and so forth. Moreover, some (but not all) “observations” in the platform context can be automated by, for example, deploying artificial intelligence to identify certain kinds of conduct. While artificial intelligence and other automated methods of observation are no panacea, they have certainly proven effective in a number of important cases. For example, companies have devised methods to automatically fingerprint and then identify duplicates of known child pornography images so that they may be removed from platforms without any human having to look at them.² It is also possible to have a kind of human-in-the-loop observation where artificial intelligence identifies behavior to be reviewed by a human. Governments can automate enforcement too (speed cameras exist), but governments still have to expend costs to translate behavior in the physical world into data, whereas companies have data all the way down.

Similarly, many forms of rule enforcement are, at least in the abstract, much cheaper for platform companies than for physical world governments. The most common tools of enforcement in the platform context are various forms of restriction or removal either of user accounts or of the specific conduct of user accounts (such as posted content). But, once again, these enforcement techniques do not require the building of prisons or the operation of courts, or the sending of heavily armed police officers to people’s houses – they simply require the design of software to create affordances for enforcers within companies to alter user privileges and content, and then the operation of those tools – the click of a button within some internal system to change a field in a database.

But the relative ease of governance of platforms is to some extent an illusion, for platform companies are still subject to two core abstract-level problems that they share with physical governance.

The first is the difference between observation and knowledge, namely interpretation or what I will describe in Chapter 3 as “legibility”: While platform companies can “see” any behavior on the platform by querying a database, interpreting that behavior, and in particular, determining whether that behavior is harmful or

² See Microsoft’s description of the PhotoDNA technology at www.microsoft.com/en-us/photodna.

harmless, compliant with company rules or noncompliant, is much more difficult. In many ways, this interpretation is likely to be more difficult for platform companies than physical world governments, for platform companies – particularly the most successful companies at the greatest scale – have to deal with a vast amount of behavioral diversity and novelty, and are much more socially and culturally distant from many of their users. A bunch of lawyers and engineers in Menlo Park are likely to have a much harder time understanding the behavior of people in, say, Luxembourg or Kenya than the local government of a municipality composed of representatives of the people they govern.

The second is the problem of incentives. Not every person within a company's decision-making structure may have the incentives necessary to actually govern problematic behavior. Personnel whose job is to keep government officials happy may be more vulnerable to external pressure by politicians who benefit from disinformation. Engineers and product designers on "growth" teams who are rewarded for increasing activity may have an incentive to build features that facilitate all activity, whether those are harmful or helpful to a company's long-term interests. Stock markets may give executives an incentive to permit behavior that promotes short-term profits at the expense of long-term profits.

Those two problems are, respectively, the subjects of Chapters 3 and 4 of this book – I contend that they are the core challenges that vex the enterprise of platform governance in the domain where company interests (properly understood, from a long-term perspective) and social interests are aligned.

2.2 WHY BUILD CAPACITY?

The global community has long wrestled with the problem of states that are unable or unwilling to regulate the harms that users of their territory inflict on the rest of the world. Piracy is a classic example: When a state fails to adequately patrol its territorial waters, it can become a haven for pirates who then prey on international shipping (e.g., Pham 2010). But there are other kinds of criminal activity that can result from governance failure (or governance unwillingness) and can predominantly affect people outside a given country, ranging from terrorism to intellectual property misuse and the cottage industry of internet scams in certain towns in Eastern Europe (Subramanian 2017; Bhattacharjee 2011).

This problem may arise because a country lacks sufficient ability to fully regulate behavior on its territory. Maybe, for example, the government lacks legitimacy with the population, or the funding to train and deploy adequate security forces. Moreover, it may lack adequate incentive, either simply in virtue of the fact that much of the cost of the harm is externalized or because leaders have mixed motives, as when the criminals pay bribes to continue their activity. The piracy example is again apt: Percy and Shortland (2013) have plausibly argued that even improvements in state capacity are ineffective in controlling some cases of modern

piracy in view of the fact that the activity provides local economic benefits and its harms are externalized.³

Nonetheless, it may be difficult for other states to directly regulate this behavior, even though it inflicts harm on their citizens or in their territories. They may not, for the simplest example of a reason, have their own physical access to the territory where the international harms are being generated. And they may lack sociological legitimacy with the people of that territory, even more so than an unpopular domestic government – it should be unsurprising, for example, that the people of Somalia would be unwilling to be governed by the United States and France.

Accordingly, a core strategy that other states have adopted to address the problem of externalities generated by under-governance is to attempt to assist the state which is failing to govern in building or rebuilding its capacity to do so, and providing it with incentives to use that capacity. On our running example, Bueger et al. (2020) describe maritime capacity-building efforts in various nations in the region to address piracy from Somalia. Similar efforts have been made in order to improve the ability of countries to fight international terrorism within their territories (e.g., Bachmann and Hönke 2010).

Observe the incentive structure underneath such capacity-building enterprises. Under-governing states create certain kinds of external harms in view of the fact that they fail to control the behavior of people operating from their territory. For convenience, I label those kinds of harms *governance failure externalities*. Moreover, such states not only fail to control harmful behavior, there's an important sense in which they provide resources with which individuals may carry out that behavior, and hence provide an incentive for that behavior. The most obvious example of such resources is a headquarters – by allowing (deliberately or inadvertently) pirates or terrorists or international fraudsters or whomever to operate from their territory, such states shield them from the direct control of others – Westphalian notions of sovereignty impose a barrier to more effective states directly acting to control the harmful behavior. Another related example of a resource that failed states provide is access to victims – by failing to control an important region of the ocean located among an important trade route, a state effectively provides pirates with a known location in which victims may be found. Regardless, the point is the externalities: other states, and international organizations, have self-interested reasons to provide capacity-building assistance and incentives to lower the cost of the home state regulating domestic behavior with negative spillover effects.

³ A more sophisticated description of the problem that captures this point may be that strong states which do not have a problem with piracy have little incentive to permit it, as they can prosper from less harmful forms of economic activity; however, once pirates get a foothold in a weak state, it may be that when such states become stronger, they still lack adequate incentive to expend the costs to shift to more nonviolent and less harmful forms of revenue generation.

I do not wish to describe this problem in a way that buys into problematic US-centric or developed country-centric representations of global problems and responsibility. Wealthier and more powerful countries can also impose governance failure externalities on the rest of the world. For the most obvious example, the failure of the United States to control the carbon emissions emanating from its territory imposes vast harms on the rest of the world. In terms of long-run global welfare, greenhouse gases are far more damaging than pirates. In principle, it would make sense for the rest of the world to help us Americans control our polluters. Unfortunately, given the vast power and wealth disparity between the United States and most other nations, our troubling belief in American exceptionalism, and the fact that our governance failures are built deep into our polarized politics and ineffective constitutional institutions, there's probably precious little that other countries could do to help us.⁴

Platforms create precisely the same kinds of governance failure externalities. Here, some care is in order: a corollary of the point which I have repeatedly emphasized that only some platform harms are within the domain of interest alignment with society is that not every harm that platforms impose on the outside world is a governance failure externality. Sometimes, the platforms themselves do things that directly harm third parties – such as when they violate people's privacy or squash their competitors. But other times, the externalities come from the independent behavior of users making use of platform affordances.⁵ In the platform context, the most obvious examples revolve around various kinds of fraud or confusion – when viral disinformation circulates on social media, for example, or when counterfeit N95s plague Amazon. And much like in the international context, the affordances of platforms both facilitate this behavior and make it harder for governments to control (primarily by increasing its quantity and speed, as well as cross-border impact).

⁴ Moreover, the politics of the United States are also creating externalities even in the domain of social media: terroristic conspiracy group QAnon, which arose in the United States (perhaps with Russian assistance) and spread in social media on behalf of Donald Trump (Bleakley 2021), has spread to other countries, notably Germany (Hoseini et al. 2021), where it was linked to a coup plot in late 2022 (Gilbert 2022).

⁵ As noted in the virality section of Chapter 1, there's an important sense in which *everything* platforms do comes from the platforms, insofar as they control the causal mechanisms that permit third-party behavior to have an impact on others, such as the distribution of content over social media. But, taking as given an economy of platforms in which user behavior is conferred some (platform-mediated) causal impact on the experience of others, we can coherently talk about categories of harms where it makes sense to attribute primary responsibility for those harms to users. For example, suppose that some method of content distribution has many socially valuable uses, such that it would sacrifice significant public well-being to ask a social media platform to disable it, but that method of content distribution also has some small number of seriously harmful uses. Under such circumstances, it seems to me most plausible to point to interventions on user behavior to end the harmful uses as the best path forward for overall social well-being, and hence to attribute responsibility for harms generated in the absence of such interventions primarily to user behavior.

More specifically, the relative inability of governance “donor” states to act directly on the people of recipient states is replicated in the platform context. As a whole, companies tend to have both access to more information about user behaviors and access to a wider array of tools (such as major interventions on their own platform design) to regulate them than do governments. By contrast, even though ordinary platform users would presumably prefer to be able to use their market power to mitigate some of the harms they receive from platform companies (especially if that sometimes means, e.g., accidentally buying fake N95s in a pandemic, or being deceived by political lies), governments have much more leverage over platform companies than the users of those platforms do, in virtue of the fact that governments, of course, have more coercive tools available (they can freeze assets and lock up executives). Users have very little power over platforms themselves (their only viable option is exit, network externalities lead to very high switching costs for individuals, and mass coordination is costly).

This suggests that a regulatory capacity-building approach dedicated to the goal of empowering companies to protect the interests of users and society – which this book argues will also require empowering ordinary people to intervene on platform choices – is an appropriate way to think about how we, as citizens of democracies wishing to control our environments, ought to bring ourselves, as users of these platforms, under control. Of course, there are dangers as well to the project of capacity building – which I will address below – most obviously the rightful fears that many of us have about enhancing the power of unaccountable private platform operators over behavior in a context where increasing shares of important human interaction happen via those platforms. But only when we enable ourselves to think directly about governance capacity do we even become capable of properly raising those questions.

2.3 HOW TO BUILD CAPACITY?

In the abstract, it is helpful to think about the enterprise of capacity building as giving a putative governor (as a first pass, a platform company – but the remainder of the book will call into question the notion that it is *companies* in particular who should be doing the governing, as opposed to a combination of companies, workers within companies, civil society, perhaps governments, and most importantly ordinary people) both the *incentive* and the *tools* to control unwanted behavior.

In the context of states, relevant examples of such capacity-building efforts include the enterprises of rule-of-law development and of security assistance. Powerful states and their associated NGOs have offered less powerful states a variety of tools for more effectively extending their authority over their territories, including military training and arms, legal personnel training, and building of physical infrastructure for legal administration such as courthouses; of particular relevance for the argument to come, some of these resources (such as training) have also been extended

to civil society actors within recipient countries rather than governments directly, with some success in achieving the goals of legal development (Blattman, Hartman, and Blair 2014). Rule-of-law development efforts also have intervened directly on state incentives, such as with constitution-drafting projects; additionally (and rather more controversially) economic development programs directed by entities such as the World Bank and the International Monetary Fund have also aimed to intervene on state incentives with “structural adjustment” programs which impose obligations on states to do things like privatize public services as a condition for receiving international loans.

Governments have analogous tools with respect to companies. The United States or Europe could, for example, impose changes on the corporate structure of platform companies or on the ways that decisions are made within that structure by alterations to corporate law. In other corporate contexts, we have seen examples such as the Sarbanes-Oxley Act, enacted after the Enron scandal, which imposed new requirements on the boards of directors and outside auditors of public companies. We can understand this as directly analogous to (but hopefully more coercive than) the constitution drafting assistance offered by the world’s liberal democracies in places such as Afghanistan. Foreshadowing the proposals articulated in Chapter 6 as well as this book’s Conclusion, governments could also directly insert third parties into the governing structures of companies, for example, by mandating certain kinds of representation on corporate boards, by giving third parties a veto on some company decisions, or even by directly acquiring company stock and exercising decision rights on behalf of empowered third parties. This too is a governance tool with which governments are familiar. For example, Article 27 of the European Charter of Fundamental Rights provides that workers have a right to “information and consultation.” As implemented by an EU directive, workers are entitled to that “information and consultation” on, among other things, “decisions likely to lead to substantial changes in work organisation.”⁶ Member states have further implemented this right by providing for worker consultative bodies within companies.⁷

While there are limits, such as the American constitutional protections for property (and, in the case of social media companies, free speech), the very brief survey above should suffice to illustrate that governments have the tools to intervene quite deeply into the decision-making processes of platform companies. In view of the significant (and legitimate) coercive power that countries exercise over those

⁶ Directive 2002/14/EC of the European Parliament and of the Council of 11 March 2002 establishing a general framework for informing and consulting employees in the European Community – Joint declaration of the European Parliament, the Council and the Commission on employee representation, Article 4(2)(c).

⁷ See, for example, French Code du travail, Articles L2311-1 to L2311-2, which sets up a social and economic committee (“comité social et économique”) in companies above a certain size to represent employees; that committee has a broad array of powers including receiving substantial amounts of information from the company, raising complaints to board level, investigating safety issues, and advancing employee concerns to government officials.

companies, we also have reason to think that the use of those powers for purposes of capacity development has greater prospects of success than capacity-building projects in the context of states.

This book concerns itself with the enterprise of governance in the abstract, that is, with the kinds of institutions (organizational forms) that ought to be adopted in order to permit platforms to effectively govern (and permit us to effectively govern them). In the context of the analogy to governance development in the international community, this is analogous to saying that the scope of concern should be the promotion of institutions like democracy and the rule of law, rather than the prescription of the particular substantive legal and economic rules which recipient countries ought to adopt.

The institutional approach recognizes that the concrete rules that platforms – or even platforms in specific sectors, such as social media – ought to adopt will potentially vary widely across platforms, across underlying social contexts, and across time (just as with countries). Moreover, many of the real-world rules problems facing platforms are sufficiently contested and potentially irresolvably difficult in their form as rules, even if there are correct answers to them. What, for example, is the normatively correct position for a platform to take on disputed claims about treatments for COVID-19, or allegations about the behavior of a US president and their family, for example? It seems to me overwhelmingly likely that the best answer to this category of problem will operate in a combination of the domains of what Rawls (1999, 73–78) called “imperfect procedural justice” and “pure procedural justice.” That is, that the answer to the question will focus on how the decision is to be made (and, by extension, who is to make it), rather than what the decision is to be, sometimes because the procedure is our best albeit imperfect way to get to an externally defined right answer, and sometimes because whatever answer follows from the procedure (as with democratic policy choices) is for that reason the right one.

Another component of the institutional approach adopted in this book is the supposition that the concrete interventions that are likely to be effective in platform governance are more in the domain of product design in terms of the high-level affordances available to users as well as in tools and methods of governance, rather than in low-level enforcement details or volumes of work carried out (such as the number of person-hours available to do content moderation). Sahar Massachi, a former Facebook software engineer in Civic Integrity and one of the founders of the Integrity Institute, sometimes refers to this as the difference between adding more traffic cops and adding speed bumps to the roads.⁸

With respect to social media platforms, free speech ideals also place important constraints on the scope of potential government regulation. While few countries have free speech traditions quite as strong as those of the United States, the value

⁸ Mostly in private conversation among those involved in the Institute, including myself; a less pithy presentation of the general idea is in Massachi (2021).

is recognized in every liberal democracy and by international human rights law. Under US law, it would be impossible to directly restrain many of the pathological speech phenomena on social media (or to force the companies to do so), such as by banning the spread of conspiracy theories and so-called “fake news” or racist speech. Countries with hate speech regulations and the like, of course, might be able to go further, but even there – as with most efforts to regulate widely dispersed behavior via law – the practical difficulty of such regulations and the cost of enforcement is widely recognized. The same point holds with respect to transactional platforms. If, for example, states and municipalities really wanted to directly regulate every person who offered up a room on Airbnb, well, they could do so, but the effectiveness and practical scope of those regulations would be limited by the sheer number of people involved and the difficulty of the government directly monitoring so much behavior.

By contrast, regulations of the decisional infrastructure of technology companies themselves, as well as their design, and the affordances they offer – as the institutional approach recommends – are somewhat more insulated from free speech challenges insofar as they do not directly target expressive activity and could also apply to non-expressive activity. To be sure, regulations of companies are not without limits. In the United States, it would probably be unconstitutional to impose something like a “fairness doctrine” directly on social media companies, requiring them to engage in politically neutral content moderation (the misguided views of the State of Texas and the Fifth Circuit in a recent bill and case notwithstanding). It would also be unconstitutional, although somewhat more difficult to prove in a specific case, to impose regulatory disadvantage on them in ways motivated by political viewpoint, such as various Republican proposals to abolish Section 230 of the Communications Decency Act in retaliation for companies’ interpretation of their rules to remove conservative speech from their platforms. But other sorts of institutional interventions are free from such constitutional barriers. For example, many interventions on the workplace structure of companies, such as a proposal given in the Conclusion to this book to give more job-related power to the offshored contractors currently serving as content moderators in social media companies, would not be directly related to speech at all. Also not directly related to free speech would be the use of antitrust law to break up companies with excessive market power (discussed later in this chapter). Similarly, government efforts to assist the voluntary building of user and public-dominated governance institutions for companies through industry-wide action like that proposed in Chapter 6 would be unlikely, because of their voluntary character, to run afoul of companies’ First Amendment rights; for similar reasons, the interposition of an uncoerced choice by companies to build such institutions (albeit with the government’s help) would at least arguably insulate the government’s role from challenges on the basis of the First Amendment rights of users as well.

2.4 OBJECTION 1: IS PRIVATE GOVERNANCE AN OXYMORON OR A DANGER?

Before we delve into the account of how we might promote the effective governance of users by platforms, some interrelated preliminary objections must be addressed. The first identifies a kind of teleological mismatch between the goals of governance development and the ends of private, profit-motivated, entities. Even in the ideal form of the so-called “new governance” models under which the state recruits or collaborates with private entities to govern the public, there’s an important sense in which the state remains in the driver’s seat as much as possible – for example, by contracting with private entities to carry out governance tasks, recruiting the support of private entities for pieces of an overall project of infrastructure or community-building, or convening private for-profit entities to participate in their own regulation rather than the regulation of third parties.⁹ By contrast, the most plausible ways to develop private platform governance involve platforms regulating their users, and, in view of the international nature of the major platforms, it’s unlikely that governments will be able to exercise close supervision over the conduct of that governance. Indeed, close supervision by governments is almost certainly undesirable, due to the oppressive nature of some governments and the hegemonic nature of others – issues to be explored further below.

We have strong reasons to be wary of the delegation of governing responsibilities from states to private companies. At least in theory, states are oriented toward (and accountable to) the good of their people, whereas companies are oriented toward the good of their shareholders. In law, for example, the whole enterprise of “consumer protection” reflects a recognition that giving companies the power to dictate the behavior of people who engage in non-shareholder economic relationships with them (e.g., with excessively empowering contractual terms) makes individuals vulnerable to exploitation.

Even when states, rather than companies, have been the beneficiaries of capacity-building efforts, those efforts have often led to unintended consequences relating to the abuse of that power. For example, antiterror capacity-building programs in Kenya have also built the government’s capacity to engage in repressive actions against local minorities (Bachmann and Hönke 2010). Similarly, with respect to companies, we have good reason to fear that increasing their control over the conduct of users on their platforms could facilitate any number of inappropriate choices

⁹ Bell and Hindmoor (2012) give an extended argument for why these sorts of enterprises still depend on the notion of the state and its authority at the core. Arguably, Bell and Hindmoor’s argument could apply to the program of platform governance development as well – we could simply understand it as a form of “metagovernance” in which states govern more effectively by arranging for some helpers. But in view of the fact that I explicitly propose institutions meant to dodge state control in a number of areas, particularly where authoritarian (or shakily democratic) states aim to turn platforms into cat’s paws to do things like political censorship, I don’t think I can avail myself of their argument.

made by the companies over the years. For perhaps the most important example: Every platform company is widely criticized for violating the privacy of users; many scholars have suggested that such violations in the form of surveillance are the core business model of companies (e.g., Zuboff 2019). But by asking companies to do more governance we might be exacerbating their surveillance by giving it a kind of normative cover. After all, the same kinds of practices that offend our notions of privacy and raise the specter of surveillance might also be useful in effective governance. Consider database merging and user profiling across platforms – many of us object to, say, Facebook’s using its infamous “pixel” to track our behavior across non-Meta-operated websites and even allegedly build profiles of users independent of their Meta-affiliated accounts (so-called “shadow profiles”) (Garcia 2017; Aguiar et al. 2022; Bekos et al. 2022). But what if that information is used to identify “inauthentic” users, such as Russians attempting to subvert elections?¹⁰ Suddenly, it seems perhaps a bit more acceptable, and perhaps even something we might demand of companies. Looking down the game tree, we may resist conferring additional governance resources (and responsibilities) on companies in order to avoid giving them further incentives and tools to conduct such surveillance.

This problem of surveillance – indeed, the general danger that capacity improvements for governance purposes (i.e., where company interests align with the public interest) may turn into capacity improvements for exploitative purposes (where company interests conflict with the public interest) – is real and serious.

However, I contend that we can avoid objections that point to the above-described problems by attending more carefully to the ways that (a) states and companies are subject to similar incentives (so state regulation of users directly, even if possible, might not be better than governance development); and (b) we can build capacity improvements for platform governance that do not necessarily empower companies as such – but rather empower other ecosystem actors with incentives aligned both with the public and with companies on governance issues – including users, non-user citizens, civil society, workers and other important constituencies.

To begin, it is notable that many of our most useful accounts of state development suppose that leaders have self-interested motives, and, historically, it is difficult in the extreme to suppose that the nascent stages in the early development of the high-capacity nation-state were purely or even mostly public oriented. Indeed, a substantial part of the point of the literature on the development of governance

¹⁰ At a Congressional hearing, Mark Zuckerberg denied any knowledge of “shadow profiles” but did claim that there were safety-related reasons to collect data on nonusers (Molitorisz, Meese, and Hagedorn 2021, 51), so this example is not hypothetical – although apparently the justification Zuckerberg gave was more focused on preventing third-party profile scraping (Hatmaker 2018), and hence on protecting privacy (I confess to some confusion about how collecting data on nonusers protects against privacy invasions by scrapers as such). More generally, Meta does use personnel and techniques from domains such as intelligence and criminal investigation to track down Russian subversion (Stamos 2018; Murphy 2019). So does Amazon (Amazon.com 2020).

capacity and the state itself is the design or evolution of institutions that can re-align the incentives of top-level and mid-level leaders to cause them to pursue the public good in their own private interest, or the identification of private interest based causal frameworks for the emergence of such institutions (e.g., Olson 1993; De Lara, Greif, and Jha 2008; North, Weingast, and Wallis 2009; sometimes this can include the development of private legal frameworks as well, e.g., L. Liu and Weingast 2018).

Unsurprisingly, the worries which I articulated above also appear frequently in the context of states. Take the problem of surveillance. It's useful to help states prevent crime, terrorism, subversion, and so forth. And it's incredibly dangerous to enable states to conduct political repression. I'd go so far as to say that the problem of surveillance in the context of states is the most classic and intractable example of the persistent tension between security and freedom. Shorter version: J. Edgar Hoover existed. Assuming that someone has to watch what people say on social media to control the spread of things like viral misinformation, there is little reason to think that the state is any less likely to abuse the power to do so than companies.

Moreover, it seems to me that the key problem with the objection to private governance is that it assumes a too-thick distinction between "public" and "private" which is, to some extent, merely a quirk of our place in history. Rather, we ought to understand that the "public," both in the sense of governments and in the sense of other ways of noncommercially organizing people (e.g., via civil society, social movements, and even purely artifactual organizations), can be integrated into the decision-making processes of private companies in ways that are *both* private and public – private because the companies remain for-profit businesses with private owners, whose decisions are in substantial part motivated by profit – yet public because those decisions can both be controlled in the public interest and improved in the private interest by integrating noneconomic forms of human organization, and hence human incentives, into them.

Critiques of private governance can also merge into critiques of the undemocratic nature of such forms of governance – and hence the implicit claim that what counts as democratic is limited by existing forms of public authority. As an example, consider two recent papers, one by Blayne Haggart and one by Haggart and Clara Iglesias Keller (Haggart 2020; Haggart and Keller 2021). The two papers criticize several major platform governance initiatives such as the Meta Oversight board and proposals by several governments, as well as leading scholarly proposals, on the basis of their lack of democratic legitimacy.¹¹ In particular, they focus on a conception of

¹¹ As so frequently happens in papers about "legitimacy," the authors seem to run together legitimacy in a normative sense, that is, decisions that ought to be respected because the right people made them, and legitimacy in a sociological sense, that is, decisions that will in fact be respected by relevant populations. For present purposes, I will assume that the core of their argument is about normative legitimacy, as that is what matters, morally speaking. (Sociological legitimacy might have moral importance as well, but only in the context of an argument according to which sociological legitimacy is required for normative legitimacy, or for some other important moral end. Otherwise, we're simply trying to get an "ought" from an "is.")

democratic deficit framed around what the Haggart/Keller paper describes as their lack of “input legitimacy,” that is, the failure of major decisions to go through democratically accountable institutions.

There is force to their objections. For example, they object to the design of the Meta Oversight Board in virtue of the fact that its decisions ultimately are still in the control of the company, both because the company makes the rules and because the company sets the terms of civil society and other participation (e.g., by choosing the initial chairs). And, they criticize David Kaye’s (2019b) human rights-focused approach for being overly concerned with protecting companies and their users against censorship efforts by governments, and failing to distinguish between democratic and authoritarian governments and the differing degrees of trust which ought to be extended to them in controlling online speech.

Yet the vision of those papers is limited by their rigid conception of the shapes democratic institutions might take. Both papers assume a fairly cramped Westphalian statist notion of democracy, in which it is impossible for a decision-making process to be legitimate unless it is associated with a specific, bounded, *demos*.¹² But if we value democracy because of the normative importance of individual and collective autonomy and the practical advantages of superior decision-making, I contend that those values can be achieved outside the context of the nation-state (back to the Deweyian point about how democratic institutions need to proceed from an understanding of the relevant underlying public and, well, its problems). Similarly, I think Kaye has the better of the argument with Haggart and Keller with respect to the danger of excessive control over platform speech even by nominally democratic states – we need only consider the threatened abuses of government power in weak democracies run by demagogues such as the United States during the Trump regime or India during the Modi regime to illustrate the risks involved. More generally, even robust democracies have long-term minorities, indigenous peoples, and other groups of people identified with, but not fully represented by, the nation-state that happens to rule the physical territory. For example, Spain is widely seen as a stable and legitimate democracy, but it is far from obvious that it can be understood to fairly speak for its Basque and Catalan citizens; how much less can the United States be understood to fairly speak for its Native American citizens.

It seems to me that the goal of achieving “input legitimacy” is a worthy one, but that we ought to focus our attention on ways to do so that include, but are not limited

¹² This is a consequence of the intellectual frameworks within which they operate: the Haggart paper relies on Dani Rodrik’s account of the difficulty of democratic “national self-determination” in the global context, and the Haggart/Keller paper relies on Vivien Schmidt’s claim that “input legitimacy” requires a kind of thick national polity with a discrete identity to which decision-making processes are accountable. Proceeding from those starting points, the arguments in those papers make perfect sense. But those starting points are controversial. (I happen to think they’re outright false, but can’t defend that claim in this book except insofar as some of the discussions of platform identity and constitutional patriotism in Chapter 5 might constitute a defense.)

to, the nation-state. Rather, this book will argue for platform governance reforms that can genuinely incorporate both democratic peoples organized into such states *and* national minorities, global civil society and social movements, and other forms of meaningful human organization not encompassed by the modern state. I will further argue that such reforms can actually make platform governance more effective at governance – that is, at writing rules that work and can be effectively enforced to control harmful externalities. The state, as well as platform companies, must drive such reforms, because it is states (and multi-state organizations like the EU) and companies that currently have the power to do so. But those empowered by the reforms cannot be limited to states and companies.

A better approach is the one described by Hannah Bloch-Wehba (2019, 67–71) in what I believe to be the closest affinity in the existing literature to this book’s governance development approach. While Bloch-Wehba does not argue for capacity building as such, she does argue for the building of global democratic institutions to govern platforms as a form of what we might call legitimacy building, to contrast with capacity building. By doing so, she recognizes in a way that Haggart and Keller do not, that novel contexts for the exercise of power call for novel forms of democratization that are not wedded to pre-existing institutional forms.

The objection to private governance as raising a kind of teleological conflict between the interests of private companies and the public interest represented by states also ignores the scope limitation given at the beginning of this book. The scope of governance development as I propose is limited to the category of problems mentioned in the Introduction, where there is a *minimal alignment of interest* between platform companies and the rest of us. “Minimal” stands in for the recognition that platform companies might have complicated and conflicting interests, especially when considering the difference between short-term and long-term perspectives. “The rest of us” stands in, more or less, for the people of reasonably liberal and democratic states with interests in socially healthy behavior, an avoidance of fraud and malicious political subversion, and the like. The notion of minimal alignment of interests is also meant to recognize that sometimes the role of public policy is in part to give companies an incentive to follow the side of their interest which they share with the rest of us. Consider the problem of viral political misinformation: This is clearly against the public interest (poisoning democratic debate with lies is bad for everyone except the malicious actors spreading them). And it is *partly* against company interests – becoming known as a place where your quirky uncle Terry learns to be afraid of vaccines is not good for the long-term health of companies (or uncle Terry), as both users and advertisers are likely to be driven away. However, it is also partly consistent with their interests, understood from a different (more short-term) perspective: Viral misinformation is still activity, and it’s activity that some users vigorously consume; in an economic model that monetizes engagement, toxic engagement still counts. Governance reforms might make a difference in those kinds of problems by giving companies the incentive and the ability to focus

on the long-run benefits of getting rid of the viral misinformation rather than the short-run benefits of allowing it.¹³

By contrast, the problems where such alignment of interest does not exist are not amenable to governance reform. Consider the surveillance problem again: critics of the platforms represent it, perhaps rightly, as a pure clash of interests between companies and the rest of us. They make more money by doing and facilitating more surveillance; we enjoy more privacy, freedom, and power over our own lives by being subject to less of it. These problems are serious and critical, and we need to find a way to address them – but they are not the subject of this book, except insofar as they are indirectly *implicated* by this book, for example, because we need to craft our governance reforms with an eye toward not making the conflict of interest between platform companies and the rest of us worse.

That being said, in an ideal world, well-crafted governance development might also change platform incentives in the right direction even with respect to those situations where there is a genuine clash of interests between companies and the rest of us. In particular, suppose I'm right that more effective governance – over issues where companies' interests are aligned with the rest of us – requires giving ordinary people more power within and over companies (as I argue in the rest of the book). Those same reforms might also give the rest of us enough influence to at least partly nudge company incentives over a little bit in places where they are *not* aligned. More bluntly, once we democratize the platforms, we might discover that our new democratic institutions can help them find revenue models that are less dependent on surveillance, the exploitation of workers, the evasion of local regulation, and all the rest of their abusive behaviors.

2.5 OBJECTION 2: DOES PLATFORM CAPACITY BUILDING OVER-EMPOWER THE ALREADY POWERFUL? IS IT COLONIAL? IMPERIAL?

Here in the United States, our problems tend to dominate not only our own but the world's minds. Because of our economic power, our ideologies also tend to dominate platform governance. Twitter's executives used to call it “the free speech wing of the free speech party,” in line with America's position on the global extremes of free speech and the fact that, at the time, its leadership was dominated by famously libertarian Silicon-Valley types. On the other end of the libertarian spectrum, the same pattern holds: America's “leadership” (if we can call it that) in hyper-aggressive intellectual property regulation and the power of its media companies have exercised an outsized influence on global intellectual property law (e.g., Kaminski 2014; Yu 2006), doubtless including its enforcement by platform companies.

¹³ At a sufficiently – and uselessly – high level of abstraction, we might say that the basic nature of any governance reform is to decrease the extent to which the ones doing the governing discount the future relative to the present.

To a limited extent, Europe has also been able to propagate its values via platforms. Because of the size and wealth of its market, a number of aggressive EU internet regulations have exercised an international influence – the General Data Protection Regulation being the most prominent, but the Digital Markets Act and the Digital Services Act have also drawn particular attention, as have a number of proposed or implemented intellectual property regulations.

Accordingly, any interventions that might make these firms more effective at governing platform behavior potentially promote the further propagation of US and European values to people and places that may not want them. US free speech ideals might, for example, be applied in regulating conduct on social media in countries with much less libertarian norms; US and European intellectual property regulations might be used to censor the speech of users from countries and cultures with a very different attitude toward things like authorship and copying.

At the limit, Americans may even be making extremely contentious political judgments about the legitimacy of the leaders of other states. At least when Donald Trump was kicked off Twitter and Facebook, the executives making the decision were Americans, in American companies. But what happens if, say, Narendra Modi or Jair Bolsonaro were to incite a similar attack on their national legislatures after losing an election? Would we be as comfortable with the likes of Dorsey and Zuckerberg (or, heaven help us, Musk) making the same call?

Consider as a concrete example of the difficulty of these determinations the Thai laws against insulting the monarchy. There's a legitimate debate to be had about whether those laws – restrictions on political speech, to be sure, but restrictions that are associated with a longstanding political tradition and that are, at least on their face, fairly narrow – are consistent with human rights principles. Platform companies find themselves thrust into such debates when, for example, the Thai government asks YouTube to take down videos committing *lèse-majesté* (DeNardis 2014, 218; Klonick 2018, 1623).¹⁴

Exacerbating this risk of colonialism is the notorious problem of cultural competence experienced by many platforms, which frequently lack political, social, and even linguistic knowledge sufficient to adequately interpret behavior or model the values of foreign cultures. Perhaps the most famous example is Facebook's habit of taking down pictures of indigenous persons with culturally appropriate bared breasts (Alexander 2016), evidently on the basis of having misinterpreted the expression as pornography.

As with so many issues in internet governance, the problem of colonialism has not begun with the phenomenon of platforms. Rather, even governance strategies for the basic technical infrastructure of the Internet, such as technical standard-setting

¹⁴ However, this example may be inapposite, insofar as the Thai *lèse-majesté* law is used for political repression, as Streckfuss's (1995) account suggests. On the history and social function of the law more generally, see Mérieau (2019, 2021). It is notable that Thai commentators have compared the law to the laws against insulting the Prophet in Muslim countries (Mérieau 2019, 57), a similar area of cross-cultural conflict about the function of free speech in the context of sacred or sacralized figures.

processes for things like IP addresses, WiFi, and the like, can easily exhibit features that could be described as having at least a semi-colonial character. DeNardis (2009, 201–5) describes how such standard-setting processes can become dominated by developed country politics, and, in doing so, embed economic advantages to developed countries, such as a preference for incumbent intellectual property holders in those countries rather than new entrepreneurs in developing countries.

Moreover, there is a distinctive historical fidelity to the worry about colonialism in private governance. A major driver of European colonialism as it actually occurred was the empowerment of companies like the Dutch East India Company, in a process that historical sociologist Julia Adams (1994, 329–34) described as “patrimonial” in virtue of the way that the Dutch state, under the influence of powerful merchant elites, granted the company legal and financial privileges and sovereign authority while structuring its operations in order to reinforce their own institutional power. Adams (1994, 335–36) argues that the organizational structure of the company facilitated colonial enterprises of military and economic domination up to and including genocide by providing a mechanism through which the state’s repressive goals could be laundered.

The enterprise of rule-of-law development too has been criticized as colonial – including in my own prior work (Gowder 2016, 170–71; see also Mattei and Nader 2008). Rule-of-law development projects often appear to rest on the assumption that developing nations lack legal systems of their own – because they don’t look exactly like the western ones – and attempt to impose the values and institutions of the countries – mostly the United States and Western Europe – doing the promotion on the nations in which such projects are carried out. Certain kinds of capacity-building assistance may be downright harmful to recipient countries, perhaps even imposed on them coercively. For example, the enterprise of capacity building has been deployed in order to “support” developing countries engaging in more vigorous intellectual property rights enforcement (May 2004) – even though local economies may benefit from a lower degree of intellectual property protection than the international norm, and even though the international norm might be understood as being imposed by the greater trading power of powerful countries like the United States. My suggestion that platform governance regulation amounts to a kind of capacity development similar to rule-of-law development in the state context might seem to invite similar worries.

The conduct of major platforms has already been criticized as colonial. For example, Kwet (2019, 12) criticizes Facebook’s “Free Basics” internet service, which provides network connectivity for free in developing countries as a way to acquire market dominance via preferring the company’s own services, in view of the way such market positions enable such companies to exercise control over communications in those countries and thereby, in Kwet’s (2019, 12) words, “undermines various forms of local governance.” Similarly, Cohen (2019, 59–61) interprets platform extraction of data from and surveillance of the global South – sometimes with the collaboration of local governments – through a colonial lens.

Moreover, other scholars have identified the dominance of US platforms with the concept of imperialism. As I read Dal Yong Jin's (2013, 2017) influential work on the subject, platform imperialism is primarily economic in nature, and revolves around the way that platforms leverage intellectual property law as well as the commodification of information created by as well as about their users to extract wealth from the global South. In this sense, it can be distinguished from colonialism, which I mean to refer to political and social rather than economic domination – that is, the undermining of the freedom of colonized peoples to determine their own fate. However, the two phenomena obviously tend to go together – the propagation of American (and European, etc.) cultural influence increases both American economic influence and American social/political influence, and the two sorts of influence can be self-reinforcing.

This imperial quality can also be a consequence simply of the use of platforms as a kind of standard due to their positive network externalities. As David Grewal (2008) has cogently argued, positive network externalities constitute a form of power in that they degrade the usefulness and availability of alternatives, and hence deprive people of access to those alternatives. Even if nobody is in control of the object of one of these kinds of positive network externality coordination goods (or “standards”), such as a language or a currency being used as a lingua franca, Grewal's work shows that nonetheless, they have the potential to suppress diversity and local autonomy. So, for example, the network externality-derived usefulness of Facebook or Google in the context of a global economy may impede the development of social media or search engines that are more responsive to local needs.

Another form of colonialism or neocolonialism (quasi-colonialism?) in the platform business revolves around their labor practices. In particular, Roberts (2019, 182–200) describes how the commercial content moderation enterprise is embedded in colonial relationships, as offshored content moderators in the Philippines work in their (exploitative) jobs under conditions both of colonial cultural influence (this work is done in the Philippines partly because American colonial legacies have left workers there with a high degree of cultural and linguistic competence relative to American culture) as well as company-driven development paths in things like infrastructure choices.¹⁵

Objections to private governance and objections to colonialism potentially come together, as one tool with which developed nations have objectionably dominated developing nations has been the integration of nonstate organizations into government. Even leaving aside historical examples such as the various East India companies, contemporary governance capacity-building efforts have been criticized for effectively replacing functions of local states with private organizations that appear

¹⁵ One of the key recommendations of this book is more substantially integrating such workers in company decision-making, a recommendation which may directly alleviate this particular colonial aspect of those relationships.

to answer more readily to the donor states than to the domestic governments they are supposed to serve (e.g., Mkandawire 2002, 16–18).

In a world in which every government was legitimate and just, some potential solutions would suggest themselves, primarily involving platform rules being modified and applied in accordance with the requests of local governments. But, of course, this is not the case, and, in fact, the quasi-colonial character of platform governance is actually valuable to many living under authoritarian regimes, particularly with respect to social media. For example, one of the stories we might tell about the role of social media in the Arab Spring is that revolutionaries and dissidents took advantage of the fact that companies like Twitter imposed American free speech norms on content in countries like Egypt. For a more recent example, Brazilian populist demagogue Jair Bolsonaro attempted (albeit unsuccessfully) to prevent social media companies from removing his COVID-19 misinformation (*BBC News* 2021; Marcello 2021). Thus there is an inherent tension between the goal to promote the self-determination of peoples with diverse cultural and value systems and the goal to not be a tool of government oppression or demagoguery.

Nor can international human rights standards, on their own, resolve the tension between platform colonialism and platform culpability in governments gone bad, for human rights standards, like all laws and other normative principles, are subject to interpretive disputes, and it may still be objectionable for Western elites to make themselves the authorities to decide whether, for example, the request from some government to censor some content from one of its citizens comports with those standards. Moreover, the content of those standards themselves is disputed, especially by states that see themselves as having a religious identity shared with the entire community – thus, for example, the Cairo Declaration on Human Rights in Islam is widely seen as a response to the Universal Declaration on Human Rights which articulates a contrasting framework meant to be specifically Islamic.¹⁶

Even among wealthy liberal-democratic nations, there is disagreement on the scope and applicability of human rights norms, and that disagreement has already leaked into the platform governance domain. Internet human rights group Article 19 has criticized Germany's "NetzDG" intermediary liability law for (among other things) imposing liability on companies who do not help enforce Germany's prohibition against glorifying the National Socialist party – a law that, in Article 19's view, violates international human rights norms with respect to free speech.¹⁷ To me, Article 19's objection is deeply misguided: Germany has prohibited the National Socialist party for a very long time and *we all know why*. It's hard to imagine a more

¹⁶ A later human rights instrument with many of the same signatories, the Arab Charter on Human Rights, has been criticized for itself violating international human rights standards, in particular, in its denial of the principle of self-determination with respect to its condemnation of Zionism (Barnidge Jr 2018; United Nations 2008).

¹⁷ Article 19, "Germany: The Act to Improve Enforcement of the Law in Social Networks," August 2017, www.article19.org/wp-content/uploads/2017/12/170901-Legal-Analysis-German-NetzDG-Act.pdf p. 15.

compelling case for deference to a local interpretation of a global human rights norm in light of a country's distinctive history, culture, values, and dangers. But this just goes to show that even the extent to which global human rights standards are subject to local variation is itself controversial.

2.5.1 *Mitigating Platform Neocolonialism*

The force of the neocolonialism objection may be limited in two respects. First, it does not apply to platform governance innovations that are focused on the sorts of wealthy Western countries which are traditionally the sources rather than the victims of colonialism. Consider the governance failures that social media companies' inability to deal with the American alt-right and the incitement that led up to the January 6 attacks revealed: US-focused governance remedies that apply only in the United States, to be implemented by other Americans, might be lots of things – potentially racist and sexist, for example – but they are unlikely to be objectionable on grounds of colonialism (with the potential exception of disagreements about how such innovations are to be applied to situations involving Native American nations).

Second, to the extent objections to capacity-building efforts rooted in *uncontroversial* applications of international human rights principles (e.g., the prohibition of genocide or states repressing nonviolent opposition parties) are rooted in objections to liberal democracy itself, it may be that such objections are less to capacity building and more to the very presence of platform companies at all. As I argue in more detail in Chapter 5, there's a sense in which the platforms are inherently oriented to liberal democracy.¹⁸ If countries don't wish to comport with minimum human rights standards, they can always prohibit their citizens from using the platforms in question, as China does. In view of the fact that platforms are not imposed on countries, if a country rejects the very idea of human rights itself it seems to me that it is fully satisfactory to say that doing so amounts to a rejection of platforms as inherently encapsulating ideas like free markets and a discursive public sphere.

However, for the reasons discussed above, the foregoing is not a full answer to objections to human rights-derived platform governance. Countries or peoples originating in different cultural contexts may accept the idea of human rights but have different legitimate interpretations of the contents of human rights, or of the relationship between the behavior of individuals, companies, and states to human rights standards or other relevant normative standards. The problem of free expression is particularly salient here, not merely because of the German example noted above but also because the major platform companies originate in the United States, a country that is widely recognized as a global outlier in its free speech absolutism. And while recent developments (such as Musk's welcoming the alt-right back to

¹⁸ For more on rooting content moderation decisions in company values, see Bruckman (2022, 205–7).

Twitter) are extreme even by American standards, the people of other countries have previously noted a kind of American bias in platform interpretations of free speech norms relative to the interpretations of other countries relating to matters such as hate speech. In the words of one “non-Anglophone European” content moderator for Facebook, as told to the *New Yorker* “If you’re Mark Zuckerberg then I’m sure applying one minimal set of standards everywhere in the world seems like a form of universalism[.] To me, it seems like a kind of libertarian imperialism, especially if there’s no way for the standards to be strengthened, no matter how many people complain” (Marantz 2020).

Part of the answer to the problem of platform colonialism from capacity-building efforts may just be to observe that the status quo ante is already terrible. Americans and (to a lesser extent) Europeans are *already* regulating the behavior of people in other countries and imposing their own values on them. It’s not as if the alternative to capacity development for platform governance is autonomous self-governance by people from all countries and cultures, or even a void of governance. Likewise, countries such as the United States are already using platforms as an instrument of their own foreign policy and have the potential to do so to a greater extent under existing law and company institutional structures. For example, United States Customs and Amazon have run at least one “joint operation” to control the import of counterfeit goods (U.S. Immigration and Customs Enforcement 2020), and Amazon has reportedly channeled counterfeiters to police agencies in numerous countries (Brodkin 2021) – in effect helping the United States enforce an international as well as domestic intellectual property regime that, of course, is fairly heavily weighted toward US interests and against the interests of developing countries who may benefit from looser enforcement. On the other end of the platform economy, there has been substantial speculation that social media companies could be held liable under American “material support” statutes for hosting content by organizations that it designates as terrorists (VanLandingham 2017; Bedell and Wittes 2016).

Rather than a world free of platform colonialism, the alternative to governance reform may simply be destructive incompetent colonial governance partly by ham-fisted corporate efforts to maintain the market appeal of their platforms, partly by the unintended consequences of product design for other purposes, and partly by the short-term political goals of leaders in powerful countries – much like we currently have. Genuine governance reforms directed at platform companies may make things better almost by default.

Some kinds of governance improvements are also inherently improvements from an anti-colonial standpoint. For example, one of the key recommendations of this book is the integration of nonstate groups in both more and less economically developed countries into platform decisions. Such processes would track emerging human rights standards; for example, Article 19 of the United Nations Declaration on the Rights of Indigenous Peoples provides that such peoples should be consulted on decisions that affect them. To the extent we understand that the right to apply to

platform companies, including indigenous peoples in company decision processes represents a direct implementation of that right; in countries with governments directly descending from settler colonialism (such as the United States, Brazil, Australia, and Canada, for the most obvious examples), including indigenous peoples in company decisions, where previously only governments largely representing non-indigenous peoples had any influence over them, would amount to a positive reduction in colonialism.

That being said, these points do not eliminate the reason to guard against colonialism in platform governance development proposals. Even if any governance development program will represent an improvement over the (terrible) status quo, there can still be more or less colonial programs, and we in countries that bear culpability for the evil of colonialism and have the power to implement platform governance development programs have compelling moral reasons to implement those programs with the aim of reducing colonialism. Accordingly, the danger of colonialism in both cultural and political respects ought to be considered in any effort to reform and advance the capacity of platforms to govern their users. And we can state a minimal criterion for a successful reform proposal that it decrease, relative both to the status quo baseline and to other viable options, the risk of the United States and other wealthy powers, as well as corporate personnel aligned with those powers, exercising colonial power over the people of the global South.

In addition, governance improvements may exacerbate imperialism in the economic sense described above even if they do not exacerbate, or even if they mitigate, colonialism, insofar as they increase the capacity of platform companies to extract wealth from less economically dominant countries and peoples. Accordingly, governmental action to incorporate non-US peoples into platform governance should also be paired with efforts to bring about fair compensation for informational resources extracted from those peoples. However, while important, this is not the core concern of this book, as existing philosophical theories of global justice already offer a critique as well as proposed remedies for global economic imperialism (e.g., Caney 2001).

2.5.2 *Facebook Probably Shouldn't Have Been in Myanmar at All*

There's one easy way to avoid colonialism: It may be that American (and secondarily European) companies simply ought not to be in some countries. If companies lack the cultural competence or legitimacy (in either a sociological or a normative sense) to make rules for user behavior without creating injustice or exacerbating existing harms in a given country, and if the government of that country cannot be trusted to give the company its marching orders, then we might simply conclude that the presence of a given platform in a given country may cause more harm than good, and it ought to withdraw from the market.

In principle, the governments of developed countries – particularly the United States – in conjunction perhaps with United Nations human rights agencies, could

implement legislation in accordance with the previous paragraph. The United States already has a substantial legal framework permitting the designation of countries with which its companies are not permitted to trade. And platform companies are already covered by some of those laws, or at least appear to in some cases be voluntarily complying with them in the absence of debates over coverage – for example, there is evidence that Facebook has made use of the US government’s “Specially Designated Global Terrorists” list in developing its dangerous individuals ban list (Biddle 2021), which may amount to compliance with individually targeted sanctions provisions against those individuals.¹⁹ In principle, legislation could be enacted prohibiting any US company from operating in a country with respect to which the State Department or the United Nations High Commissioner for Human Rights has certified that the company in question has a severe, ongoing, negative human rights impact. And while the United States First Amendment may pose some challenges to the application of such a rule to social media, it may be defensible, particularly if the legislation in question is not specifically targeted at speech and applies to many other US companies as well. Such breadth, however, likely makes any legislation along these lines politically impossible: consider its application to US resource extraction companies which have generated numerous terrible human rights impacts, and the lobbying power of such companies (e.g., of possible application, see Giuliani and Macchi 2014, 485).

2.5.3 *Dispersing Power: Simultaneously a Governance Reform and an Anti-colonial Measure?*

Beyond outright withdrawal, another strategy for avoiding colonialism in a context in which the American state as well as American (and secondarily European) cultural values currently exercise a substantial governing role over people in less developed countries, is to design governance reforms that *reduce* and *disperse* that power, not to other Americans (or to local elites allied with particular elements in the American political/economic system), but to the actual people in the countries in question. In the later parts of this volume, I argue that dispersing power to the actual people whose interests are at stake (as opposed to a bunch of Americans) is actually likely to improve the efficacy of platform governance as well as to make it less colonial.

In the international rule-of-law development field (the area with which I am most familiar), similar goals have been integrated into some projects which have thereby avoided, or at least reduced, the danger of neocolonialism. Such projects have been distinctive for focusing on local empowerment: identifying sociologically legitimate

¹⁹ See, for example, 31 CFR §§ 594.201, 204, 310, describing prohibition on providing services to specially designated global terrorists, although exceptions for “informational materials” and related items may arguably cover some Facebook services. I hedge with the “may” in the text because I lack expertise in sanctions law.

local leaders, including those independent of formal state institutions, and giving them the tools they need to participate in local dispute resolution.²⁰ Such projects have succeeded, not by attempting to impose American or European conceptions of law and of adjudication but by consciously accepting the legitimacy of legal pluralism and so-called “informal” adjudication. By, in effect, meeting people where they are, such projects simultaneously are (at least potentially) more successful at actually bringing the benefits associated with a legal order to the recipients of development assistance while reducing colonialism.

A first pass solution to the problem of platform colonialism may also involve similar forms of local empowerment. Rather than assuming that governments – which might be dictatorial, or simply unrepresentative – are the right entities to introduce local variation in either the content, interpretation, or enforcement of platform policies, it may be possible to turn to, for example, civil society leaders, religious leaders, activists, and other representatives of diverse local groups to carry out these functions. And while identifying the appropriate people to whom to turn and recruiting them into the governance framework can be challenging, again strategies from international development, such as hiring experts with deep local knowledge and social science training to identify relevant leaders, might be imported into the platform context.

However, there’s still a degree of colonial imposition in such a plan. One of the things that colonizer powers actually did in the existing history of colonialism was identify and elevate local leaders whom the colonizers perceived, because of their own biases, preferences, and interests, to be most suitable. Some of the grimmest consequences of colonialism have their roots in this history; most infamously, it has been argued that the Rwandan genocide was at least in part attributable to the decision of Belgian and German colonizers to promote the Hutu/Tutsi ethnic distinction and make use of the Tutsi as a local aristocratic class (Newbury 1998). Obviously, similar practices must be avoided in the platform context. Fortunately, one advantage of the current time as opposed to the period of active colonialism is that there is an active transnational human rights community that can assist in identifying civil society organizations capable of participating in governance institutions, where human rights organizations – while doubtless still aligned with the interests of the powerful – are at least *more* neutral than companies and governments themselves.

Ultimately – and this is a point to which I will return in Chapter 6 – the only real answer to this problem is to accept imperfection and adopt a system of iterative inclusion, in which those who participate in governance institutions are subject to continuing critique and modification. The real challenge in institutional design is to create sufficiently open-ended modes of inclusion, which can receive feedback from excluded groups and respond to that feedback by effectively incorporating new people into decision-making processes, as well as structures for dispute resolution

²⁰ See discussion in Gowder (2016), 171–175.

between different groups who have a stake in any given decision.²¹ This is all unhelpfully abstract right now, but this chapter is directed simply at describing the general character of the problem – concrete elucidation of the sorts of designs that might satisfy these criteria will have to await Part III of this book.

Moreover, the attention to mitigating colonialism must also include attention toward mitigating other injustices, and sometimes there is a trade-off between the two – or at least an apparent trade-off. Consider the problem of subordinated persons within colonized peoples. For example, Bina Agarwal (2001, 1629) observes that when women are included in local forestry groups in India, it is often because of external pressure by “a local NGO, forest official, or international donor.” When women are included, resources are managed more effectively – but is it colonial for international development “donors” to impose the inclusion of women on local institutions which traditionally operate under more patriarchal norms?²² Perhaps, but there is a moral trade-off to be made (as well as an effectiveness tradeoff, since Agarwal also observes that gender-inclusive forestry groups work better), and ultimately I shall argue that the goal should be to include everyone – states in the global South *and* the indigenous and colonized peoples whom such states often inadequately represent *and* those, including often women, whom indigenous and colonized peoples themselves often inadequately represent. The only way to do so, as I said a moment ago, will be via an iterative learning process in which claims for inclusion by those who are left out of any pre-existing institutional arrangement are sought out and prioritized for action.

However, it may be that solutions to injustices within colonized peoples can only emerge from the peoples themselves. This is how I read Fanon’s (1994, 35–67) famous essay on the veil: while veiling may be an injustice to the extent it is imposed by tradition on (as opposed to voluntarily chosen by) women in colonized cultures,²³ a reconciliation between the autonomy of a colonized people and the remedy of the injustices within a pre-colonized culture can only occur through (and after) the process of shaking off the chains of colonization, and from the self-assertion of the colonized.²⁴

Particularly important in Fanon’s account is the way the veil becomes operationalized by the oppressed in the course of the collective action of the Algerian revolution.

²¹ If we believe Dewey (1927, 198–99), speaking of education but with logic that applies to many other social institutions, this is in a sense characteristic of all human social design: Using our knowledge of human sociality to intervene on the social world changes the underlying facts, and hence changes the interventions we may wish to make. Similarly, developing inclusive institutions, if done right, changes who is making the decisions about the shape of those very institutions, and hence changes who ought to be included. It’s iterative all the way down.

²² Cf. Okin (1999).

²³ I confess to not being sure about Fanon’s view on the pre-colonial injustice of the veil.

²⁴ Moreover, it is easy – because of the psychology of the colonizer – to represent efforts to destroy the colonized culture as efforts to “free” those whom it allegedly oppresses. If that’s right – and it seems right to me – then a project such as this book can never be untainted by colonialism. Even the judgment of which precolonial practices are unjust is likely to be inaccessible except from within colonized cultures.

Originally unveiling is a tool of the colonizer, but when women begin to participate in armed resistance, it becomes a tactical tool of that resistance, to be dropped or adopted as necessary in the broader goal of liberation. What this illustrates to me is that the important fact is the autonomy of those alleged to be oppressed even within a colonized people – autonomy that of course is necessarily conditioned both by the system of colonialism as well as by the restraints imposed by the colonized culture, but that can nonetheless be *taken* by their own self-assertion. Thus, returning to Agarwal's account of women's participation in the forest management process, she also notes that women had been observed to seize their own forms of participation through parallel organizations which were neither imposed on their communities from the outside nor merely licensed by the men (Bina Agarwal 2001, 1629–30). While she notes that these organizations are suboptimal as a method of including women insofar as the "official" organizations have formal control of the underlying forests, I submit that it might be a foundation for a kind of postcolonial self-assertion along Fanon's lines.

For present purposes, this suggests a modification to the concept of iterative inclusion, in which the ultimate goal of governance reforms is not to be understood merely as participation by colonized peoples but as *self-determination*, at least to the extent that such a thing is possible within the limited field of action encompassed by the book (for true self-determination would require changes well beyond the scope of the platform economy). That self-determination should be capable of interacting with any novel governance structure in both causal directions – that is, governance structures should be open to receiving demands of self-determination generated from the outside, but they should also be organized so as to promote the implementation of demands for self-determination generated from the inside.²⁵

2.6 OBJECTION 3: THE WHOLE INDUSTRY IS BAD – RADICALLY REMAKE IT BY FORCE

One question that immediately arises as we consider the harms platform companies have caused is: Should we (as citizens of democratic states) simply put a stop to them? Should we ban their services (consistent with constitutional and international free speech and other limitations)? Should we tax them into oblivion?

Those who write about the major platforms almost universally emphasize the social harms they create. And, doubtless, those harms are significant. Amazon crushes small businesses and abuses its workers, while facilitating a seemingly endless supply of counterfeits and scams. Facebook and Twitter and Instagram and the like promote viral conspiracy theories, political polarization (maybe, the empirical

²⁵ This issue may be intrinsic to human rights in general. Abdullahi Ahmed An-Naim (2021) makes an argument along these lines, suggesting that state-centric conceptions of human rights entail either relying on the very states that are violating those rights, or relying on external coercive intervention which itself creates human rights violations – instead, the protection of human rights should arise from political and social movements.

jury is still out (Tucker et al. 2018)), and psychologically unhealthy forms of social self-presentation and “news feed” addiction. Airbnb and Uber and the like make it dramatically more difficult for us to democratically regulate industries like hotel-keeping and livery services, and have engaged in troublingly exploitative work practices under the guise of independent contracting.

But Facebook and (pre-Elon) Twitter and Instagram have also provided many of us the services listed on the tin: They have allowed us to maintain and build real connections with other people which we might not otherwise have maintained.²⁶ There are friends from the many places I’ve lived with whom I keep in touch almost exclusively over social media; and the fact that the pictures of their dog or kid or whatever occasionally scroll through my feed makes it much easier to maintain the rewarding social relationships. Those relationships have social and emotional value that is real, and that was much harder to maintain before the existence of social media.²⁷ Those relationships have economic value too: as has been well understood at least since Granovetter (1973), things like jobs and other valuable opportunities are often learned about through “weak ties” – that is, people with whom one is acquainted, but less close than, say, one’s intimate friends or the colleagues whom one sees daily – precisely the ties that social media most helps us maintain.²⁸

²⁶ One caution that must be given is that revealed preference theories of social value cannot do the work here. A naive defense of platforms would take the fact that people choose to use them in the face of alternatives as evidence that they provide a benefit. For example, one study suggested that people would require compensation of just under fifty dollars a month to give up Facebook (Brynjolfsson, Collis, and Eggers 2019). The problem is that revealed preference theories of value are not sound, because people’s choices are endogenous to the market in which they find themselves. Thus, many critics of social media rightly worry that these products are addictive, and that people would be better off if they were protected from the urge to use them – a supposition supported by a number of empirical studies (e.g., Tromholt 2016; Dhir et al. 2018; Hunt et al. 2018). We also know that the entrance of new options in a market can affect the available alternatives – a prominent example being the way that social media engagement algorithms have degraded the quality of journalism by making available to people articles which reward attention to headlines rather than those which reward critical attention to content. (For a helpful discussion of endogenous preferences in markets, see Satz 2010, 180–81.) Finally, of course, we know that platforms, like all other forms of social activity, are sometimes characterized by negative externalities: individuals may personally benefit from their activity on platforms, but those activities may harm uninvolved third parties; revealed preference theories of social value only capture the benefit to participants and not the harm to bystanders.

²⁷ Of course, not every social media platform provides a meaningful amount of social value. It’s unlikely that Parler and Gab create any benefit to society, due to their function as focal points for right-wing extremist gathering. In a more nonpartisan context, many hyper-specialized kinds of social media can be hyper-specialized in things that are socially pernicious. The most striking example as of this writing is a downright dystopian application called “Citizen,” which is essentially a many-to-many social crime reporting platform that drives engagement by whipping up people into irrational levels of fear with pervasive “incident” notifications (that can be anything from an assault or a fire to a mere crowd of people), and, unsurprisingly, facilitates racial profiling as well as risking (or sometimes downright attempting to promote) vigilantism (Cox and Koebler 2021).

²⁸ There are sound reasons for this: people with strong ties are likely to occupy many of the same social spaces as oneself, and hence to learn about the same opportunities that one already is aware of. Weak ties are likely to bring to one’s attention opportunities arising in social contexts that one does not occupy.

And, indeed, I have learned of important publication and consulting opportunities from weak ties via social media.²⁹

The same can be said for the transactional platforms. For all the harm that businesses like Uber and Airbnb have caused, the fact of the matter is that they meet (or at least met at the time of their creation) a real need. For perhaps the most obvious example, the taxi market was egregiously anticompetitive if not outright corrupt until Uber and Lyft came along. At one point, New York taxi medallions were selling for up to 1.3 million dollars (Harnett 2018). That's blatantly corrupt: drivers would go into life-ruining levels of debt (Rosenthal 2019) to buy the privilege of participating in an artificially supply-constrained market; not all that different from student loans to go to medical school or law school, except that those professional degree programs provide an actual education necessary to do the job, whereas the taxi medallion system was just a pure government-sponsored cartel.³⁰ Or they would have to rent their medallions from those who could afford to own them – so that the government monopolies gave way to investor monopolies in which a rentier class who (until Uber and Lyft came along) effectively took no risk (thanks to the artificially constrained supply) and could simply print money.³¹ What possibly legitimate policy reason could there have been to create a system in which investors could speculate on occupational licenses? The level of total policy failure represented by the medallion system is just astonishing: by creating an artificially supply-constrained cartel and then putting membership in that cartel on the market, it managed to combine the worst aspects of capitalism with the worst aspects of central planning. Uber and Lyft did a good thing by blowing it up; the taxi drivers who were harmed in the process should seek recompense from the corrupt politicians and medallion oligarchs who built the predecessor system in the first place.

Similarly, Amazon, for all its brutal market power, in all likelihood helped save lives during the COVID-19 pandemic as its massive and painstakingly built logistics network meant that it could actually deliver necessities to people at semi-affordable prices rather than send them out into retail stores to infect one another with a deadly disease (and yet, at the same time, they probably killed a few as well by providing a platform for fraudulent sellers to distribute counterfeit N95 masks). Similar claims could easily be made of food delivery companies like Doordash and Instacart. In the beginning of 2021, as the federal government struggled to manage the distribution of the COVID-19 vaccine, President Biden reportedly considered recruiting companies like Amazon to help – and rightly so, since the kind of pervasive logistics

²⁹ My experience consulting with Facebook itself came from a then-weak tie who worked there, and who announced that her team was looking for someone matching my profile over Facebook (although ultimately working together for a period transitioned us from weak to strong ties).

³⁰ Even if you accept the so-called “signaling” theory of education, at least education actually provides that signal – whereas taxi medallions were nothing at all but an artificial supply constraint.

³¹ For example, one investor in 2006 bought an entire block of city-issued medallions at auction (Olshan 2006).

network that can somehow reliably deliver a potholder to the door of almost any American in two days had the potential to be immensely helpful in getting vaccines into arms (Scola and Ravindranth 2021).³²

The big mainstream platforms also cause undeniable harms. I take no position on whether the overall balance of social benefits and harms caused by platforms is in the black or the red. But I wish to urge caution here: nobody else is likely to be able to carry out that balancing act either. The harms they have inflicted on the world are probably incommensurable with their benefits, and that evaluation embarrasses the case for taking a prohibitionist approach to the companies.

For example: During the COVID-19 pandemic Amazon, Instacart, and similar delivery companies helped protect the health of those who could afford their services at the expense of workers, such as warehouse employees and delivery “contractors” (really, misclassified workers), who were unjustly placed at greater risk for the virus without being fairly compensated for taking that risk or, in many cases, having any real choice about the matter. This is a moral stain on the companies, on everyone who benefited from the harms inflicted on those workers, and on society at large. It is also a fairly common experience in capitalism, no different in any meaningful sense from the moral wrong committed for decades when Americans benefited from cheap and readily available energy to drive economic growth on the backs of brutally endangered coal miners. Critics from the left have been pointing out that workers are exploited under capitalism and endangered for the benefit of others for centuries, this has nothing to do with the platform economy. But can we honestly say that the world would be better off if Instacart had not existed during the pandemic? Or is the genuinely honest thing to say a variant of “something like Instacart should exist, but combined with massive redistributive taxation to ensure that those who benefit fairly compensate those who take the risk, and those who take

³² At this point, I suspect that some readers will accuse me of peddling “technology solutionism.” But I fear that the concept is easy to misuse. Sometimes technological solutions – or at least improvements, mitigations to problems – actually do exist. For example, Wymant et al. (2021) estimate that several hundred thousand COVID-19 cases in the UK over a three-month period were prevented by the National Health Service’s contact-tracing app, built on a joint Apple/Google API. If those estimates are correct, then – subject to potential (but hard-to-estimate) caveats about alternative ways to have expended the resources that the UK government spent on app development and perhaps sociological effects according to which contact tracing apps undermined the impetus for more traditional and effective public health measures – real human lives were saved, and that is unquestionably morally valuable. At its best, the concept of technology solutionism through the insights of thoughtful (if sometimes polemical) commentators like Evgeny Morozov (2013) serves to draw our attention to the risks of oversimplifying complex social problems and imagining that simple technological changes can fix them without attention to their underlying causes or networks of effects. When deployed that way, the critique of technology solutionism is compatible with the argument raised in this book (which, after all, directly calls for the incorporation of messy humanity and complexity into the governance of internet platforms). But sometimes “solutionism” can be used to stand in for the denial that technological innovations can improve anything, or for the claim that technological changes contribute no social value at all – and in that form, I think it’s just a mistake.

the risk do so freely?”³³ And in the absence of an obvious route to that alternative, does it make a lot of sense to try to engage in global evaluations of the with-Instacart and without-Instacart worlds?

Again, the same is true of the social platforms. Facebook and Twitter and Reddit and the like have unquestionably facilitated immense harms to our societies. Many people plausibly believe that Donald Trump wouldn't have been president without them, and that a mound of corpses of unknown size (related, e.g., to mismanagement of the COVID-19 pandemic and brutal immigration policies) is thereby attributable to their existence. Maybe so, but none of us have any clue how to actually evaluate that counterfactual in a world in which, for example, Fox News and talk radio exist, and in which the existence of the Internet itself makes it almost inevitable that spaces for dangerous political groups will exist outside the scope of big technology companies.

At any rate, heavy-handed regulation is a notoriously ineffective strategy for controlling products that produce real benefits to their users, even if they also produce social harms. Consider Uber again. In principle, municipalities could ban un-medallioned transport-for-hire altogether, and could subject people who offered or used those services to criminal sanctions. But this kind of prohibition strategy as to products and services that are already popular does not, it is fair to say, have a terribly good record. In all likelihood, now that the ridesharing genie is out of the bottle, any such prohibitory efforts would both increase the pervasiveness and harms caused by the intrusion of law enforcement in day-to-day life (you think the corporate surveillance done by platforms is bad? imagine the government surveillance necessary to control ride-sharing!) and be unsuccessful at controlling the conduct; in fact, exacerbate the harms of the conduct by driving it to black markets. See, for example, alcohol prohibition, the war on drugs, anti-prostitution policing, and countless other examples.

2.7 A CAUTIONARY APPROACH TO ANTITRUST LAW

A major policy lever that both the American left and the American right seem eager to pull – ranging from Elizabeth Warren (Kolhatkar 2019) to Josh Hawley (Bartz 2021) – is the use of antitrust law to break up the platform companies on the order of Standard Oil or AT&T, or at least to prevent additional corporate acquisitions. In some sense, this is the opposite of a capacity development strategy, as the goal becomes not to improve companies' capacity to govern user behavior on platforms but at least partly to defang it – and defanging it is at least part of the reason for such proposals (certainly it motivates Hawley's, as the US political right have often decried companies' alleged censorship of their views – more on this in Chapter 4). Nonetheless, to the extent that some of the social harms from platforms come from their market dominance, it is worth a brief discussion of this alternative strategy to close this chapter.

³³ On the unfree character of terrible jobs, see Gowder (2014a).

The premise of the use of antitrust law against platforms is that many of their pathologies are the consequence of market power: There are no realistic alternatives to Google for search or Facebook for social media (although surely the latter is plainly untrue), and hence the bad decisions of those companies are amplified across society at large, while network externalities make it unreasonably difficult for market discipline to control those decisions.

The political relevance of what many platforms – especially social media platforms – do adds fuel to the case for their breakup. As Tim Wu (2018, 132) has suggested, there may be a democratic case for increasing competition in the companies that control substantial shares of the public sphere similar to the case for efforts against traditional media concentration.³⁴ This is certainly a point that would be welcomed by the American political right, as there is increasing reason to think that at least Facebook intervenes (however ineffectually) to stop far-right movements from getting a foothold on the platform (Horwitz and Scheck 2021) – and, alas, in the United States, there appears to be very little distance between the extreme right and the mainstream Republican Party. Similarly, Zephyr Teachout (2020, 43–44) deploys the fundamental premise of this book – that platforms are behaving like governments – to suggest that we ought to see companies like Meta as in effect competing against the democratic authority of the United States, and hence as particularly amenable to being broken up.

However, conventional strategies of corporate dissolution may, unless paired with other substantial policy interventions, entail the sacrifice of important public interests that are actually served by the companies in their present forms. With respect to Amazon, notwithstanding all the company's serious misbehavior, it is eminently reasonable to suppose that there are genuine economies of scale in a national logistics network of the kind that Amazon built.

With respect to social media, some breakup strategies seem to risk – perhaps deliberately, as in the Hawley case noted above – undermining the goal of governance to the extent they might encourage companies that serve niche markets, which, in turn, may cater to socially harmful interests. Consider that, for all the harms they inflict, Meta and Alphabet have strong incentives rooted in their scale to be compatible with mass consumption markets: If, for example, either Facebook's News Feed or YouTube's recommended videos became pervasively filled with Nazi-adjacent content, both large numbers of ordinary users and advertisers would likely flee the platforms. The slow-rolling destruction of Twitter after the Musk acquisition, with users and advertisers fleeing the platform as the extreme right is welcomed back, stands as a clear example.³⁵

³⁴ This is in addition to the general democratic case articulated by Wu and others for breakups of companies that exercise excessive political power through the ordinary means of wealth and influence.

³⁵ While it is in principle possible for companies like Meta and Alphabet, aided by their industry-leading machine learning expertise, to segment their user and advertiser markets sufficiently fine that they can match users who are disposed to Nazi-adjacent content with that content and with advertisers who

By contrast, consider the actual niche social networks that have arisen in recent years: companies like Gab and Parler outright designed to appeal to Nazi-adjacent users. And consider their niche predecessors in enterprises like 4chan or its even scarier descendant Kiwi Farms. Lacking the constraints of market operations at scale, like mass-market users, mass-market advertisers, press attention, and publicly traded stock, these companies had no reason to control the incitement found on their platforms; indeed, it took a large and hence publicly accountable company, Amazon (in its infrastructure rather than platform business), to bring one of them at least temporarily under control, to wit, by kicking Parler off its Amazon Web Services infrastructure following its role in facilitating the organization of January 6, 2021, coup attempt at the US Capitol (Romm and Lerman 2021; Guynn 2021). In a vivid illustration of the effects of social isolation on extremism, at least a handful of the insurrectionists spent some of their time in the Capitol showing off for the presumptively sympathetic audience to be found on Parler (Klein and Kao 2021). According to legal pleadings filed by Amazon in its defense against a suit by Parler over the termination of its hosting service, Amazon had been urging Parler to control the threats of violence against Democratic officials such as Nancy Pelosi at least since November.³⁶ While Amazon (and Apple, and Google) may certainly be criticized for not putting a stop to Parler well before January 6, there is zero reason to believe that Parler's operators would have put a stop to its incitement on its own initiative *ever*.

This point draws on a larger worry about antitrust: to the extent there's a normative justification for antitrust law in consumer services, it revolves around the idea that a more competitive market will provide a wider variety and better prices of those services to the general public. But where services might be specialized to socially harmful market segments, it might actually be a bad thing to encourage the growth of niche service providers, and a good thing to permit the homogenization associated with big centralized providers.

One might imagine a number of alternative strategies for platform company breakups that could avoid these problems. For example, different services could be disaggregated: Facebook and Instagram could be separated; Amazon's delivery services and its network infrastructure (AWS) could be separated; Google search and YouTube could be separated. There might be some possibility of geographically

don't particularly care (or themselves are interested in purveying or exploiting Nazi-adjacent dispositions), the public prominence of those companies means that their recommendation algorithms are under intense press scrutiny, so that when this has in fact happened (presumably by accident), the news has come out and the companies have tried (however ineffectually) to remedy the problem to avoid being punished by their users and advertisers. For example, Facebook acted swiftly to get rid of algorithmically created anti-Semitic ad targeting categories after ProPublica discovered them (Angwin, Varner, and Tobin 2017).

³⁶ Defendant Amazon Web Services, Inc.'s Opposition to Parler LLC's Motion for Temporary Restraining Order, p. 4, filed January 12, 2021, in *Parler LLC v. Amazon Web Services, Inc.*, No. 2:21-cv-00031-BJR, United States District Court, Western District of Washington, at Seattle; brief available at https://storage.courtlistener.com/recap/gov.uscourts.wawd.294664/gov.uscourts.wawd.294664.10.o_1.pdf.

centered breakups – such as Amazon logistics across US regions – or along the lines of proxies for geography such as language.³⁷

On the social network side, larger companies might be able to achieve socially beneficial economies of scale in content moderation, for example, training machine learning models to identify certain kinds of content on much larger and more diverse datasets. However, in principle, antitrust strategies could be devised to avoid undermining such economies of scale, for example, by breaking up product sides of companies while leaving governance sides (such as “integrity” entities) intact. Voluntary industry cooperation could also ameliorate some of the difficulties of disaggregation in that respect.

On the whole, it seems to me – again, quite tentatively – that while antitrust breakups may be useful to ameliorate other harms that platforms impose on the world – such as invasions of privacy and impeding innovation by blocking market entrants – it is on balance unlikely to be effective as a solution to those problems that fall under the category of “governance” as used in these pages, within the domain of this book’s concern, that is, where long-term platform company interests and the interests of the rest of us are aligned.

The question of antitrust, however, fits more closely into the overall project of this book in two respects. First, the democratic reforms which this book defends may also be a component of a program to mitigate the public dangers of antitrust action of the sorts I question above. The core proposal of Chapter 6 is to create democratic institutions of platform governance that interact with, but are independent of, individual companies and operate across them. Creating such institutions might also make it less dangerous to break the companies up because those institutions themselves could embody some of the beneficial effects of scale, such as giving companies pressure to moderate content in accordance with the needs and desires of vast bodies of people (rather than wicked niches) and aggregating both practical knowledge as well as more brute tools like machine learning models or training data. In the world I envision in Chapter 6, governance economies of scale might be capable of being spread across the industry as a whole (depending on how much infrastructure to intervene directly in content moderation and other kinds of rule enforcement could be built around the participatory institutions I describe), and pressure from empowered ordinary people could be directed at the industry as a whole. Under such circumstances, perhaps even small platforms serving niche markets could benefit from overall governance reforms designed for the big companies.

From the other direction, antitrust law could be a *vehicle* for implementing those reforms. As Chapter 4 describes in some more detail, some of the governance

³⁷ Wu (2018, 83–91) argues that at least some economies of scale in ordinary commercial markets may be illusory. We may well find out after, for example, breaking up Amazon’s logistics network into regional sub-units that the smaller companies can still deliver things like affordable and fast delivery. It also might be that logistics is simply a public good which government ought to supply – that is, adequately fund the United States Postal Service.

pathologies of platforms result from poor internal governance, such as the excessive authority of social media company lobbyists over content policy decisions regulating the behavior of the same politicians whose favor they seek. In a recent article, Herb Hovenkamp (2021, 2021–32), one of the United States’ most prominent antitrust scholars, suggests using antitrust law to internally reorganize the decision-making structure of platform companies, for example, to empower different groups of stakeholders in order to prevent abuses of power.³⁸ Hovenkamp’s approach is quite similar to my own – he draws, as do I, on Ostrom’s governance theory,³⁹ and offers the compulsory integration of stakeholders in corporate governance structure as a tool to address platform abuses of power similar to – and in some cases overlapping with – the ones I address, such as Amazon’s abuse of data to discriminate against third-party sellers.⁴⁰ Accordingly, it may ultimately be the case that the reforms offered in this book are less an alternative to antitrust law than a potential product of it.

At this point, I consider the broad approach of a governance development framework adequately defended. The rest of this book turns toward how such a framework would work – the characteristic problems of governance and institutional solutions to those problems shared both by states that struggle to govern and by platforms that struggle to govern (Chapters 3 and 4) and the sorts of concrete institutions that might be built to alleviate them in the platform context (Chapters 5 and 6, and the Conclusion).

³⁸ It might also be used, for example, as a regulatory lever to force platform companies to permit the third-party creation of certain kinds of “middleware” (D. Keller 2021).

³⁹ Compare Hovenkamp (2021, 2022); Chapter 3 this volume.

⁴⁰ Compare Hovenkamp (2021, 2029); Chapter 4 this volume.