

No cover
image
available

The Oxford Handbook of AI Governance

Justin B. Bullock (ed.) et al.

<https://doi.org/10.1093/oxfordhb/9780197579329.001.0001>

Published: 2022

Online ISBN: 9780197579350

Print ISBN: 9780197579329

CHAPTER

Harms of AI

Daron Acemođlu

<https://doi.org/10.1093/oxfordhb/9780197579329.013.65> Pages C65P1–C65N5

Published: 18 July 2023

Abstract

This essay discusses several potential economic, political, and social costs of the current path of AI technologies. I argue that if AI continues to be deployed along its current trajectory and remains unregulated, it may produce various social, economic, and political harms. These include: damaging competition, consumer privacy, and consumer choice; excessively automating work, fueling inequality, inefficiently pushing down wages, and failing to improve worker productivity; and damaging political discourse, democracy's most fundamental lifeblood. Although there is no conclusive evidence suggesting that these costs are imminent or substantial, it may be useful to understand them before they are fully realized and become harder, or even impossible, to reverse, precisely *because* of AI's promising and wide-reaching potential. I also suggest that these costs are not inherent to the nature of AI technologies, but are related to how they are being used and developed at the moment—to empower corporations and governments against workers and citizens. As a result, efforts to limit and reverse these costs may need to rely on regulation and policies to redirect AI research. Attempts to contain them just by promoting competition may be insufficient.

Keywords: [AI](#), [automation](#), [communication](#), [competition](#), [data](#), [democracy](#), [inequality](#), [information](#), [monitoring](#), [politics](#), [social media.](#), JEL Classification: L13, L40, J23, J31, O33, P16.

Subject: [Political Institutions](#), [Politics](#)

Series: [Oxford Handbooks](#)

Collection: [Oxford Handbooks Online](#)

1 Introduction

To many commentators, artificial intelligence (AI) is the most exciting technology of our age, promising the development of “intelligent machines” that can surpass humans in various tasks; create new products and services; or even build other machines that can improve themselves, perhaps beyond all human capabilities. The last decade has witnessed rapid progress in AI, based on the application of modern machine learning techniques and huge amounts of computational power to massive, often unstructured data sets (e.g., Russell & Norvig, 2009; Neapolitan & Jiang, 2018; Russell, 2019).¹ AI algorithms are now used by almost all online platforms and in industries that range from manufacturing to health, finance, wholesale, and retail (e.g., Ford, 2015; Agarwal et al., 2018; West, 2018). Government agencies have also started relying on AI, especially in the criminal justice system and in customs and immigration control (e.g., Thompson, 2019; Simonite, 2020).

Whether AI will be everything its enthusiastic creators and boosters dream or not, it is likely to have transformative effects on the economy, society, and politics in the decades to come. Some of these effects are already visible in AI algorithms’ impact on social media, data markets, monitoring of workers, and work automation. Like many technological platforms (or “general purpose technologies”) that can be used for the development of a variety of new products, services, and production techniques, there are a lot of choices about how AI technologies will be developed. This, combined with the pervasive effects of AI throughout society, makes it particularly important that we consider its potential dark side as well.

In this chapter, I will focus on three broad areas in which the deployment of AI technologies may have economic and social costs if not properly regulated. I want to emphasize at the outset that the arguments I will present are *theoretical*—currently, there is insufficient empirical evidence to determine whether the mechanisms I isolate are important in practice. The spirit of the exercise is to understand the *potential harms* that unregulated AI may create so that we have a better understanding of how we should track and regulate its progress.

The areas I will focus on are:

Collection and control of information.

I will argue that the combination of the demand from AI technologies for data and the ability of AI techniques to process vast amounts of data about users, consumers, and citizens produces a number of potentially troubling downsides. These include: (a) *privacy violation*: companies and platforms may collect and deploy excessive amounts of information about individuals, enabling them to capture more of the consumer surplus via price discrimination or violate their privacy in processing and using their data; (b) *unfair competition*: companies with more data may gain a strong advantage relative to their competitors, which both enables them to exercise market power to extract surplus and also relaxes price competition in the marketplace, with potentially deleterious effects; and (c) *behavioral manipulation*: data and sophisticated machine learning techniques may enable companies to identify and exploit biases and vulnerabilities that consumers themselves do not recognize, thus pushing consumers to lower levels of utility and simultaneously distorting the composition of products in the market.

Labor market effects of AI.

I will argue that, even before AI, there was too much investment in cutting labor costs and wages in the United States (and arguably in other advanced economies as well). Such efforts may be excessive either because, in attempting to cut costs, they reduce production efficiency, or because they create non-market effects (for example, on workers losing their jobs or being forced to take lower-pay work). In principle, as a broad technological platform, AI could have rectified this trend—for example, by promoting the creation of new labor-intensive tasks or by providing tools for workers to have greater initiative. This does not seem to have taken place. Instead, automation is a quintessential example of efforts to cut labor costs, and like other efforts, it may be excessive. Many current uses of AI involve automation of work or the deployment of AI in order to improve monitoring and keep wages low, motivating my belief that AI may be exacerbating the excessive efforts to reduce labor costs. In this domain, I focus on four broad areas: (a) *automation*: I explain why automation, a powerful way to reduce labor costs, can be part of the natural growth process of an economy, but it can also be excessive because firms do not take into account the negative impact of automation on workers; (b) *composition of technology*: problems of excessive automation intensify when firms have a choice between investing in automation versus new tasks—I explain why this margin of technology choice may be severely distorted and how AI technologies may further distort this composition; (c) *loss of economies of scope in human judgment*: in contrast to the hope that AI will take over routine tasks and, in the process, enable humans to specialize in problem-solving and creative tasks, AI-human interplay might gradually turn humans into worse decision-makers as they hand over more and more decisions to machines, especially when there are economies of scope across tasks; and (d) *monitoring*: I also explain how technologies like AI that increase the monitoring ability of employers are very attractive to firms, but may at the same time generate significant social inefficiencies.

AI, communication and democracy.

Finally, I will suggest that AI has also exacerbated various political and social problems related to communication, persuasion, and democratic politics that once again predate the onset of this technology. The main concern here is that democratic politics may have become more difficult, or even fundamentally flawed, under the shadow of AI. I focus on: (a) *echo chambers in social media*: how AI-powered social media generates echo chambers that propagate false information and polarize society; (b) *problems of online communication*: online social media, which is interwoven with AI, may disadvantage fruitful political exchange relative to offline communication because there is less accumulated trust between communicators, which favors non-political speech and “broadcasting” rather than bilateral informative communication; (c) *big brother effects*: AI increases the ability of governments to closely monitor and stamp out dissent; and (d) *automation and democracy*: finally, I suggest that the process of automation may further damage democracy by making workers less powerful and less indispensable in workplaces. Each of these outcomes is damaging, but their most consequential effects are mediated by impairing democratic discourse.

In each one of the above instances, I discuss the basic ideas about potential costs, their root causes and how they will exhibit themselves in practice. Throughout, my approach will be informal, attempting to communicate the main ideas. For this reason, I present some of the details of the models that clarify these mechanisms in the [Appendix](#). Some of these models are based on existing work and others are offered as ideas for future exploration.

In addition, I will point out the relevant context and evidence in some cases, even though, as already noted, we do not have sufficient evidence to judge whether most of the mechanisms I explore here are likely to be important in practice. I then discuss the common aspects of the potential harms from AI and explore their common roots. I also argue that these costs, if proved important, cannot be avoided in an unregulated

market. In fact, I will suggest that in many of these instances, greater competition may exacerbate the problem rather than resolving it.

The aforementioned list leaves out several other concerns that experts have expressed (e.g., AI leading to “misaligned” super intelligence or AI’s effects on war and violence), because of space restrictions and because they are further away from my area of expertise. I mention them briefly toward the end of the chapter.

The long list of mechanisms via which AI could have negative economic, political, and social effects may create the impression that this technological platform is *bound* to have disastrous social consequences, or it may suggest that some of these problems are solely created by AI. Neither is true. Nor am I particularly opposed to this technology. I believe that AI is a hugely promising technological platform. Furthermore, with or without AI, our society has deep problems related to the power of corporations, automation and labor relations, and polarization and democracy. AI exacerbates these problems *because* it is a powerful technology and, owing to its general-purpose nature and ambition, it is applicable in a wide array of industries and domains, which amplifies its ability to deepen existing fault lines. These qualities make the potential negative effects of AI quite difficult to foresee as well. Perhaps even more than with other technologies and platforms, there are many directions—with hugely different consequences—in which AI can be developed. This variability makes it doubly important to consider the costs that AI *might* create. It is also vital to think about the *direction* of development for this technology.

Indeed, my point throughout is that AI’s costs are avoidable. If they were to transpire, this would be because of the choices made and the direction of research pursued by AI developers and tech companies. They would also be due to the lack of appropriate regulation by government agencies and societal pressure to discourage nefarious uses of the technology and to redirect research away from them. This last point is important: again, like most other technologies, but only more so, the direction of research of AI will have major distributional consequences and far-ranging implications for power, politics, and social status. It would be naïve to expect that unregulated markets would make the right trade-offs about these outcomes—especially because, at the moment, major decisions about the future of AI are being made by a very small group of top executives and engineers in a handful of companies. Put differently, AI’s harms are the harms of unregulated AI. But in order to understand what needs to be regulated and what the socially optimal choices may be, we first need to systematically study the potential downside of this technology. It is in this spirit that the following essay is written.

The rest of this essay is organized as follows. To begin, I start with the effects that AI creates via the control of information. The chapter then moves to discuss AI’s labor market implications. I then turn to the effects of this technological platform on social communication, polarization and democratic politics. Next, the chapter briefly touches upon a few other potential unintended consequences of AI technologies. I will then step back and discuss the role of choice in this process. I explain why the direction of technological change in general, and the direction of AI research in particular, is vital, and how we should think about it. This discussion also reiterates that many of the costs are the result of choices made about the development and use of AI technologies in specific directions. I then build on the mechanisms discussed in the chapter to emphasize that unregulated markets are unlikely to internalize AI’s costs, that greater competition may sometimes make things worse, and that unfettered markets are unlikely to direct technological change towards higher social-value uses of AI. In this spirit, I also provide some ideas about how to regulate the use of AI and the direction of AI research. The Appendix contains the details of the models discussed informally in the text.

2 AI and Control of Information

Data is the lifeblood of AI. The currently-dominant approach in this area is based on turning decision problems into prediction tasks and applying machine learning tools to very large data sets in order to perform these tasks. Hence, most AI researchers and economists working on AI and related technologies start from the premise that data creates positive effects on prediction, product design, and innovation (e.g., Brynjolfsson & McAfee, 2014; Jones & Tonetti, 2020; Varian, 2009; Farboodi et al., 2019). However, as emphasized by several legal scholars and social scientists, data and information can be misused—deployed in exploitative ways that benefit digital platforms and tech companies at the expense of consumers and workers (e.g., Pasquale, 2015; Zuboff, 2019). Zuboff, for example, argues that such exploitative use of data is at the root of the recent growth of the tech industry, which “claims human experience as free raw material for hidden commercial practice process of extraction, prediction, and sales.” (2019, p. 8).

First we must discuss the social costs of AI related to the control of data and information, with a special emphasis on exploring when data can become a tool for excessive extraction and prediction.

2.1 Too much data

Concerns about control and misuse of information become particularly important when there are benefits to individuals from “privacy”. Individuals may value privacy for instrumental or intrinsic reasons. The former includes their ability to enjoy greater consumer surplus, which might be threatened if companies know more about their valuations and can charge them higher prices. The latter includes various characteristics and behaviors that individuals would prefer not to reveal to others. This could be for reasons that are economic (e.g., to avoid targeted ads), psychological (e.g., to maintain a degree of autonomy), social (e.g., to conceal certain behaviors from acquaintances), or political (e.g., to avoid persecution).

Standard economic analyses tend to view these privacy-related costs as second-order for two related reasons: if individuals are rational and are given decision rights, then they will only allow their data to be used when they are compensated for it adequately, and this would ensure that their data will be used by companies only when the benefits exceed the privacy costs (e.g., Varian, 2009; Jones & Tonetti, 2020). Secondly, in surveys, individuals appear to be willing to pay only a little to protect their privacy, and hence the costs may be much smaller than the benefits of data (e.g., Athey et al., 2017). Yet these arguments have limited bite when the control of data has a “social” dimension—meaning that when an individual shares her data, she is also providing information about others. This social dimension is present, by default, in almost all applications of AI because the use of data is specifically targeted at learning from like-cases in order to generalize and apply the lessons to other settings.

How does this social dimension of data affect the costs and benefits of data? This question is tackled in a series of papers, including MacCarthy (2011), Choi et al. (2019), Bergemann et al. (2021), and, serving as the basis of this discussion, Acemoglu, Makhdoumi et al. (2022). The social dimension of data introduces two interrelated effects. First, there will be *data externalities*: when an individual shares her data, she reveals information about others. To the extent that data is socially valuable and individuals do not internalize this, data externalities could be positive. But if indirect data revelation impacts the privacy of other individuals, these externalities could be negative. The second effect is what Acemoglu, Makhdoumi et al. call *submodularity*: when an individual shares her data and reveals information about others, this reduces the value of others’ information both to themselves and to potential data buyers (such as platforms or AI companies). This is for the simple reason that when more information is shared about an individual, the individual’s own data becomes less important for predicting his or her decisions.

The Appendix provides a bare-bones model useful for understanding the role of data externalities and submodularity. Here I communicate the main ideas by discussing two simple examples (see the [Appendix](#) for details). Consider a platform with two users and suppose that one of them, user 1, has a relatively low value of privacy, while the other, user 2, has a high value of privacy. Data externalities are rooted in the fact that the two users have data that are very highly correlated, so the platform can learn a lot about user 2 from user 1's data. Given the low value that user 1 attaches to her privacy, the platform will always purchase her data. But this also implies that it will indirectly learn about user 2 given the correlation between the two users' data. If user 2's value of privacy is sufficiently large, it would be socially optimal to shut down data transactions and not allow user 1 to sell her data either. This is because she is indirectly revealing information about user 2, whose value of privacy is very large. This illustrates how data externalities lead to inefficiency. In fact, if user 2's value of privacy is very large, the equilibrium, which always involves user 1 selling her data, can be arbitrarily inefficient.

More interesting are the consequences of submodularity, which can be illustrated using the same example as well. To understand these, let us consider the edge case where the information of the two users is very highly correlated, so that the platform can learn everything relevant about user 2 from user 1's data, or vice versa. The important observation is that this data leakage about user 2 undermines the willingness of user 2 to protect her data. In fact, since user 1 is revealing almost everything about her, she would be willing to sell her own data for a very low price. In this extreme case with very highly correlated data, therefore, both the willingness of the platform to buy user 2's data and benefits user 2 receives from protecting her data are very small, and thus this price becomes approximately zero. But here comes the disturbing part for data prices and the functioning of the data market in this instance: once the second user is selling her data, this also reveals the first user's data almost perfectly, so the first user can only charge a very low price for her data as well. As a result, the platform will be able to acquire both users' data at approximately zero price. This price, obviously, does not reflect users' value of privacy. They may both wish to protect their data and derive significant value from privacy. Nevertheless, the market will induce them to sell their data for close to zero price. Imagine once again that user 2's value of data is also very high. Then, despite this high value of privacy to one of the users, there will be a lot of data transactions, data prices will be near zero, and the equilibrium will be arbitrarily inefficient. These consequences follow from submodularity.

As a second example, consider the case in which both users have the same value of privacy and suppose that these are sufficiently high that the platform could not fully compensate them and still make profits from acquiring their data. However, if the two users' data are sufficiently correlated, the platform may still be able to acquire their data, once again because of submodularity. The reasoning is similar to what we encountered in the previous example: when user 1 is expected to sell her data, this reduces the value of user 2's data, such that user 2 would rather cheaply sell her own data than be uncompensated for the platform's tacit use. Likewise, when user 1 expects user 2 to sell her data, user 1 would do likewise. This "coordination failure" between the two users is created and exploited by the platform to generate value but, critically, this is at the expense of the users. It is straightforward to see that this outcome is a consequence of submodularity as well: when each user expects the other one to sell their data, they become less willing to protect their own data and more willing to sell it relatively cheaply. This locks both users into an equilibrium in which their data is less valuable than they would normally assume and as a result there is, again, too much data transaction.

One final conclusion is worth noting. In addition to the results of excessive data transactions and use, these externalities also shift the distribution of surplus in favor of the platform. Even when it is socially optimal for data to be shared with and used by the platform, data externalities and correlation between the data of the two users will mean that economic benefits from data will be captured by platforms—not by the users. In particular, when the two users' data is very highly correlated, data prices will again be driven to zero, and thus all of the benefits from the use of data will be captured by the platform.

Are data externalities and the inefficiencies they create empirically relevant? Like many of the channels I discuss in this essay, the answer is that we do not know for sure. If, as industry insiders presume, benefits of data are very large, then they will outweigh the costs from data externalities I have highlighted here. Even in this case, the market equilibrium will not be fully efficient due to disincentives of platforms to compensate users nontrivially for their data, though the use of data by platforms and corporations may be welfare-increasing overall. However, there are reasons to believe that privacy considerations may be quite important in practice. First, many digital platforms have a monopoly or quasi-monopoly situation (such as Google, Meta, or Amazon), and thus their ability to extract rents from consumers can be significant. Second, some of the intrinsic reasons for consumers to care about privacy—related to dissent and civil society organizing—are becoming more important, as I discuss further in this chapter.

In summary, the general lessons in this case are clear: when an individual's data is relevant to others' behavior or preferences (which is the default case in almost all applications of data), then there are new economic forces we have to take into account, and these can create costs from the use of data-intensive AI technologies. In particular:

1. The social nature of data—enabling companies to use an individual's data for predicting others' behavior or preferences—creates externalities, which can be positive or negative. When negative externalities are important, corporations and platforms will tend to overuse data.
2. The social nature of data additionally generates a new type of submodularity, which makes each individual less willing to protect their own data when others share theirs. This submodularity adds to the negative externalities, but even more importantly, it implies that data prices will be depressed and will not reflect users' true values of data and privacy.
3. In addition to the result of excessive data use, both of these economic forces have first-order distributional consequences: they shift surplus from users to platforms and companies.

If these costs of data use and AI are important, they also call for regulating data markets. Some regulatory solutions are discussed in Acemoglu, Makhdoumi et al. (2022), and I return to a more-general discussion of regulation of AI technologies and data toward the end of this chapter.

2.2 Data and unfair competition

AI technologies amplify the ability of digital platforms and companies to use data to predict consumer preferences and behavior. On the upside, this might enable firms to design better products for customers (after all, this is one of the main benefits of AI). But the use of such data can also change the nature of competition. These effects become even more pronounced when some firms are much better placed to collect and use data relative to their competitors, and this is the case I will focus on in this subsection. Specifically, one firm's collection and use of data that others cannot access may create a type of "unfair competition", enabling this firm to capture much of the consumer surplus in the market and relax price competition.

A model elucidating how this may happen is provided in the Appendix. Here I will provide the high-level overview. In the pre-AI environment, two firms compete to attract customers. This forces both to charge relatively low prices, enabling consumers to capture most of the benefits from this market.

Next consider the post-AI environment in which one of the firms, say firm 1, can use machine learning, big data, and other methods in order to learn the exact preferences of the consumers it is supplying. This has an economic benefit: firm 1 can customize its products for each consumer. But it also has a dark side: because firm 1 now knows much more about its customers' exact preferences and willingness to pay, it can now price

these products so as to capture all or most of the consumer surplus. As a result, the direct effect of the use of AI technologies by firm 1 will be a transfer of economic value from consumers to the firm.

There is an indirect effect, potentially harming consumers as well. If firm 1 now focuses on providing customized products to its consumers, it may cease competing for the rest of the consumers (for whom it does not have detailed data). This will relax competition in the market, enabling the other firm, say firm 0, to also charge higher prices and thus capture some of the surplus previously obtained by consumers.

These two effects, taken together, tend to harm consumers and benefit firms—especially firms that can use AI technology to collect and process extensive data about their customer bases.

In summary, the general lessons from this discussion are complementary to the ones from the previous subsection:

1. The use of AI technologies and detailed consumer data for prediction may improve the ability of firms to customize products for consumers, potentially improving overall surplus.
2. However, it also increases the power of these companies over consumers.
3. This has direct distributional implications, enabling AI-intensive firms to capture more of the consumer surplus.
4. The indirect effect of the better collection and processing of data by one firm is to relax price competition in the market, increasing prices and amplifying the direct distributional effects.

Although in this model the overall surplus in the economy increases after the introduction of AI technologies, in the previous discussion we saw that this is not necessarily true in the presence of other data-related externalities. We will next encounter a new economic force distorting the composition of products offered by platforms.

2.3 Behavioral manipulation

The previous section discussed how even the beneficial use of improved prediction about consumer preferences and behavior might have a downside. But improved prediction tools can also be put to nefarious uses, with potentially far-ranging negative effects. Platforms that collect and effectively process huge amounts of data might be able to predict consumer behavior and biases beyond what the consumers are aware of themselves. Anecdotal examples of this concern abound. They include the chain store Target successfully forecasting whether women are pregnant and sending them hidden ads for baby products, or various companies estimating “prime vulnerability moments” and sending ads for products that tend to be purchased impulsively during such moments. They also include marketing strategies targeted at “vulnerable populations”, such as the elderly or children. Less extreme advertising strategies also have elements of the same type of manipulation: for example, when websites favor products such as credit cards or subscription programs with delayed costs and short-term benefits or when companies like YouTube and Meta use their algorithms to individuate, favoring more-addictive videos or news feeds based on each user’s preferences. As legal scholars Hanson and Kysar have noted, “Once one accepts that individuals systematically behave in nonrational ways, it follows from an economic perspective that others will exploit those tendencies for gain.” (1999, p. 630).

Although these concerns are as old as advertising itself, economists and policy-makers hope that consumers will learn how to shield themselves against abusive practices. The sudden explosion in the capabilities of digital platforms to use AI technologies and massive datasets to improve their predictions, however, undercuts this argument. Learning dynamics that had made consumers well-adapted to existing

practices would be quickly outdated in the age of AI and big data. Throughout this chapter, I discuss the main issues informally, though they are formally developed in Acemoglu, Makhdoumi, et al. (2023) and I provide an exposition in the [Appendix](#).

Suppose that a platform can find out more about what types of products a consumer is likely to enjoy, including which products she would be tempted to consume in the short run. In a world where the platform has interests aligned with those of the consumer, this additional information may improve welfare. The platform could nudge the consumer (either via which options it presents to her or through its pricing) to purchase the products that she is likely to enjoy. This is the optimistic take about how platforms could use their massive informational advantage in a world of AI and abundant data.

However, in most realistic scenarios, the interests of the platform and the consumers are likely to be misaligned. Consider a simple example. Suppose that the consumer in question may have some “tendencies” or biases that can be exploited in the short run—for example, she may be tempted to purchase products that provide instant gratification or have longer-term costs. From the viewpoint of overall welfare, it would be better for the consumer not to purchase these products, which ultimately might be very costly for her. And yet, if the platform can forecast which products create such temptations for the consumer, it will have a profit incentive to steer her towards these products. Thus, the greater informational advantage of the platform may lead to a type of “behavioral manipulation”, meaning that the platform will manipulate the behavior of the consumer in a direction that is deleterious to consumer welfare.

It is also worth noting that, in contrast to the pattern in the previous subsection, this type of behavioral manipulation would not only increase platform profits at the expense of consumers, but it would also distort consumption as it lures consumers towards lower-quality products, reducing overall welfare.

The general lessons in this case are complementary to, but different than, the ones I previously highlighted:

1. AI technologies can enable platforms to know more about consumers’ preferences than the consumers do, themselves.
2. This creates the potential for behavioral manipulation, whereby the platform can extract more surplus from consumers while distorting their choices.
3. This type of behavioral manipulation tends to do more than just shift surplus from consumers to the platform; it also distorts the composition of consumption, creating new inefficiencies.

3 Labor Market Effects of AI

US labor markets have not been doing well for workers over the last 40 years. Wage growth since the late 1970s has been much slower than during the previous three decades, while the share of capital in national income has grown significantly (Acemoglu & Autor, 2011; Autor, 2014). Additionally, wage growth, such as it is, has been anything but shared. While wages for workers at the very top of the income distribution—those in the highest tenth percentile of earnings or those with postgraduate degrees—have continued to grow, workers with a high school diploma or less have seen their real earnings fall. Even college graduates have gone through lengthy periods of little real wage growth.

Many factors have contributed to this sluggish average wage growth and real wage declines at the bottom of the distribution. The erosion of the real value of the minimum wage, which has fallen by more than 30 percent since 1968, has been clearly important for low-wage workers (Lee, 1999). The decline in the power of trade unions and much of the private sector may have played a role as well. The enormous increase in

trade with China also likely contributed, by forcing the closure of many businesses and large job losses in low-tech manufacturing industries such as textiles, apparel, furniture, and toys (Autor et al. 2013).

My own work with Pascual Restrepo (Acemoglu & Restrepo, 2019, 2021) emphasizes and documents the importance of the direction of technological progress in this process. In the four decades after World War II, automation went hand-in-hand with new task creation, which contributed to labor demand. A very different path of technological development emerged starting in the 1980s, however—this one exhibiting more automation and fewer advances in worker-friendly technologies to create new tasks (Acemoglu & Restrepo, 2019). Automation eliminated routine tasks in clerical occupations and on factory floors, depressing the demand and wages of workers specializing in blue-collar jobs and clerical functions. Meanwhile professionals in managerial, engineering, financial, designing, and consulting occupations flourished—both because they were essential to the success of new technologies and because they benefited from the automation of tasks that complemented their own work. As automation gathered pace, the wage gap between the top and the bottom of the income distribution magnified. In Acemoglu and Restrepo (2021), we estimate that automation has been, possibly, the most important factor in reshaping the US wage structure, explaining approximately 50 to 70 percent of the variance of changes in wages between 1980 and 2016.

All of this predates AI. In Acemoglu, Autor et al. (2021), we find that AI activity across US establishments picks up speed only after 2016. Nevertheless, this background is useful because AI may be the next phase of automation. There is evidence that it is already being used both for automation and for tighter monitoring of workers, thus further depressing wages and the labor share of income. In this section, I first explain how automation works and why we may be concerned about excessive automation in general, and how AI may exacerbate these concerns. I then discuss how AI could be used to generate new tasks and technologies that complement humans, but whether this will be the case or not depends on technology adoption and the research and development choices of companies. In this context, I suggest reasons for being concerned that the composition of AI research may be heavily distorted. I also discuss why the most benign view of AI's role in the labor market—automating routine jobs, so that workers have time for more creative, problem-solving tasks—may need to be qualified. Finally, I explore how AI may have pernicious effects when it is used for monitoring.

3.1 Excessive Automation and AI

In order to situate the role of AI in the broader context of automation technologies, I start with an informal review of the framework from Acemoglu and Restrepo (2018, 2019), which models the automation of tasks (as well as the creation of new tasks).

In this framework, production is modeled as the completion of a range of tasks. For example, textile production requires cleaning, carding, combing, and spinning the fibers; weaving, knitting, and bonding the yarn; designing, dyeing, chemical processing, and finishing the textile good; and various non-production tasks, including marketing, advertising, transport, wholesale, retail, and so on.

The key economic decision centers on the allocation of these tasks to capital and labor (and to different types of labor). Automation corresponds to expanding the set of tasks performed by capital. In addition to automation and other types of technologies, Acemoglu and Restrepo model the introduction of new tasks. New tasks create new employment opportunities for labor, and I return to a discussion of these new tasks in the next section.

It is also useful to allow for labor market imperfections: for example, modeled as a wage floor, w , below which equilibrium wages cannot fall. Without this wage floor, the labor market would be competitive and all workers who want to work will find employment, although—depending on which tasks are automated—all

employed workers may end up specializing in a narrow set of tasks. When there is such a wage floor and it is binding, some workers may become unemployed, and the extent of unemployment will depend on the extent of automation.

As I explain in the [Appendix](#), automation creates opposing productivity and displacement effects. The productivity effect helps labor, because it reduces production costs and subsequently increases labor demand. The displacement effect harms labor because workers are displaced from the tasks they used to perform and may need to be reallocated to perform other tasks where their marginal product may be lower. Even when the productivity effect is substantial, automation always reduces the labor share—and thus increases inequality between capital and labor. Worse, when the productivity effect is small, automation may reduce wages and employment. The productivity effect, in turn, will be small when automation happens in tasks where labor is already fairly productive and capital is not very productive. Hence, the technologies that are worse for labor are those that displace tasks through automation, but are themselves not very productive—what Acemoglu and Restrepo call “so-so technologies”. When the wage floor is binding, the introduction of so-so automation technologies will reduce employment.

What are the welfare consequences of employment-reducing automation? In a perfectly competitive market—where workers are at the margin, indifferent between leisure and work, and when there are no other distributional concerns—an automation-induced decline in employment does not have first-order welfare consequences. In fact, it is straightforward to see that the competitive equilibrium would always maximize net output (defined as total production minus what is used up to produce capital). However, when there are labor market imperfections, such as those captured by the wage floor \underline{w} , then low-productivity automation reduces welfare—thus motivating the term “excessive automation”. This can be seen with the following argument: because the productivity effect is approximately zero, gross output and profits do not increase (workers in marginal tasks are replaced by machines, but total costs have not changed). Yet, capital usage, which is costly, increases, and this reduces net output. At this point, reallocating marginal tasks away from capital towards labor—thus reducing automation—would increase net output.

Why is the equilibrium misaligned with social welfare maximization? The answer is related to the wage floor. Firms, when making their hiring and automation decisions, are responding to the market wage, \underline{w} , whereas a utilitarian social planner—seeking to maximize net surplus—should take into account the opportunity cost of labor, which is zero. This argument establishes that, when productivity effects are limited, there will be excessive automation. It also pinpoints one of the channels for this type of inefficiency: in economies with labor market imperfections, firms base their automation decisions on the higher wage rate, rather than the lower social opportunity cost of labor.

This argument also clarifies that automation is likely to be excessive and potentially welfare-reducing *especially when* it generates small or negligible productivity effects. If the productivity gains from automation had been large, net output would have increased, even if it displaced workers. Moreover, with a large productivity effect, there may not have been a decline in labor demand in the first place (see the [Appendix](#)).

The case for excessive automation is strengthened if there are other considerations favoring higher levels of employment. For example, if employed individuals generate positive external effects (on their families and communities or for democracy) relative to the unemployed, then the social planner may want to increase employment beyond the equilibrium level. Distributional concerns would also weigh in the same direction because, in general, automation helps firms and firm owners, while reducing the labor share. In addition, as shown in Autor et al. (2003) and Acemoglu and Restrepo (2021), automation boosts inequality across worker groups, creating another distributional cost.

What does this imply for AI? AI is a broad technological platform, and can be used for developing many different types of technologies. Automation, especially automation of various white-collar tasks and jobs with significant decision-making components, is one of these applications. If AI is used for automation, then the arguments outlined above would also imply that low-productivity AI may reduce welfare. Two key questions are, thus, whether AI technologies are likely to be deployed as capital and algorithms substituting for labor in various tasks and whether this will generate small or large productivity gains. The evidence in Acemoglu, Autor et al. (2021) suggests that there has been a significant uptick in AI activities since 2016, and that much of this has been associated with task displacement. That paper also finds reduced hiring in establishments that adopt AI technologies, so the evidence is consistent with, though does not prove, the idea that new AI technologies may not sufficiently improve productivity. There are other reasons why productivity gains from AI may be small. Most importantly, AI technologies are being used in some tasks in which humans are quite good (e.g., natural language processing, facial recognition, problem-solving; see Acemoglu, 2021).

In summary, the general lessons from this section are:

1. Automation reduces the labor share and may also reduce the (average) wage and/or employment, and the latter outcome is more likely when productivity gains from automation are small.
2. When labor market imperfections create a wedge between the market wage and the social opportunity cost of labor, automation tends to be excessive and welfare-reducing, particularly when it also impacts employment negatively. This too is more likely to be the case when the productivity gains are small. The same considerations apply when there are non-market reasons for preferring high levels of employment (for example, because employed workers contribute more to their families, communities or society in general).
3. Because it increases the capital share and reduces the labor share, and because it boosts inequality among workers, automation may also be excessive from a welfare point of view due to distributional concerns
4. If AI is used predominantly for automation, it will have similar effects to other automation technologies, and depending on its productivity effects and relevant welfare criteria, it may have a negative impact on social welfare.

3.2 Direction of AI technology and its labor market consequences

Some implications of AI used for automation were previously discussed, but it was also noted that AI, as a broad technological platform, can be used for creating new tasks or increasing labor productivity as well. Acemoglu and Restrepo (2018, 2019) show that the introduction of new (labor-intensive) tasks also has two effects: a productivity effect (because it raises the productivity of labor) and a “reinstatement effect”, rooted in the fact that it reinstates labor centrally into the production process.

The productivity effect is positive as usual (even if the exact sources of productivity gains from new tasks are different than those from automation). The reinstatement effect is also positive, because new tasks generate new employment opportunities for labor. As a result, new tasks always increase employment and/or wages. Moreover, the presence of the reinstatement effect implies that the wage bill increases proportionately more than the productivity gains, pushing up the labor share—the converse of the impact of automation. Acemoglu and Restrepo (2019) have argued that the reason why wages grew robustly during the decades following World War II was that rapid automation in certain tasks went hand-in-hand with the introduction of sufficiently many new tasks, counterbalancing the labor market implications of automation.

Returning to the implications of AI, with the same argument, using AI for new tasks would be welfare-improving, especially when there are labor market imperfections or other considerations favoring higher levels of employment than in equilibrium. Furthermore, if AI boosts the creation of new tasks and improves human productivity, it could counterbalance some of the adverse effects of automation by other technologies (e.g., robotics or specialized software).

When AI can be used both for automation and for new task-creation, the pivotal question becomes how the balance between these two activities is determined; that is, the direction of technological change. Acemoglu and Restrepo (2018) provide a framework for the analysis of the equilibrium direction of technology. This framework clarifies how the direction of technology depends on factor prices and the labor share, and also emphasizes that the equilibrium and the optimal directions of technology will often differ. In particular, labor market imperfections not only promote too much automation—as we saw in the previous subsection—but also tend to generate an unbalanced composition of AI research between automation and new tasks (see also Klinova, 2023, in this *Handbook*).

There may also be reasons for distortions in the direction of technological change that go beyond the purely economic. In Acemoglu (2021) I emphasize that the direction of technology is partly shaped by the business models of leading firms and the aspirations of researchers. If these favor automation, the equilibrium may involve too much automation, even absent economic distortions. A related argument is that US corporations may have become too focused on cost-cutting, which might also encourage excessive automation. Acemoglu et al. (2020), on the other hand, show that the US tax code imposes a much higher marginal tax rate on labor than on equipment and software capital, thus favoring automation. This policy channel triggers both excessive adoption of automation technologies and disproportionate emphasis on new automation technologies in research and development.

AI as a technological platform could, in principle, boost efforts to create new tasks. Take education as an example. Current investments in this area are focused on using AI technologies for automated grading and the development of online learning tools to replace various tasks performed by teachers. Yet, AI can be deployed for creating new tasks and directly increasing teacher productivity as well. It can be used for adapting teaching material to the needs and attitudes of diverse students in real time, overcoming a major problem of classroom-based teaching—the fact that students have diverse strengths and weaknesses and find different parts of the curricula challenging (see the discussion in Acemoglu, 2021). Likewise, AI has many diverse applications in health that can personalize care and empower nurses and general practitioners to make more and better decisions in care delivery. These promising directions notwithstanding, AI may be more likely to aggravate excessive automation. The current trajectory in AI research is shaped by the visions of large tech companies, who are responsible for the majority of the spending on this technology. Many of these companies have business models centered on substituting algorithms for humans, which may make them focus excessively on using AI for automation. At the same time, many AI researchers focus on reaching “human parity” in narrow tasks as the main metric of success, which could create another powerful force towards automation, rather than towards using this platform to create new tasks. Like other automation technologies, AI may also appeal to executives intent on cost-cutting. If there are additional tax breaks and favorable treatments for software in general and AI-related technologies, specifically, these may exacerbate the focus on automation (Acemoglu & Johnson, 2023).

Overall, even though there is no definitive evidence on this question, it is possible that the direction of technological change was already tilted too much towards automation even before AI, which may have exacerbated these trends. If so, one of the major, harmful effects of AI could be its labor market implications.

The general lessons from this discussion are, therefore:

1. AI could, in principle, be used to increase worker productivity and to expand the set of tasks in which humans have a comparative advantage, rather than focusing mainly on automation. If used in this way, AI may counterbalance some of the negative effects of automation on labor and may generate positive welfare effects and beneficial distributional outcomes.
2. But there is no guarantee that the composition of technological change, in general, and the balance of AI between automation and more worker-friendly task creation will be optimal. In fact, there are many possible distortions, some of them economic and some of them social, that encourage excessive automation using AI.

3.3 AI and human judgment

The arguments in the previous two sections are partly predicated on the notion that AI-based automation may not generate sweeping productivity gains, which could compensate for—or even undo—the displacement effects it creates. AI's most-enthusiastic boosters, on the other hand, believe that AI can bring huge productivity gains. One of the most powerful arguments for this outcome is that, as AI helps automate and improve (both cognitive and noncognitive) tasks that do not require human judgment and creativity, it will increase the demand for problem-solving tasks that require creativity and judgment and also free workers to focus on these tasks. Although seemingly plausible, I now suggest a potential reason why this expectation may be too optimistic and argue that, even when such reallocation takes place, AI-based automation may be excessive.

The main idea is simple (and a more-detailed treatment is provided in the [Appendix](#)). In many activities there are “economies of scope”, meaning that individuals acquire knowledge from performing certain tasks that can then be used in other tasks. When there are such economies of scope, automation of a subset of tasks may reduce potential human productivity in the remaining tasks as well. Moreover, in this case, firms will often adopt AI in order to reduce their costs in the AI-suitable tasks—ignoring the effects that this will have on the productivity of workers in other tasks.

More generally, a finer division of labor and the reallocation of some tasks away from humans can be cost-reducing, but to the extent that human judgment improves when workers gain experience from dealing with a range of problems and recognizing different aspects of those problems, it may also come at a cost. When some aspects of the problem are delegated to AI, workers may lose their fluency with, and ability to understand, the holistic aspects of relevant tasks, which can then reduce their productivity—even for the tasks in which they specialize. An extreme example of this phenomenon can be given from the learning of mathematical reasoning. Calculators are much better than humans in arithmetic. But if students stopped learning arithmetic altogether, delegating all such functions to calculators and software, their ability to engage in other type of mathematical and abstract reasoning may suffer. For this reason, most mathematical curricula still emphasize the learning of arithmetic. If delegating certain tasks to AI becomes similar to the hypothetical cessation of learning arithmetic, it may have significant costs.

In the presence of economies of scope, the market equilibrium may involve the automation of tasks that could reduce worker productivity and adversely affect the overall efficiency of the economy.

In summary, the general lessons from this short discussion are:

1. In addition to the costs of worker displacement, discussed earlier in this section, economies of scope across tasks may create additional costs from the use of AI technologies. In particular, the deployment of AI in various cognitive tasks that do not require a high degree of human judgment and creativity may enable workers to reallocate their time towards tasks that involve judgment and creativity. But if economies of scope are important for human productivity, AI may have additional costs.

2. The cost-minimization incentives of firms may encourage them to use AI technologies in ways that forgo economies of scope.

3.4 AI and Excessive Monitoring

Another use of AI-powered technologies is in worker monitoring, as exemplified by Amazon's warehouses and new monitoring systems for delivery workers. Here, too, employers' incentives to improve monitoring and collect information about their employees predates AI. But once again, AI may magnify their ability to do so. Some amount of monitoring by employers may be useful to improve worker incentives. However, I argue that increasing employer flexibility in this activity can also lead to inefficiently high levels of monitoring, for a very simple reason: at the margin, monitoring is a way of shifting rents away from workers towards employers, and thus is not socially valuable. But precisely because it shifts rents to firms and will often increase their profits, firms would have an incentive to engage in monitoring even when it is not socially efficient. In such situations, new technologies that extend employers' abilities to engage in monitoring may be socially harmful.

The economic force here is, potentially, quite general: AI, by enabling better control and use of information, provides one more tool to employers that can shift rents away from workers and towards themselves, leading to inefficiently high levels of rent-shifting activities.

Is this potential inefficiency relevant? Once again, there is little systematic evidence to suggest one way or another, but the fact that the US labor market has a "good jobs" problem, and that wages at the bottom of the distribution have fallen in real terms over the last several decades (Acemoglu, 2019; Acemoglu & Restrepo, 2021), suggests that it may be.

In summary, the general lessons from this analysis are:

1. AI technologies also create new opportunities for improved monitoring of workers. These technologies have first-order distributional consequences because they enable better monitoring and, thus, lower efficiency wages for workers.
2. Because, at the margin, the use of monitoring technologies transfers rents from workers to firms, monitoring will be excessive in equilibrium. By expanding monitoring opportunities, AI may thus create an additional social cost.

4 AI, Political Discourse, and Democracy

The 1990s witnessed a rapid strengthening of democracy around the world, in a pattern the political scientist Samuel Huntington (1991) called "The Third Wave". During this process, many Latin American, Asian, and African countries moved from nondemocratic regimes towards democracy and several others strengthened preexisting democratic institutions (Markoff, 1996). The last 15 years have witnessed a pronounced reversal of this process, however. Several countries have moved away from democracy, and perhaps even more surprisingly, Democratic institutions and norms have come under attack in numerous Western nations (Levitsky & Ziblatt, 2018; Snyder, 2017; Mishra, 2017; Applebaum, 2020) and the citizenry appears to be more polarized than in the recent past (Abramowitz, 2010; Judis, 2016). Some have pointed to social media and online communication as major contributing factors to these headwinds (e.g., Marantz, 2020). I now turn to a discussion of these issues. I focus on the effects of AI on communication, political participation, and democratic politics. As indicated at the beginning of this chapter, I will highlight several distinct, but related, mechanisms via which AI might degrade democratic discourse.

4.1 Echo chambers and polarization

Social media is often accused of promoting echo chambers in which individuals communicate with others who are like-minded, which might prevent them from being exposed to counter-attitudinal viewpoints, consequently exacerbating their biases. Cass Sunstein noted the potential dangers of echo chambers as early as 2001. He stated that encountering individuals with opposing opinions and arguments is “important partly to ensure against fragmentation and extremism, which are predictable outcomes of any situation which like-minded people speak only with themselves”, and emphasized that “many or most citizens should have a range of common experiences. Without shared experiences, a heterogeneous society will have a much more difficult time in addressing social problems.” (Sunstein, 2001, p. 9). The recent documentary *The Social Dilemma* describes the situation as, “[t]he way to think about it is as 2.5 billion Truman Shows. Each person has their own reality with their own facts. Over time you have the false sense that everyone agrees with you because everyone in your news feed sounds just like you.” AI is crucial to this new reality on social media. For example, the algorithms that sites like Facebook and Twitter use to decide what types of news and messages individuals will be exposed to are based on applying AI techniques to the massive amount of data that these platforms collect (Alcott & Gentzkow, 2017; Guriev et al. 2022; Mosleh et al., 2021). Recent studies document that these algorithmic approaches are exacerbating the problem of misinformation on social media, for example by creating algorithmic “filter bubbles”, whereby individuals are more frequently exposed to news that agrees with their priors and biases than to news that challenges those priors (Levy, 2021). As a result, Vosoughi et al. (2018) conclude that on social media there is a pattern of “falsehood diffused significantly farther, faster, deeper, and more broadly than the truth in all categories of information”.

I will discuss these issues, building on the approach in Acemoglu, Ozdaglar, et al. (2021), which suggests that social media platforms may endogenously create such echo chambers and more so in the presence of extremist content. Imagine an online community in which each individual receives a news item, which may contain misinformation. This misinformation may take the form of completely fake news or it may involve some misrepresentation of facts. The primary, individual decision is whether to further share the news.

Let us assume that if an agent is found out to have shared misinformation, she incurs a cost. To avoid this, she may instead decide not to share the article or even call it out for containing misinformation.

Suppose that there are two sub-communities, one is left-wing and the other is right-wing. Suppose also that each member of one of these communities is more likely to be connected to and sharing items with other members of her community, but there are also some cross-community links. The extent of these cross-community links determines how much “homophily” there is. With high homophily, there are almost no cross-community links. With low homophily, cross-community links are more common. Finally, suppose that each news item contains some “message”, which is relevant for political decisions and beliefs, and also a reliability score, which measures the likelihood that the news contains misinformation. How do individuals make their sharing decisions?

The answer depends on the type of message they receive and the extent of homophily in the online community. Take a right-wing agent receiving a left-wing message. Given her political beliefs, she is much more likely to think that this message contains misinformation than a right-wing message. Therefore, all else being equal, she is much more likely to dislike and not share it. This effect is further magnified when she expects to share it with other fellow right-wingers because they are also more likely to dislike a left-wing message and call it out for containing misinformation.

In contrast, consider this right-wing agent receiving a right-wing message. Now she is less suspicious of the message and is thus more likely to share it. Interestingly, homophily now has the opposite effect on

sharing decisions: if the individual is in a homophilic network, she expects other right-wingers to share the news item, and as a result, she is more likely to share it herself, creating a viral spread.

These observations lead to the first, basic result of the setup: a news item is more likely to become viral when (1) it reaches individuals who have congruent beliefs, and when (2) this concordant news item is being shared with others with similar beliefs.

Now consider the problem of the platform on which this community is situated. Suppose that, via its choice of algorithms, the platform influences the degree of homophily (as in Acemoglu, Ozdaglar, et al., 2021). Suppose also that the platform's objective is to maximize engagement and hence, greater virality of the article is better for the platform. First suppose that most of the news items arriving from the outside have high reliability scores and are also being distributed roughly equally between the two sub-communities (so that it is not only left-wing messages going to the left-wing community and likewise for the right-wing community). Then the platform's engagement-maximizing policy is likely to be to introduce between-community links, thus exposing each individual to news items from the other side, since these high-reliability items will not be immediately disliked and may even reach a large audience in the two communities.

In contrast, suppose we have a situation in which there are many low-reliability news items and, moreover, left-leaning news items from the outside go to the left-wing community and right-leaning news items go to the right-wing community. If the platform were interested in stopping misinformation, it could choose low homophily (so that right-wing articles go to the left-wing community as well, for example), and this would induce less sharing among the right-wing agents. It would also cause interruptions to virality when left-wingers discover the right-wing news item to contain misinformation (which is likely, given that the news item has low reliability). These considerations imply that when news items have lower reliability, the platform would prefer to induce extreme homophily via its algorithms and propagate misinformation in order to maximize engagement.

This result, mimicking the main finding of Acemoglu, Ozdaglar, et al. (2021), is in some ways quite striking. It highlights that the platform has incentives to create endogenous echo chambers (or filter bubbles). Worse, this happens precisely when there are low-reliability news items, which are likely to contain misinformation and be distributed within the sub-communities in a "polarized fashion" (e.g., left-wing messages going to left-wing groups, etc.).

The role of AI technologies is, again, crucial. Without these technologies, the platform would not be able to determine users' biases and create relatively-homogeneous communities. It would also not be able to support the rapid dissemination of viral news items.

Broader social implications of these types of filter bubbles are easy to see as well. Suppose that individuals also update their beliefs on the basis of the news items. When left-wingers receive right-wing news items and the relevant news items have reasonable reliability scores, they will tend to moderate their beliefs, exactly as Sunstein (2001) envisaged in the quote above. In contrast, when left-wingers receive only left-wing news items in a filter bubble, this might lead to further polarization. On the basis of these considerations, Acemoglu, Ozdaglar, et al. (2021) suggest various interventions, including regulation and outside inspections, in order to discourage such filter bubbles and reduce polarization. I return to the issue of regulation later in this chapter.

In summary, this section leads to the following general lessons:

1. AI-powered social media presents a variety of new opportunities for connecting individuals and information sharing.

2. However, in the process, it may also distort individuals' willingnesses to share unreliable information. When social media creates echo chamber-like environments, in which individuals are much more likely to communicate with like-minded others, they become less careful and more likely to share news items that are consistent with their existing views and more willing to allow the circulation of misinformation.
3. Centrally, social media platforms that are focused on maximizing engagement have an incentive to create echo chambers (or "filter bubbles"), because interruptions of the circulation of news items with unreliable messages reduces engagement. As a result, especially when there are more items with misinformation, platform incentives are diametrically opposed to social objectives.

4.2 Perils of online communication

The previous section argued that political discourse may be hampered because of the algorithmic policies of digital platforms. I now suggest that there might be reasons more-endemic to the nature of online communication that disadvantage political communication (see also Lanier, 2018; Tirole, 2021). The main argument is a version of the ideas developed in Acemoglu, Kleinberg et al. (2022), and here I provide a brief discussion (with more details in the [Appendix](#)).

The main ingredient for the effects discussed here: there is more mutual information about, and greater trust between, participants in a real-world social network than in an online social media network. This trust makes it more likely that the individual will be able to influence the beliefs of their acquaintances in the real-world social network via political discourse. When they expect that political communication is feasible, individuals may prefer it to gossip. In contrast, in online social media, most interactions are between people who have less knowledge about each other and less trust in the communicators. With the same reasoning, this will tilt the balance towards gossip and other non-political forms of communication. As a result, online communication may become inundated with gossip at the expense of useful political exchange.

The situation may be worse when we take into account that online interactions typically take the form of broadcast (rather than bilateral) communication. This would exacerbate the situation I have outlined here for two reasons. First, there may be many agents who may want to broadcast their views, competing for attention. Because broadcasting may be particularly attractive to extremists, they may be overrepresented among those engaged in such competition. This would endogenously increase the assessment of all the agents that political news is coming from extremists. Second, if there is heterogeneity in the utility of gossip across agents, those who value gossiping most may be the ones monopolizing the channels. In both cases, online communication becomes less effective as a way of sharing politically or socially relevant information. If political communication and news sharing in social networks are an important aspect of democratic politics, then the forces identified in the subsection also create new challenges for democracy, again rooted in the use of AI technologies.

The general lessons from this brief analysis can be summed up as follows:

1. Bilateral, offline communication, especially when the subject matter is political or social, relies on trust between parties. The trust that naturally exists in the context of in-person social networks may enable this type of communication.
2. When communication is taking place online and in multi-lateral settings, such as in modern social media platforms powered by AI technologies, this type of trust-based communication becomes harder. This may favor non-political messages, such as gossip, which then drive out political communication, with potentially deleterious effects for political discourse and democracy.
3. This potential barrier to online communication is exacerbated when there is competition for

attention, which is encouraged by the broadcast or multi-lateral nature of online communication.

4.3 Big Brother effects

The previous two sections focused on how AI-powered social media and online platforms change the nature of communication, with potentially negative effects on the sharing of political information, which is the bedrock of democratic participation by citizens. In this section, I suggest that the other crucial pillar of democratic institutions, citizen protests, is likely to be hampered by AI technologies.

I have argued in Acemoglu and Robinson (2000, 2006) that protests, riots, and uprisings are critical for the emergence of democratic regimes (because the threats that they pose for power-holders in nondemocratic regimes induces democratization). This argument is relevant for democratization in currently authoritarian governments, such as China, Russia, and Iran. Many AI-based technologies, including facial recognition, tools for identifying online and offline dissent, and data collection software that underpins citizen-score schemes (such as the Chinese social-credit system) can significantly strengthen authoritarian governments and even make it near-impossible for opposition groups to organize. Such technologies can significantly increase the longevity of nondemocratic regimes.

The problem is not confined to these nondemocratic nations, however, and applies to the US and other countries with democratic institutions as well. A similar argument suggests that protests and civil disobedience are often critical for the functioning of democratic regimes. The civil rights movement in the US illustrates this vividly. Even though the US was democratic at the federal level, the Jim Crow South routinely violated the political, social, and economic rights of Black Americans. Democratic institutions in the North and, to the extent that they existed, in the South did not create a natural impetus for these discriminations to cease. The turning point came with civil disobedience organized by various Black (and later multi-ethnic) civil society groups, such as the NAACP. Vivaly, even federal politicians opposed to Jim Crow were not in favor of these protests initially, viewing them as disruptive political drawbacks, especially given that any federal action against Jim Crow practices would trigger backlash from Southern politicians (see the discussion in Acemoglu & Robinson, 2019). Without civil disobedience, protest, and other sources of bottom-up pressures, it is likely that reform of voting, civil, education, and discrimination laws in the American South would have been further delayed.

In the [Appendix](#), I provide a simple model illustrating how AI technologies prevent political dissent and protest activities. When there is the threat of protest, civil disobedience, and other types of dissent, a nondemocratic elite (or even an elite within an imperfect democracy) will be tempted to make various compromises in order to prevent adverse reactions from civil society. But when AI enables governments and corporations to shut down dissent, it also discourages elite compromises. As a result, political decisions become more elite-biased and political distortions multiply.

Overall, the general lessons from this discussion are:

1. AI technologies can be used to improve governmental monitoring of, and preventative action against, protest activities.
2. Because the threat of protests has a disciplining role on nondemocratic governments, and even on some democratic governments, the shift of power away from civil society towards governments will weaken democracy and aggravate policy distortions.

4.4 Automation, social power and democracy

The previous subsection explained how the use of AI as a tool for controlling society and political dissent can have harmful effects on democratic politics. I argue that AI-powered automation can further weaken democracy and undermine social cohesion.

To develop this argument, let us first go back to the framework in Acemoglu and Robinson (2006), which emphasizes two types of conditions for the emergence and survival of democratic institutions. First, there must be enough discontent with nondemocratic regimes to generate sufficient demand for democracy. Second, democracy should not be too costly for the elite, who would otherwise prefer to use repression or other means to avoid sharing political power with the broader population. One aspect of this problem, which could be important, but was not analyzed in Acemoglu and Robinson (2006), is cooperation from workers. For example, if workers become disenchanted with the current regime or decide that they need to take action against current (economic or political) power-holders, they may choose not to cooperate with capital in their workplaces. When labor is essential for production, this withdrawal of cooperation could be very costly for capital. When capital owners are influential in the political system, they may then push for democratization or redistribution in order to placate labor.

The main argument in this section is that automation may make workers less indispensable in workplaces, and as such, it will tend to reduce their political power. Moreover, when technology enables the automation of most of the tasks previously performed by humans, workers could be made completely unnecessary in production, thereby obliterating their political power (see also Boix, 2023, in this *Handbook*).

I present a model formalizing this argument in the [Appendix](#). The main lessons from these ideas and the model are:

1. Automation can generate an indirect negative impact on democracy and redistributive politics when ensuring cooperation from labor in workplaces is an important motivation for elites to make concessions to labor.
2. When automation brings only small productivity gains, it encourages the elite to reduce redistribution and make fewer democratic concessions. This will make policies less responsive to the majority's wishes and may further raise inequality.
3. Productivity benefits of automation may soften this effect, because an automation-driven increase in output raises the opportunity cost of losing labor's cooperation. There may also exist a sufficiently high level of automation such that, once we reach this level, labor becomes sufficiently irrelevant for production that the withdrawal of workers' cooperation ceases to be very costly. After this threshold, the elite may prefer to abandon democratic institutions and withhold any concessions—proceeding without the cooperation of workers—once again, with harmful effects on democracy, redistribution, and social cohesion.

5 Other potential costs

I will now briefly list a few other areas that may be important, but without providing as much detail or formal analysis.

The threat of AI and bargaining. Even when the actual labor market effects of AI discussed previously are not realized, the threat of adopting AI technologies may influence wages and inequality. Specifically, if there is bargaining and rent-sharing, employers may use the threat of AI-based automation as a way of increasing

their bargaining power, which could have some of the same effects as actual automation and AI-powered monitoring.

Discrimination. Bias in AI has already received considerable attention. As AI gains greater importance in our social and economic lives, ensuring that popular algorithms are fair and unbiased has become vital. Existing studies show that simple AI algorithms can improve important public decisions, such as bail or sentencing, without increasing discrimination (e.g., Kleinberg et al., 2018). However, in most such applications, algorithms use data generated from biased agents and potentially discriminatory practices (e.g., Thompson, 2019). For example, significant portions of the police and the judiciary in the United States are generally thought to be biased against certain groups, such as Black Americans. In such situations, there is a danger that preexisting biases will become a fundamental part of AI algorithms. This may not only promote persistent bias and discrimination, but may in fact cement these biases more deeply in society via a process similar to the “signaling role of laws” (e.g., Posner, 2002). Indeed, if society starts trusting AI algorithms, their discriminatory choices may come to be accepted as more justifiable than when they were made by individual decision-makers.

Technocracy versus democracy. Advances in AI may create the temptation to delegate more and more public decisions—or even political decisions—to algorithms or, rather, to the technocrats designing and using these algorithms. Although reliance on these “experts” may be justifiable for certain decisions, an excessively relied-upon technocracy, without citizen input, may also start encroaching into political decisions—such as the extent of redistributive taxation or how much we should protect disadvantaged groups (Sandel, 2020). In this case, reliance on AI may further undermine democracy, amplifying the concerns highlighted in the previous section.

AI-powered weapons. AI technologies have already started to be incorporated into weapons and are advancing towards autonomous weapon systems. These new technologies will cause a host of ethical and social dilemmas, and should perhaps be regulated before prototypes are deployed or even fully developed. In addition to these ethical and social issues, AI-powered weapons may further strengthen governments against civil society, protesters, and even some opposition groups, adding to the concerns we discussed in the previous section.

The alignment problem. The potential downside of AI technologies that has received the most attention is the “alignment problem”: the problem of ensuring that intelligent machines have objectives that are aligned with those of humanity. Although all of the harms of AI emphasized so far can be thought of as rooted in a misalignment between the current crop of AI technologies and broader societal objectives, many researchers and public intellectuals have been concerned with another aspect of misalignment: machines reaching super-human capabilities and then, implicitly or explicitly, turning against humans (e.g., Bostrom, 2014; Russell 2019; Christian, 2019). My own view is that these concerns are somewhat overblown and often distract from shorter-term problems created by AI technologies (on which this chapter has focused), but naturally, they deserve careful consideration, monitoring, and preparation. The general misalignment problem is studied through the lens of economic externalities by Balwit and Korinek (2023) in this *Handbook*.

The international dimension: The current development of AI technologies is intertwined with international competition, especially between the US and China (Lee, 2018). A discussion of AI regulation has to take into account this international dimension. For example, it may not be sufficient for the US and Europe to start regulating the use of data or excessive automation when these issues remain almost completely unregulated in China. This suggests that the regulation of AI needs to have a fully-fledged international dimension and we may need to build new international organizations to coordinate and monitor deregulation of AI across the globe.

6 The Role of Technology Choice and Regulation

In the preceding discussions, I consider a number of theoretical arguments which suggest that the deployment of AI technologies may generate economic, political, and social costs. In this section, I highlight that, in all of these cases, the problems are not inherent to AI technologies, per se. Rather, the harms I have emphasized are caused by corporate and societal choices about how these technologies are deployed. Even though these costs are far-ranging, taking place in product markets, in labor markets, and in the realm of politics, they exhibit a number of commonalities, which I explore in this section. I also discuss some possible remedies. The general emphasis in this section will be on three main ideas:

1. The importance of choices, both on the use of existent AI technologies and on the direction of AI research. The costs I have modeled are not intrinsic to AI technologies, and are instead based upon how this new technological platform is currently being developed: to empower corporations and governments against citizens and workers.
2. The inadequacy of market solutions that mainly rely on increasing competition.
3. The need for regulation.

Let me start with the mechanisms discussed earlier. All three potential costs of AI turn on how AI technologies enable the use and control of data. In each one of these cases, a different way of distributing control rights over data would ameliorate or prevent most of the costs (Posner & Weyl, 2019). Let us start with data markets. The source of inefficiency in this case is the ability of platforms to find out information about others from the data that an individual shares. This then opens the way to potential misuses of data—for example, to reduce the surplus of consumers or to violate their privacy in other ways. Effective regulation in this case could take one of two forms. First, as suggested in Acemoglu, Makhdoumi et al. (2022), it may be possible to strip away parts of the data of an individual in order to prevent or minimize information about others being leaked (although the details matter here, as simply anonymizing data is not sufficient). Second, more systematic regulations on how platforms can use the information they acquire would lessen the harmful effects working through privacy.

In contrast, increasing competition may not be sufficient, and not even useful, in this case. My analysis focused on a monopoly platform. Acemoglu, Makhdoumi et al. (2022) show that if there are two platforms competing to attract users, this may exacerbate the pernicious effects of data externalities.

Let me then turn to the problem that was next considered. The source of inefficiency in this case is the ability of platforms to use the data that individuals reveal about themselves in order to manipulate their weaknesses. If this misuse of data can be prevented or if consumers can be made more-aware of how data are being used, some of these costs could be prevented. Suppose, for example, that consumers are informed frequently that platforms know a lot about their preferences and may use this knowledge to (very effectively) market low-quality products to their users. There is no guarantee that such informational warnings will work for all consumers, but if they are displayed saliently and are specific (e.g., calibrated according to the group of individuals and relevant class of products), they may prevent some of the harms identified above. In this case, too, increasing competition would not be an effective solution. If two platforms are competing for consumers, but consumers continue to be semi-behavioral and fail to recognize the increasing platform capabilities, both platforms may exploit their abilities to offer products that have apparent short-term benefits but disproportionate long-term costs.

The issues are similar when we turn to economic forces, but now the implications of competition are more nuanced. In this case, effective regulation would prevent one of the firms from using the additional information it acquires to capture all of the consumer surplus. Methods like price controls and limits on

price discrimination might achieve these ends, although clearly such regulation is far from straightforward. What happens if we can increase competition in this case? Greater competition that results from firm 0 also using AI methods to estimate its own past consumers' preferences and customize its services accordingly would not be necessarily useful. Now both firms become local monopolists, capturing all of the consumer surplus. However, if each firm can also acquire information about the other's customers and if collusion can be prevented, then they can be induced to compete fiercely, with better outcomes for consumers. This case thus highlights that, in some scenarios, fostering competition might have benefits—though only to a limited extent and only when some case-specific conditions are satisfied.

I would also like to emphasize the implications for the direction of AI research in this case. Suppose that AI researchers can devote their time to developing alternative applications of this broad technological platform. For example, some of them may be able to use AI to create tools that empower citizens or consumers, or develop new technologies for preserving privacy. All the same, if any one of the mechanisms related to the control and misuse of information are relevant, then this mechanism will also produce a powerful demand for technologies that enable corporations to acquire and better-exploit this type of information. These effects are exacerbated when the ability of consumers to pay for alternative technologies is limited, relative to the resources in the hands of corporations. In such scenarios, the demand for “misuse of AI” will be transmitted to AI researchers, who may then respond by devoting their time to developing the AI technologies that corporations demand and by moving away from technologies that have greater social value or that empower consumers and citizens. This is a general point, which applies whether the harmful effects of AI are on the control of information, labor markets, or politics. It is for this reason that innovation, when unregulated, is unlikely to produce self-correcting dynamics. To the contrary, the demand for misuse of AI will typically distort the allocation of AI research across different applications, amplifying the social and economic costs of this research.

The same considerations apply even more-evidently in the models presented in the discussion of automation. If automation is excessive, increasing competition in the labor market would not be particularly useful. On the other hand, the demand for automation technologies from firms will tend to be strong, encouraging researchers to double down on using AI to develop automation technologies. Regulatory solutions are feasible, but may be more difficult to design and implement in this case. In theory, when automation is excessive and AI research is not being directed to creating new tasks, welfare-promoting regulation should discourage automation at the margin and encourage the creation of new labor-intensive tasks.

However, distinguishing marginal (low-productivity) and infra-marginal (higher-productivity) automation is difficult. Even more challenging is that regulators might have a hard time separating the AI that is used to create new tasks from the AI that is used to automate low-skill tasks and empower higher-skilled or managerial workers. But it is also possible to view these problems not as absolute barriers, but instead as measurement challenges. More research might shed light on how to distinguish different uses of AI in the labor market and might reveal new regulatory approaches for influencing the direction of AI research.

Finally, there are similar lessons from the models we discussed when considering politics and democracy, although there are also some new challenges related to the fact that the effects are now on political and democratic outcomes. For one, increasing competition is unlikely to be a very effective way of dealing with misaligned platform incentives. For example, if there are multiple social media platforms trying to maximize engagement, each may have incentives to create filter bubbles. Pro-competitive solutions may also be less effective when there are systemic issues, such as the widespread malfunctioning of democratic institutions.

The effects of these new technologies for democratic politics raises new conceptual issues as well. Most importantly, if the incorrect deployment of AI technologies is weakening democratic politics, developing after-the-fact regulatory solutions might become harder because democratic scrutiny of those who benefit from the distortionary use of AI technologies would also become more difficult. These considerations suggest a “precautionary regulatory principle”—an *ex ante* regulation slowing down the use of AI technologies, especially in domains where redressing the costs of AI become politically and socially more difficult after large-scale implementation. AI technologies impacting political discourse and democratic politics may be prime candidates for the application of such a precautionary regulatory principle.

7 Conclusion

In this chapter I have explored several potential economic, political, and social costs of the current path of AI technologies. I have suggested that if AI continues to be deployed along its current trajectory and remains unregulated, it could harm competition, consumer privacy, and consumer choice. It may excessively automate work, fuel inequality, inefficiently push down wages, and fail to improve productivity. It may also make political discourse increasingly distorted, cutting one of the lifelines of democracy. I have also mentioned several other potential social costs from the current path of AI research.

I should reemphasize that these potential harms are theoretical. There is much evidence indicating that not-all is well with the deployment of AI technologies and that the problems of increasing market power, disappearance of work, inequality, low wages, and meaningful challenges to democratic discourse and practice are all real—though we do not have sufficient evidence to be certain that AI has, thus far, been a serious contributor to these troubling trends.

Nevertheless, precisely because AI is a promising technological platform, which aims to transform every sector of the economy and every aspect of our social lives, it is imperative for us to study the possible downsides, especially given its current trajectory. It is in this spirit that I have discussed the potential costs of AI in this chapter.

My own belief is that several of these costs are real and we may see them multiply in the years to come. Empirical work exploring these issues is, therefore, greatly needed.

Beyond empirical work, we also need to understand the nature and the sources of these potential costs and how they can be prevented. It is for this reason that I have suggested various policy responses, in each case emphasizing that the costs are rooted in the way that corporations and governments choose to develop and use these technologies. Therefore, my conclusion is that the best way of preventing these costs is to regulate AI and redirect AI research away from these harmful endeavors and towards areas where AI can create new tasks that increase human productivity and new products that can empower workers and citizens. Of course, I realize that such a redirection is challenging. The regulation of AI is probably more difficult to design than the regulation of many other technologies—both because of its fast-changing, pervasive nature and because of the international dimension. We must also be careful because history is replete with instances of governments and powerful interest groups opposing new technologies with disastrous consequences for economic growth (e.g., Acemoglu & Robinson, 2012). Redirecting technological change does not mean withholding technological advances, though it clearly requires government policy and societal pressure to discourage net-detrimental investments and deployment at the margin, while encouraging the development of more-optimal AI technologies. The centrality of these technologies to our future, and their potential harms, justify the need for these conversations.

Appendix

In this Appendix, I present sketches of several models aimed at clarifying the economic mechanisms discussed in the text.

Model for Too Much Data

Following Acemoglu, Makhdoumi et al. (2022), consider a community consisting of n agents/users interacting on a (monopoly) digital platform. Each agent i has a type denoted by x_i which is a realization of a random variable X_i , where the vector of random variables $X = (X_1, \dots, X_n)$ has a joint normal distribution $\mathcal{N}(0, \Sigma)$, with covariance matrix $\Sigma \in \mathbb{R}^{n \times n}$ (and $\Sigma_{ii} = \sigma_i^2 > 0$ denoting the variance of individual i 's type). Each user has some personal data, S_i , which are informative about her type. Personal data include both characteristics that are the individual's private information (unless she decides to share) and also data that she generates via her activity online and offline. Suppose $S_i = X_i + Z_i$ where Z_i is a normally-distributed independent random variable, $Z_i \sim \mathcal{N}(0, 1)$.

Although Acemoglu, Makhdoumi et al. (2022) discuss various metrics, here I suppose that the relevant notion of information is mean square error (MSE). Then we can define *leaked information* about user i as the reduction in the mean square error of the best estimator of her type:

$$\mathcal{I}_i(\mathbf{a}) = \sigma_i^2 - \min_{\hat{x}_i} \mathbb{E} \left[(X_i - \hat{x}_i(\mathbf{S}_\mathbf{a}))^2 \right],$$

where S is the vector of data the platform acquires, $\hat{x}_i(\mathbf{S})$ is the platform's estimate of the user's type given this information, and $a = (a_1, \dots, a_n)$ is the data-sharing action profile of users (with $a_i = 0$ denoting no direct data-sharing and $a_i = 1$ corresponding to data-sharing). Then, the objective of the platform is to maximize:

$$\sum_i [\eta \mathcal{I}_i(a) - a_i p_i],$$

where p_i denotes payment ("price") to user i from the platform, which is made only when the individual in question shares her data directly (i.e., $a_i = 1$), and $\eta > 0$. The price could take the form of an actual payment for data shared or an indirect payment by the platform, for example, the provision of some free service or customization. This specification embeds the idea that the platform would like to acquire data in order to better forecast the type/behavior of users.

User i 's objective is different. She may wish to protect her privacy and she obviously benefits from payments she receives. Thus her objective is to maximize:

$$\gamma \sum_{i' \neq i} \mathcal{I}_{i'}(a) - v_i \mathcal{I}_i(a) + a_i p_i.$$

The first term represents any positive direct externalities from the information of other users (for example, because this improves the quality of services that the individual receives and does not fully pay for) and thus $\gamma \geq 0$. The second term is the loss of privacy (capturing both instrumental and intrinsic values of privacy). Hence $v_i \geq 0$ here denotes the value of privacy to user i . Finally, the last term denotes the payments she receives from the platform.

This framework allows data to create positive or negative total benefits. To illustrate this point, suppose that $v_i = v$. In that case, data creates aggregate (utilitarian) benefits, provided that $\eta + \gamma(n - 1) > v$. In contrast, if $\eta + \gamma(n - 1) < v$, the corporate control and use of data is socially wasteful (it creates more damage than good). But even in this case, as we will see, there may be data transactions and extensive use of data. In general, because v_i differs across agents, data about certain users may generate greater social benefits than the costs, while the revelation of data about others may be excessively costly.

In terms of market structure, the simplest option is to assume that the platform makes take-it-or-leave-it offers to users in order to acquire their data.

A key result, proved in Acemoglu, Makhdoumi et al. (2022), is that $\mathcal{J}_i(\mathbf{a})$ is monotone and submodular. The first property means that when an individual directly shares her data, this weakly increases the information that the platform has about all individuals, i.e., $\mathcal{J}_i(\mathbf{a}') \geq \mathcal{J}_i(\mathbf{a})$ whenever $\mathbf{a}' \geq \mathbf{a}$. Mathematically, the second implies that, for two action profiles \mathbf{a} and \mathbf{a}' with $\mathbf{a}'_{-i} \geq \mathbf{a}_{-i}$, we have:

$$\mathcal{J}_i(\mathbf{a}_i = 1, \mathbf{a}_{-i}) - \mathcal{J}_i(\mathbf{a}_i = 0, \mathbf{a}_{-i}) \geq \mathcal{J}_i(\mathbf{a}_i = 1, \mathbf{a}'_{-i}) - \mathcal{J}_i(\mathbf{a}_i = 0, \mathbf{a}'_{-i}).$$

Economically, it means that the information transmitted by an individual who directly shares her data is less when there is more data-sharing by others.

I now illustrate the implications of this setup for data sharing and welfare using two simple examples. Consider first a platform with two users ($i = 1, 2$) and suppose that $\gamma = 0$, $\eta = 1$, and $v_1 < 1$ so that the first user has a small value of privacy, but $v_2 > 1$, implying that, because of strong privacy concerns, it is socially beneficial to not share user 2's with the platform. Finally, suppose that the correlation coefficient between the data of the two users is $\rho > 0$. Because $v_1 < 1$, the platform will always purchase user 1's data. But this also implies that it will indirectly learn about user 2, given the correlation between the two users' data. If v_2 is sufficiently large, it is easy to see that it would be socially optimal to close-off data transactions and not allow user 1 to sell her data either. This is because, by selling her data, user 1 is indirectly revealing information about user 2, whose value of privacy is very large. This illustrates how data externalities lead to inefficiency. In fact, if v_2 is sufficiently large, the equilibrium—which always involves user 1 selling her data—can be arbitrarily inefficient.

Let us illustrate this possibility using the same example. Consider the edge case where the information of the two users is very highly correlated, i.e., $\rho \approx 1$. In this example, the platform will know almost everything relevant about user 2 from user 1's data. The important observation is that this data leakage about user 2 undermines the willingness of user 2 to protect her data. In fact, since user 1 is revealing almost everything about her, user 2 would be willing to sell her own data for a very low price. In this extreme case with $\rho \approx 1$, both the willingness of the platform to buy user 2's data and benefits user 2 receives from protecting her data are very small, and thus this price becomes approximately zero. But the disturbing part for data prices and the functioning of the market in this instance: once user 2 is selling her data, this also reveals the user 1's data almost perfectly, such that user 1 too can only charge a very low price for her data (despite the fact that she values her privacy as well, albeit less than user 2). As a result, the platform will be able to acquire both users' data at approximately zero price. This price, obviously, does not reflect users' value of privacy. They may both wish to protect their data and derive significant value from privacy. Nevertheless, the market will induce them to sell it for close to zero price. Imagine once again that v_2 is sufficiently high. Then, despite this high value of privacy to one of the users, there will be a lot of data transactions, data prices will be near zero, and the equilibrium will be significantly (arbitrarily) inefficient. These consequences follow from submodularity.

As a second example, consider the case in which again $\gamma = 0$ and $\eta = 1$, but now there is no heterogeneity between the two users, so that $v_1 = v_2 = v > 1$. This configuration implies that neither user would like to sell their data (because their privacy is more important than the value of data to the platform). Nevertheless, it can be shown that as long as v is less than some threshold \bar{v} (which is, itself, strictly greater than 1), there exists an equilibrium in which the platform buys the data of both users relatively cheaply. This is also a consequence of submodularity: when each user expects the other one to sell their data, they become less willing to protect their own data and more willing to sell it relatively cheaply. This locks both users into an equilibrium in which their data is less valuable than they would normally assume and, partly as a result, there is again too much data transaction.

In addition to leading to excessive data use and transactions, the externalities also shift the distribution of surplus in favor of the platform. To see this, suppose $v_1 = v_2 = v \leq 1$ and $\rho \approx 1$, so that it is now socially optimal for data to be used by the platform. It is straightforward to verify that, in equilibrium, data prices will again be equal to zero and, thus, all of the benefits from the use of data will be captured by the platform.

Model for Data and Unfair Competition

Suppose that, as in the basic Hotelling model, consumers are located uniformly across a line of length 1 and incur a cost—similar to a transport cost—when they purchase a product further away from their bliss point, represented by their location (see, e.g., Tirole, 1989). I assume that the utility of consumer i with location (or bliss point) i can be written as:

$$\alpha - \beta(x_i^f - i)^2 - p_i^f,$$

where $x_i^f \in [0, 1]$ is the product of firm $f \in \{0, 1\}$ and p_i^f is its (potentially) customized price for this consumer. Throughout, we normalize the cost of production to zero for both firms (regardless of whether they produce a standardized or customized product).

Let us interpret the two firms as two different websites, which consumers visit in order to purchase the good in question. Before AI, firms cannot observe the type of consumer and I assume that they cannot offer several products to a consumer that visits their websites. Thus, they will have to offer standardized products. This description implies that, in terms of timing, they first choose their product, and then, after observing each other's product choice, they set prices. Because each firm is offering a standard product and cannot observe consumer type, it will also set the same price for all consumers. This makes the pre-AI game identical to a two-stage Hotelling model, in which firms first choose their product type (equivalent to their location) and then compete on prices. Throughout, I assume that:

$$5\beta < 4\alpha, \tag{1}$$

which is sufficient to ensure that the market is covered and that the firms will not act as local monopolies. As usual, I focus on subgame perfect equilibria, but with a slight abuse of terminology, I refer to these as “equilibria”.

It is straightforward to see that the unique equilibrium in this model, as in the baseline Hotelling model with quadratic transport costs, is maximal product differentiation (Tirole, 1989). In this setting, the two firms will offer products at the two ends of the line ($x^0 = 0$ and $x^1 = 1$) and set equilibrium prices given by

$p^0 = p^1 = \beta$, sharing the market equally. For future reference, I also note that, in this equilibrium, total firm profits are equal to $\Pi^{\text{pre-AI}} = \pi^0 + \pi^1 = \beta$, and consumer surplus is:

$$\begin{aligned} CS^{\text{pre-AI}} &= \alpha - 2\beta \int_0^{1/2} x^2 dx - \beta \\ &= \alpha - \frac{13}{12}\beta, \end{aligned}$$

where the first line of this expression uses the symmetry between the firms and consumers on the two sides of $1/2$.

After advances in AI, one of the firms (say firm 1) can use data from its previous customers (those with $i \geq 1/2$) to predict their type and customize their products and prices.² In particular, I assume that, in this post-AI environment, firm 1 can observe the type of any consumer $i \geq 1/2$ that visits its website and offers a customized bundle (x_i^1, p_i^1) to this consumer. For simplicity, let us assume that firm 0 cannot do so and also that firm 1 cannot simultaneously offer customized and standardized products. Now in equilibrium, firm 1 will offer each consumer with $i \geq 1/2$ a customized product $x_i^1 = i$. It will also charge higher prices. The exact form of the equilibrium depends on firm 0's product choice, which, given its inability to use the new AI technology, cannot be customized. It is straightforward to see that firm 0 will also change its product, because it no longer needs as much product differentiation (since firm 1 will be charging higher prices). The unique post-AI equilibrium is one in which firm 0 changes its standardized product to $x^0 = 1/4$. It then sets a price that makes the consumers that are farthest away from it indifferent between buying its product and not doing so, i.e.,³

$$p^0 = \alpha - \frac{\beta}{16}.$$

It is also straightforward to see that it is optimal for firm 1 to set:

$$p_i^1 = \alpha \text{ for all } i \geq \frac{1}{2},$$

thus capturing all of the consumer surplus from the consumers about whom it has data.⁴ In this equilibrium, we have $\Pi^{\text{post-AI}} = \alpha - \frac{\beta}{32} > \Pi^{\text{pre-AI}}$ (which is guaranteed by (1)), while consumer surplus is now:

$$\begin{aligned} CS^{\text{post-AI}} &= \frac{\alpha}{2} - 2\beta \int_0^{1/4} x^2 dx - \frac{1}{2} \left(\alpha - \frac{\beta}{16} \right) \\ &= \frac{1}{48}\beta. \end{aligned}$$

As a consequence, consumer surplus is much lower in this case. This can be seen most-clearly by considering the limit, where $\beta \rightarrow 0$, in which case the pre-AI consumer surplus is maximal (approaching α), while the post-AI consumer surplus becomes minimal (approaching zero). The negative impact of AI technologies on consumer surplus has two interrelated causes. First, firm 1 now uses its better prediction power to capture all the surplus from the consumers, even though it is, in principle, offering a better

product and could have increased consumer welfare. Second, given firm 1's more-aggressive pricing, firm 0 is also able to capture more profits, reducing even the surplus of consumers whose data is not being used.

It is worth noting that, in the present model, there is no intensive margin of consumer choice and the market is covered (under (1)). As a result, AI does not affect quantity purchased, and even when it reduces consumer welfare, it increases utilitarian welfare—in particular, greater customization reduces “transport costs”. The logic of the model highlights that this need not be the case when there is a quantity/intensive margin, because higher markups may inefficiently reduce quantity purchased. We will see, there are other reasons for inefficiency in similar environments.

Model for Behavioral Manipulation

I now present a model that shows how behavioral manipulation can take place, thanks to AI technologies and data about consumers. This model draws on Acemoglu, Makhdoumi et al. (2023), but instead of the continuous-time learning model developed in that paper, here I use a much simpler setting with two periods.

Consider the following dynamic setting with two periods ($t = 0, 1$) and no discounting. Consumers have a choice between two products, x^1 and x^2 , in both periods. They are initially uncertain about which one will yield higher utility. Suppose in particular that the true utility that a consumer gets from the product is either H or $L = 0$. The prior belief of individual i is that these two products will yield high utility for them is, respectively, q_i^1 and q_i^2 .

Both products are produced and offered by a digital platform, which again has zero cost of production and can offer personalized prices. To start with, the platform and the consumer have symmetric information, and thus the platform knows and shares the consumer's prior beliefs. Once an individual consumes one of the two products, she obtains an additional piece of information about her utility from the product. I assume, in particular, that if the true quality is H , the consumer receives a positive signal, denoted by σ^H (with probability 1). However, if the true quality is L , the product might still have deceptively high instantaneous utility (but long-term costs). Thus, with probability λ , the consumer will receive the high signal (or will receive the low signal, σ^L , with complementary probability). The most relevant interpretation of this “false-positive” signal is that there are certain types of products that (predictably) appear more attractive to consumers—for example, because of their tempting short-term benefits or because of hidden negative attributes.

Let us assume that the platform perfectly observes the consumer's experience with the product that she has consumed, and can change its pricing and product offering in the next period. The game ends at the end of the second period.

The pre-AI equilibrium takes a simple form. The platform will offer whichever product has higher q_i for consumer i , say product j , and will set the price:

$$p_{i,0}^j = q_i^j H,$$

capturing the full surplus. If the signal after consumption is σ^L , then in the next period, it will offer the other product, \tilde{j} , charging the lower price:

$$p_{i,1}^{\tilde{j}} = q_i^{\tilde{j}} H,$$

once again capturing the full surplus. If, on the other hand, the signal is σ^H , then, in the second period, the same product will be offered, but now there will be a higher price. I assume that, in the pre-AI environment, consumers have sufficient experience with such products and the signals they generate that they can correctly anticipate the likelihood of a high-quality product given a positive signal. As a result, the price following a positive signal will not increase all the way to H . Rather, it will be given by the expected value of the product's quality conditional on a positive signal. A simple use of Bayesian updating gives this price as:

$$\begin{aligned} p_{i,0}^j &= \frac{q_i^j}{q_i^j + (1 - q_i^j)\lambda} H \\ &= \Xi_i^j H, \end{aligned}$$

which again captures the full surplus from the consumer and also defines the expression Ξ_i^j , which is convenient for the remainder of this section. (There is no option-value term in prices, because the full surplus is being captured by the platform.)

The deployment of AI technologies, once again, improves the platform's ability to predict consumer preferences and behavior because it has access to the data from many similar consumers and their experiences with similar products. As pointed out above, I assume that this goes beyond what the consumer herself knows. In particular, I suppose that the platform can now forecast whether the consumer will receive the high signal from a truly low-quality product. More generally, this captures the ability of the platform to predict whether the individual will engage in an impulse purchase or make other choices with apparent short-term benefits and long-term costs.

Post-AI, therefore, the relevant state for consumer i at time $t = 0$ becomes $\left(\{q_i^j, \xi_i^j\}_{j=1,2}\right)$, where $\xi_i^j = 1$ designates the event that product j will generate a false-positive signal—which means that, in reality, it is low-quality for the consumer, but still the signal σ^H will be realized if the consumer purchases it. Critically, the platform observes ξ_i^j , but the consumer does not. Following Acemoglu, Makhdoumi et al. (2023), I assume that consumers are “semi-behavioral” and do not fully take into account that, in the post-AI world, the platform actually knows ξ_i^j . This captures the more-general economic force mentioned above: in the pre-AI, business-as-usual world, consumers may have learned from their repeated experiences and purchases, accurately estimating the relevant probabilities. The post-AI world is new and it is less plausible to expect that the consumers will immediately understand the superior information that the platform has acquired. Note also that although it can forecast ξ_i^j , the platform cannot observe consumer preferences perfectly, and thus, even when $\xi_i^j = 0$, it does not know whether the product itself is high or low-quality.

What does equilibrium look like in the post-AI world? The key observation is that, while before AI the platform's prediction was aligned with the prior of the household, this is no longer the case in the post-AI world. In particular, suppose that we have $q_i^1 > q_i^2$, but $\xi_i^1 = 0$, while $\xi_i^2 = 1$. Then the platform may prefer to offer the second product. To understand this choice, let us compute the profits from consumer i when the platform is using these two strategies. When it offers product 1, its total profits are:

$$\pi_i^1 = [q_i^1 + q_i^1 \Xi_i^1 + (1 - q_i^1) q_i^2] H.$$

This expression follows by noting that the platform is, at first, offering product 1 and charging q_i^1 . Because $\xi_i^1 = 0$, the consumer will receive a positive signal only if the product is truly high quality, which happens with probability q_i^1 . However, as indicated by the above discussion, in this case, the consumer does not know whether this is a false-positive or a truly high-quality product, and thus her valuation will be $\Xi_i^1 H$, which

explains the second term. Finally, if she receives a negative signal (probability $1 - q_i^1$), in the second period, the platform will offer product 2, charging q_i^2 .

On the other hand, when it initially offers product 2, the platform's profits are:

$$\pi_i^2 = (q_i^2 + \Xi_i^2)H,$$

because in this case there will be a positive signal for sure.

It is straightforward to see that offering the second good is more profitable for the platform when:

$$\frac{q_i^2}{q_i^2 + (1 - q_i^2)\lambda} > q_i^1(1 - q_i^2) + \frac{(q_i^1)^2}{q_i^1 + (1 - q_i^1)\lambda}. \quad (2)$$

Condition (2) is always satisfied whenever q_i^2 is sufficiently close to q_i^1 . Intuitively, the platform is willing to sacrifice a little bit of revenue in the first period for the certainty of getting the consumer to experience a good that it knows she will like—even though this is not a truly high-quality good.

What about consumer welfare? Perhaps paradoxically, in the first case, the consumer actually has a positive welfare. This is because, in this case, we have a high-quality product (and the platform indirectly recognizes this following the realization of signal σ^H , because it knows that $\xi_i^1 = 0$) and hence the positive signal can come only from a truly high-quality good. It is then straightforward to compute the user's welfare as:

$$\begin{aligned} U_i^1 &= q_i^1 \left(1 - \frac{q_i^1}{q_i^1 + (1 - q_i^1)\lambda} \right) H \\ &= \frac{q_i^1(1 - q_i^1)\lambda}{q_i^1 + (1 - q_i^1)\lambda} H > 0. \end{aligned}$$

This positive surplus may appear as the good side of our behavioral assumption. But the platform's second strategy shows the dark side. With this strategy, the consumer will overpay in the second period (because, given $\xi_i^2 = 1$, the product is, in reality, low-quality). Hence her utility is:

$$U_i^2 = -\frac{q_i^2}{q_i^2 + (1 - q_i^2)\lambda} H < 0.$$

Therefore, the ability of the platform to predict the consumer's preferences and vulnerabilities leads to a situation in which the platform can increase its profits by marketing low-quality products that are likely to appeal to the consumer in the short run.

Model for Labor Market Consequences

Suppose there is a single good in the economy, Y , whose production requires the combination of a measure 1 of tasks:

$$Y = \left(\int_{N-1}^N Y(z)^{\frac{\sigma-1}{\sigma}} dz \right)^{\frac{\sigma}{\sigma-1}}, \quad (3)$$

where $Y(z)$ denotes the output of task z and $\sigma \geq 0$ is the elasticity of substitution between tasks. The key economic decision is the allocation of tasks to factors. Let me focus on just two factors, capital and labor, and suppose, as in Acemoglu and Restrepo (2018, 2019), that each factor has task-specific productivities, determining its comparative advantage, and only tasks $z \leq I$ can be automated, given the current level of automation technology. This implies:

$$Y(z) = \begin{cases} A^L \gamma^L(z) l(z) + A^K \gamma^K(z) k(z) & \text{if } z \in [N-1, I] \\ A^L \gamma^L(z) l(z) & \text{if } z \in (I, N]. \end{cases}$$

Here $l(z)$ and $k(z)$ denote the total labor and capital allocated to producing task z . The state of technology is captured by the following: factor-augmenting terms, A^L and A^K , which increase the productivity of the relevant factor uniformly in all tasks; task-specific productivities, $\gamma^L(z)$ and $\gamma^K(z)$, which increase the productivity of a factor in a specific task; the threshold for tasks that are feasible to automate, I ; and the measure of new tasks, N . Let us assume that $\gamma^L(z)/\gamma^K(z)$ is increasing in z , so that labor has a *comparative advantage* in higher-indexed tasks. Suppose that capital is produced from the final good, with marginal cost R , which also gives its rental rate. Labor is inelastically supplied, with total supply given by L , and the equilibrium wage is denoted by w .

Acemoglu and Restrepo (2018, 2019) characterize the competitive equilibrium in this economy. Here I allow both competitive and rigid labor markets, by assuming that the wage cannot fall below some level \underline{w} . In this case, the equilibrium wage can be written as:

$$w = \max \{ \underline{w}, \text{MPL}(L) \},$$

where $\text{MPL}(L)$ is the marginal product of labor when there is full employment at L . The wage floor may be a consequence of regulations, such as minimum wages and union-imposed minima, or may result from other labor market imperfections, such as efficiency wage considerations.

Let us first focus on how the marginal product of labor changes (without any wage floor). Following Acemoglu and Restrepo (2018), this is given by:

$$\begin{aligned} \frac{\partial \ln \text{MPL}(L)}{\partial I} &= \frac{\partial \ln Y(L, K)}{\partial I} && \text{(Productivity effect)} \\ &+ \frac{1}{\sigma} \frac{1-s^L}{1-\Gamma(N, I)} \frac{\partial \ln \Gamma(N, I)}{\partial I} && \text{(Displacement effect)} \end{aligned} \quad (4)$$

where s^L denotes the labor share and $\Gamma(N, I) = \frac{\int_I^N \gamma^L(z)^{\sigma-1} dz}{\int_{N-1}^I \gamma^K(z)^{\sigma-1} dz + \int_I^N \gamma^L(z)^{\sigma-1} dz}$ is a measure of labor's *task content of production* (capturing what fraction of tasks are assigned to labor). In the special cases where $\sigma = 1$ or where $\gamma^K(z) = \gamma^L(z)$, we have $\Gamma(N, I) = N - I$, but more generally, $\Gamma(N, I)$ is always increasing in N and decreasing in I . The first line of (4) represents the *productivity effect*, which is driven by

the fact that automation reduces costs and, thus, increases productivity—by an amount equivalent to the cost difference between producing the marginal tasks by labor vs. capital:

$$\frac{\partial \ln Y(L,K)}{\partial I} = \frac{1}{\sigma-1} \left[\left(\frac{R}{A^K \gamma^K(I)} \right)^{1-\sigma} - \left(\frac{\text{MPL}(L)}{A^L \gamma^L(I)} \right)^{1-\sigma} \right].$$

The second line of (4) represents the *displacement effect* created by automation: as tasks are allocated away from labor and towards capital, the marginal product of labor declines. This displacement effect, which reduces the range of tasks employing workers, is always negative.

When we are at full employment, (4) gives the impact of automation on wages. When, instead, the wage floor at \underline{w} is binding, then the same effects now impact employment. The only differences are that on the left-hand side of (4), we now have the proportional change in employment, and on the right-hand side, $\text{MPL}(L)/A^L \gamma^L(I)$ is replaced by $\underline{w}/A^L \gamma^L(I)$.

Let us first consider full employment. What happens to the labor market equilibrium following additional automation? Equation (4) first shows that the labor share will always decline—because of the displacement effect, the wage will increase less than proportionately with productivity. Equally importantly, the wage level may fall as well. This is because the displacement effect can be larger than the productivity effect. In particular, when the productivity effect is small—for example, in the edge case where $\underline{w}/A^L \gamma^L(I) \approx R/A^K \gamma^K(I)$ —there is no productivity effect, and the equilibrium wage will necessarily decline. When there is more than one type of labor, the same argument also implies that, though the wage of some groups may increase, the average wage may still fall (see Acemoglu & Restrepo, 2021).

This framework further clarifies why automation could reduce employment. Suppose the wage floor \underline{w} is binding and, again, take the edge case where $\underline{w}/A^L \gamma^L(I) \approx R/A^K \gamma^K(I)$ so that the productivity effect is approximately zero. Then, automation necessarily reduces employment. By continuity, the same happens when the productivity effect is positive but not too large—the case that Acemoglu and Restrepo (2019) refer to as “so-so technologies”, because they are good enough to be adopted, but not so good as to have a meaningful impact on productivity.

Finally, let us consider the effects of the introduction of new (labor-intensive) tasks, captured by an increase in N . As in Acemoglu and Restrepo (2018, 2019), we now have:

$$\begin{aligned} \frac{\partial \ln \text{MPL}(L)}{\partial N} &= \frac{\partial \ln Y(L,K)}{\partial N} \text{ (Productivity effect)} \\ &+ \frac{1}{\sigma} \frac{1-s^L}{1-\Gamma(N,I)} \frac{\partial \ln \Gamma(N,I)}{\partial N}. \text{ (Reinstatement effect)} \end{aligned} \tag{5}$$

The productivity effect is positive as before and, in fact, introducing new tasks may increase productivity much more than automation. In addition, the reinstatement effect is also positive because, by creating new employment opportunities, it raises the labor share and increases employment and wages. When employment is suboptimally low because of labor market imperfections, new tasks and the reinstatement effect will be welfare-increasing.

Model for AI and Human Judgment

Suppose that there are two tasks to be performed, 1 and 2. Overall output in the economy is given by:

$$Y = \min \{y^1, y^2\},$$

where y^i is the output of task i , and the Leontief production function imposes that these tasks are strongly complementary.

Before AI, both tasks have to be performed by humans. Suppose that there is a measure 1 of humans, each with 2 units of time. Suppose also that, for reasons I will explain below, we start with humans allocating half of their time to task 1 and the other half to task 2. In this case, they have equal productivity in both tasks, which I normalize to 1. As a result, before AI, the economy produces a total of one unit of the final good. We can think of each worker as a “yeoman-producer”, consuming his or her production. Equivalently, we can think of this economy as consisting of firms hiring workers in a competitive labor market. In this case, the per-hour wage of each worker in each task will be $1/2$, ensuring that the entire output is paid to workers.

Now imagine that there are advances in AI algorithms that produce the first task at per-unit cost $c < 1/2$. This cost is paid in terms of the final good, and the fact that it is less than the equilibrium wage before AI implies that these algorithms are cost-saving and will be adopted. If this were the end of the story, AI would improve net output because workers would be reallocated from task 1 to task 2, enabling the economy to increase its total output.

However, suppose that there are also economies of scope: individuals learn from performing both tasks at the same time (and that is why the pre-AI allocation involved each worker devoting half of their time to each task). Suppose, in particular, that if a worker does not learn from task 1, his or her productivity in task 2 declines to $1 - \beta$. The post-AI allocation will involve all workers working in task 2 and whatever their total production is in this task, the economy will also produce exactly the same amount of task 1, using AI algorithms. As a result, net output in this economy will be:

$$2(1 - \beta) - \text{spending on AI} = 2(1 - \beta)(1 - c).$$

It can be verified that in the special case where there are no economies of scope ($\beta = 0$), $c < 1/2$ is sufficient for net output to increase—in particular, from 1 to $2(1 - c) > 1$. However, as soon as $\beta > 0$, this is no longer guaranteed. For example, when $c \approx 1/2$, even small economies of scope imply that the use of AI would reduce net output.

Would the market adopt AI when there are such economies of scope? The answer depends on the exact market structure. Because the adoption of AI technologies is associated with a finer division of labor, there is no guarantee that firms will internalize the economies of scope. For example, in the pre-AI equilibrium, the cost of one unit of task 1 is $1/2$. So, a new firm can enter with the AI technology and make profits in this equilibrium. The entry of these firms would then create a pecuniary externality, discouraging other workers from working on task 1. In particular, even though there are economies-of-scope benefits from performing task 1, workers may not be allocated to task 1 because the price of this task is now lower due to the use of AI technologies.⁵ This type of entry could then destroy the pre-AI equilibrium and drive the economy to the post-AI equilibrium characterized above, even when it is inefficient because β is large.

Model for AI and Excessive Monitoring

I now develop the main ideas discussed in the text using a model based on Acemoglu and Newman (2002). Consider a one-period economy consisting of a continuum of measure N of workers and a continuum of measure 1 of firm owners, each with a production function $AF(L_i)$ where L_i denotes the number of workers employed by firm i who exert effort (the alternative, exerting zero effort, leads to zero productivity). Firms are large, so the output of an individual worker is not observable. Instead, employers can directly monitor effort in order to determine whether an employee is exerting effort and being productive.

Specifically, as in Shapiro and Stiglitz (1984), a worker exerting effort is never mistakenly identified as a shirker, and a shirking worker is caught with probability $q_i = q(m_i)$ where m_i is the extent of monitoring per worker by firm i , with cost Cm_iL_i , where $C > 0$. Suppose that q is increasing, concave and differentiable with $q(0) = 0$ and $q(m) < 1$ for all m . Suppose also that, because of the limited liability constraints, workers cannot be paid a negative wage and will simply receive a zero wage. This implies a simple incentive compatibility constraint for workers,

$$w_i - e \geq (1 - q_i)w_i,$$

where e denotes the cost of effort. Rearranging this equation we obtain:

$$w_i \geq \frac{e}{q(m_i)}. \quad (6)$$

In addition, the firm has to respect the participation constraint,

$$w_i - e \geq \underline{u}, \quad (7)$$

where \underline{u} is the worker's *ex ante* reservation utility, given by what he could receive from another firm in this market.

The firm maximizes its profits, given by $\max_{w_i, L_i, q_i} \Pi = AF(L_i) - w_iL_i - Cm_iL_i$, subject to these two constraints. As shown in Acemoglu and Newman (2002), the solution to this problem takes a simple form because the incentive compatibility constraint always binds—if it did not, the firm would reduce monitoring, increasing its profits. By contrast, the participation constraint (7) may or may not bind.

The main result from this framework relevant for my discussion, is that regardless of whether the participation constraint is binding, the equilibrium never maximizes utilitarian social welfare (total surplus), given by $Y = AF(L) - CmL - eL$, because there is always too much monitoring. The intuition is straightforward: at the margin, monitoring is used to transfer rents from workers to firms, and is thus used excessively. Mathematically, this can be seen in the following way: start from the equilibrium and consider a small decline in monitoring, coupled with a small increase in the wage, such that (6) remains binding. This will have only a second-order impact on the firm's profits because the firm was already maximizing them. But it will have a first-order benefit for workers, whose wages will increase. Thus, a utilitarian social planner would like to increase wages and reduce monitoring. This is true, *a fortiori*, if we care about income inequality (presuming, of course, that firm owners are richer than workers).

How does AI affect things in this equilibrium? Suppose that before AI, there was an upper bound on monitoring, so that $m \leq \bar{m}$, and AI lifts this constraint. If the equilibrium level of monitoring without this constraint, m^* , is above \bar{m} , then improvements in AI will lead to higher monitoring. If, in addition, \bar{m} is not too low, then AI would reduce welfare; the qualifier that \bar{m} should not be too low is included because, if it were very low, the initial equilibrium could be very inefficient.

Model for Perils of Online Communication

Let us first focus on bilateral communication and consider a pre-AI world in which large-scale social media is not possible. Suppose that, in such a world, individuals communicate bilaterally in a social network. For simplicity, we can imagine a social network that takes the form of a directed line, in which we start with individual 0, who can communicate with individual 1, and so on, all the way up to the end of the line, individual n . Suppose that each individual has a piece of gossip which they can share with their neighbor, which will give utility v^G . In addition, the individual may have some news relevant to the political/social beliefs of the community. If she decides to share this news item, which then alters the beliefs of her neighbor, she also receives utility from such persuasion, as I specify below (for simplicity, I am ignoring potential utility benefits from non-neighbors indirectly changing their beliefs as this information is shared further). Because there are “channel constraints” (for example, limited ability to communicate), the individual can share either political news or gossip, but not both.

Specifically, suppose that the state of the world is 0 or 1 (Left or Right), and all individuals start with prior $\mu_0 = \frac{1}{2}$ about the underlying state being $s = 0$. Suppose that individual 0 receives a piece of news that shows that the underlying state is in fact $s = 0$. If she shares this information with her neighbor (individual 1) and her neighbor believes it, then her neighbor’s belief will also shift to $\mu = 1$. However, each individual is also concerned that some agents in society may have ulterior motives and try to convince them that the state is the opposite of the true state, s . Suppose the probability that each individual attaches to their neighbor being of this extremist type is q . Then, the posterior of individual 1 that the state is $s = 0$ after receiving this news will be:

$$\mu_1 = \frac{\frac{1}{2}}{\frac{1}{2} + \frac{1}{2}q} = \frac{1}{1+q} > \frac{1}{2}.$$

Finally, I assume that individuals receive additional utility from shifting their neighbor’s beliefs towards the truth (or her own belief), so the overall utility of individual i is:

$$v^G x_i^G + v^N |\mu_{i+1} - \mu_i|,$$

where $x_i^G = 1$ denotes whether this individual gossips, μ_i is her belief, and μ_{i+1} is her neighbor’s belief (given the line network).

Let us assume that $v^N > 2v^G$, such that, if an individual is convinced that her information will be believed and can thus shift her neighbor’s belief from $1/2$ to her views, she prefers to share the political information rather than gossip. Therefore, there exists \bar{q} such that if $q < \bar{q}$, the individual will share the political news. Let us suppose that in the real-world social network this inequality is satisfied.

In this scenario, individual 0 will start sharing the political news. This will convince individual 1, who will then attach a sufficiently high probability to the underlying state being $s = 0$. With the same argument, she would also like to convince her neighbor to the right, individual 2. If q is sufficiently small, then individual 2,

even when she is worried about the possibility that either individual 0 or individual 1 is an extremist, would still believe this information. In general, there will exist some $n(q)$, such that the political news will be communicated up to individual $n(q)$, and then after that there will be pure gossip on the network. For q sufficiently small, the entire network might share the political news.

Next suppose that we go to online communication through social media networks. This offers a larger network with less personal contact. As a result, it is plausible to assume that the probability that each agent attaches to the event that the person communicating with them is an extremist is now higher, say $q' > q$. If $q' > \bar{q}$, then in online communication, there will be no news exchange and all communication will be gossip. The further distortions that arise in the presence of multi-lateral (broadcast) communication are discussed in the text.

Model for Big Brother Effects

Here I briefly outline a simple model capturing some of these ideas. Consider a society consisting of $\lambda < 1/2$ elites and measure 1 of regular citizens. All citizens and all elites have the same economic preferences, but citizens are heterogeneous in terms of their cost of participating in protest activities, denoted by c_i for individual i . I assume that c is distributed uniformly over $[0, 1]$ in the population.

The political system is an imperfect democracy or an autocracy in which political choices are biased in favor of the preferences of the elite. In particular, suppose that there is a unique, one-dimensional policy, and the preferred policy choice of the citizens is 0, while the most-preferred policy of the elite is $p^E > 0$. Consider the following reduced-form political game. The elite decide a policy, p , and then protests take place. If some number, q , of the citizens protest and engage in civil disobedience, then there is probability $\pi(q)$ that the policy will switch from p to 0. With the complementary probability, the policy stays at p . This political structure ignores the influence of the citizens via democratic institutions, which is for simplicity. If this is incorporated (e.g., Acemoglu & Robinson, 2008) this would not affect the main message of the model presented here.

I also assume that the government can impose a punishment on those engaged in protests. Suppose, in particular, that the state has the capacity to detect at most a measure ψ of protesters. If the total amount of protest, q , is less than ψ , then all protesters are detected and can be punished. I assume that the punishment imposed on protesters is a constant, Γ , independent of the number of protesters.

The two key economic decisions are, therefore, policy choice by elites and protests by citizens. Let me first describe the utility of the citizens. Suppose that when the policy choice of elites is p and there are q protesters in total, individual i has the following utility as a function of her protest decision $x_i \in \{0, 1\}$:

$$U_i^C(p, q, x_i) = \left[(v^C |p| - c_i) - \min \left\{ \frac{\psi}{q}, 1 \right\} \Gamma \right] x_i,$$

where, for centrality, I have ignored components of the utility that depend on policy choice but are independent of the individual's process decision. Intuitively, the utility from protesting is increasing in the distance between the actual policy and the bliss point of citizens, as captured by $v^C |p|$ (recall that their bliss point is at zero). In addition, the individual incurs the cost of participating in protests, given by c_i . The second term in square brackets captures the expected punishment from protesting, taking into account that, when $q \leq \psi$, protesters will be punished with probability 1.

Clearly, there exists a threshold value \bar{c} , such that only individuals with $c_i \leq \bar{c}$ will participate, and thus $q = \bar{c}$, since the distribution of c is uniform between 0 and 1.

Let us next turn to the elite's utility. Suppose that this is given by:

$$\begin{aligned} U^E(p, q) &= -\mathbb{E} [|\hat{p} - p^E|] \\ &= -(1 - \pi(q))v^E |\hat{p} - p^E| - \pi(q)v^E |p^E|, \end{aligned}$$

where \hat{p} denotes the realized policy and the expectation is over the uncertainty concerning whether protests will force the elites to change policy. As usual, the subgame perfect equilibrium can be solved by backward induction. In the second stage, \bar{c} is determined such that, given the policy choice of elites, p , we have:

$$(v^C |p| - \bar{c}(p)) - \min \left\{ \frac{\psi}{\bar{c}(p)}, 1 \right\} \Gamma = 0. \quad (8)$$

In general, there can be multiple equilibria in this stage, because this equation might have multiple solutions for $\bar{c}(p)$. Note, in particular, that its left-hand side may be non-monotonic. In what follows, I focus on the case in which it is monotonically decreasing, which ensures a unique equilibrium; if there were multiple equilibria, however, we could pick the one with the highest amount of protest, which will necessarily be one where the left-hand side is decreasing, yielding the same results. It is then straightforward to see that $\bar{c}(p)$ is increasing in p , meaning that a more pro-elite policy will induce more protest.

Now, turning to the elite's maximization, we first rewrite the elite's utility function, taking into account the reaction of the citizens to their policy choice:

$$U^E(p, \bar{c}(p)) = -(1 - \pi(\bar{c}(p))) |p - p^E| - \pi(\bar{c}(p)) |p^E|,$$

which simply substitutes $q = \bar{c}(p)$. We can, therefore, maximize elite utility by choosing the initial policy p . Taking into account that $p < p^E$, this maximization problem yields a standard first-order condition:

$$(1 - \pi(\bar{c}(p))) - \pi'(\bar{c}(p))\bar{c}'(p)(2p^E - p) = 0, \quad (9)$$

and under suitable assumptions, we can ensure that the second-order condition for maximization is satisfied. In this first-order condition, $\bar{c}'(p)$ is given from (8) and, under the assumption that $\psi < \bar{c}(p)$, it can be written as:

$$\bar{c}'(p) = \frac{v^C}{1 + \psi\Gamma/\bar{c}(p)^2}.$$

Next, consider the introduction of AI, modeled as an increase in ψ to some ψ' . From the previous expression, this increase will reduce $\bar{c}'(p)$ and, from (9), this reduction in $\bar{c}'(p)$ will increase p away from the citizenry's bliss point and towards the elite's preferences. Intuitively, AI-induced government monitoring of protests weakens citizens' collective ability to force the elites to make concessions, so the elite respond to the deployment of the AI technology by withdrawing concessions. As a result, AI makes policies less-responsive to the citizens' wishes and to the extent that these policies impact the distribution of resources, it will also tend to favor the elite's economic interests and increase inequality.

Model for Automation, Social Power and Democracy

Let us return to the model for Labor Market Consequences and, for simplicity, suppose that $N = 1$ and $\sigma = 1$ in the production function (3) and that labor markets are competitive (i.e., there is no wage floor). This implies that the equilibrium level of production, as a function of capital and labor, can be written as:

$$Y(K, L) = K^I L^{1-I},$$

and, thus, with competitive labor markets, the labor share is $s^L = 1 - I$ and:

$$w = (1 - I) \frac{Y(K, L)}{L}.$$

Note also that, in this case, the impact of automation on output can be written as:

$$\frac{\partial \ln Y(K, L)}{\partial I} = \ln K - \ln L.$$

Therefore, automation is output-increasing, when $\ln K > \ln L$, or $K > L$, which is also equivalent to the competitive rental rate of capital being less than the wage, ensuring that automation is cost-saving. Conversely, low-productivity (“so-so”) automation now corresponds to the case in which $K \approx L$.

Consider a political system, as in Acemoglu and Robinson (2006), where all capital owners (capitalists) are elites and all workers are non-elites, with no within-group heterogeneity. To start with, let us consider a nondemocratic regime in which the capitalists hold power or, alternatively, a democratic regime in which they have disproportionate power. For brevity, I am going to ignore any threat of revolution or protests along the lines of the models considered in Acemoglu and Robinson (2006) and will also abstract from the considerations discussed in the previous subsection.

Suppose, in addition, that there is a lump-sum tax on capitalists, which can be redistributed to workers, and let us denote the per-worker transfer on the basis of this by τ . Suppose that the workers have an aspiration for a level of net income, given by w^A , and if $w + \tau < w^A$, they withdraw their cooperation, and as a result, the effective productivity of labor declines from 1 to $\delta < 1$. In this reduced-form model, I interpret the transfer τ from the elite both as a measure of redistributive politics and also as a general concession to democratic politics—for example, allowing workers to have more voice or making less use of lobbying and other activities that distort democratic politics.

The key question is whether the elite will make the necessary transfers to convince workers to continue cooperating in workplaces. This boils down to the comparison of the following two options for the elite: redistribute via τ so that the aspirations of the workers are met, or make due with lower labor productivity as a result of broken cooperation. Let us write the payoffs to the elite from these two strategies. Suppose that there is no within-elite heterogeneity, so it is sufficient to look at the overall income of capital. When the capitalists choose to meet the aspiration constraint of workers, their payoff is:

$$U_1^K = K^I L^{1-I} - \max \{ (1 - I) K^I L^{1-I}, w^A L \},$$

where the first term in the max operator applies when the market wage is already greater than w^A , while the second term is when they have to make transfers in order to bring workers to this level. Clearly, the

relevant case for the discussion here is the latter, so I assume $(1 - I)K^I L^{1-I} < w^A L$, and thus:

$$U_1^K = K^I L^{1-I} - w^A L.$$

The alternative is to not make the necessary transfer, in which case the elite simply receive the market return to capital when the productivity of labor has been reduced to δ , i.e.,

$$U_2^K = K^I (\delta L)^{1-I}.$$

What is the effect of automation on the comparison of these two strategies? Differentiating the difference in their payoffs, $U_1^K - U_2^K$, with respect to I yields:

$$\frac{\partial(U_1^K - U_2^K)}{\partial I} < 0 \text{ if and only if } \ln K - \ln L + \frac{\ln \delta}{1 - \delta^{1-I}} < 0.$$

This expression has two immediate implications. First, if we have low-productivity automation, such that $\ln K \approx \ln L$, then, because $\ln \delta < 0$, automation always makes the second strategy of stopping redistribution and foregoing labor's cooperation more attractive. Intuitively, automation makes labor less-central for production and, thus, losing its cooperation becomes less costly for capital. The economic force going against this calculation is that when automation increases productivity (i.e., when $\ln K > \ln L$), the output loss due to lack of cooperation from labor becomes more costly. Second, however, we can also see that, for any K and L , there exists a threshold level I^* such that once $I > I^*$, the second strategy is preferred, even taking into account the productivity gains from automation. Therefore, automation makes cooperation from workers less important and, to the extent that securing this cooperation was an important part of the motivation for redistribution and democratic politics, automation may make the elites turn their backs on, or even become hostile to, democracy.

References

Abramowitz, Alan I. (2010). *The disappearing center: Engaged citizens, polarization, and American democracy*. Yale University Press.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Acemoglu, Daron. (2002). Technical change, inequality, and the labor market. *Journal of Economic Literature* 40(1), 7–72.

[Google Scholar](#) [WorldCat](#)

Acemoglu, Daron. (2019). It's good jobs stupid. *Economics for Inclusive Prosperity Research Brief*. June.

<https://econfip.org/policy-briefs/its-good-jobs-stupid>.

[Google Scholar](#) [WorldCat](#)

Acemoglu, Daron. (2021). AI's future doesn't have to be dystopian. *Boston Review*, May 20.

<https://www.bostonreview.net/forum/ais-future-doesnt-have-to-be-dystopian>.

[Google Scholar](#) [WorldCat](#)

Acemoglu, Daron, & Autor, David H. (2011). Skills, tasks and technologies: Implications for employment and earnings. In O. Ashenfelter and D. Card (Eds.), *Handbook of labor economics, volume 4* (pp. 1043–1171). Elsevier.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Acemoglu, Daron, & Johnson, Simon. (2023). *Power and progress: Our thousand-year struggle over technology and prosperity*. PublicAffairs.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Acemoglu, Daron, & Newman, Andrew F. (2002). The labor market and corporate structure. *European Economic Review* 46(10), 1733–1756.

[Google Scholar](#) [WorldCat](#)

Acemoglu, Daron, & Pischke, Jorn-Steffen. (1999). The structure of wages and investment in general training. *Journal of Political Economy* 107(3), 539–572.

[Google Scholar](#) [WorldCat](#)

Acemoglu, Daron, & Restrepo, Pascual. (2018). The race between man and machine: Implications of technology for growth, factor shares and employment. *American Economic Review* 108(6), 1488–1542.

[Google Scholar](#) [WorldCat](#)

Acemoglu, Daron, & Restrepo, Pascual. (2019). Automation and new tasks: How technology changes labor demand. *Journal of Economic Perspectives* 33(2), 3–30.

[Google Scholar](#) [WorldCat](#)

Acemoglu, Daron, & Restrepo, Pascual. (2021). Tasks, automation and the rise in US wage inequality. NBER Working Paper No. 28920.

Acemoglu, Daron, & Robinson, James A. (2000). Why did the west extend the franchise? Democracy, inequality and growth in historical perspective. *Quarterly Journal of Economics* 115(4), 1167–1199.

[Google Scholar](#) [WorldCat](#)

Acemoglu, Daron, & Robinson, James A. (2006). *Economic origins of dictatorship and democracy*. Cambridge University Press.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Acemoglu, Daron, & Robinson, James A. (2008). Persistence of power, elites and institutions. *American Economic Review* 98(1), 267–293.

[Google Scholar](#) [WorldCat](#)

Acemoglu, Daron, & Robinson, James A. (2012). *Why nations fail: The origins of power, prosperity, and poverty*. Crown Business.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Acemoglu, Daron, & Robinson, James A. (2019). *The narrow corridor: States, societies, and the fate of liberty*. Penguin Press.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Acemoglu, Daron, Autor, David H., Hazell, Jonathon, & Restrepo, Pascual. (2021). AI and jobs: Evidence from online vacancies. NBER Working Paper No. 28257.

Acemoglu, Daron, Kleinberg, Jon, & Mullainathan, Sendhil. (2022). Problems of online communication. Work in progress.

Acemoglu, Daron, Makhdoumi, Ali, Malekian, Azarakhsh, & Ozdaglar, Asu. (2022). Too much data: Prices and inefficiencies in data markets. *American Economic Journal: Micro* 14(4), 218–256.

[WorldCat](#)

Acemoglu, Daron, Makhdoumi, Ali, Malekian, Azarakhsh, & Ozdaglar, Asu. (2023). *A model of behavioral manipulation*. Work in progress.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Acemoglu, Daron, Manera, Andrea, & Restrepo, Pascual. (2020). Does the US tax code favor automation? *Brookings Papers on Economic Activity* 2020(1), 231–285.

[Google Scholar](#) [WorldCat](#)

Acemoglu, Daron, Ozdaglar, Asu, & Siderius, James. (2021). Misinformation: Strategic sharing, homophily and endogenous echo chambers. NBER Working Paper No. 28884.

Agarwal, Ajay, Gans, Joshua S., & Goldfarb, Avi. (2018). *Prediction machines: The simple economics of artificial intelligence*. Harvard Business Review Press.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Allcott, Hunt, & Gentzkow, Matthew. (2017). Social media and fake news in the 2016 election. *Journal of Economic Perspectives* 31, 211–236.

[Google Scholar](#) [WorldCat](#)

Applebaum, Ann. (2020). *Twilight of democracy: The seductive lure of authoritarianism*. Signal.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Athey, Susan, Catalini, C., & Tucker, Catherine. (2017). The digital privacy paradox: Small money, small costs, small talk. NBER Working Paper No. 23488.

Autor, David H. (2014). Skills, education and the rise of earnings inequality among the other 99 percent. *Science* 344(6186), 843–851.

[Google Scholar](#) [WorldCat](#)

Autor, David H., Dorn, David, & Hanson, Gordon H. (2013). The China syndrome: Local labor market effects of import competition in the United States. *American Economic Review* 103(6), 2121–2168.

[Google Scholar](#) [WorldCat](#)

Autor, David H., Levy, Frank, & Murnane, Richard J. (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics* 118(4): 1279–1333.

[Google Scholar](#) [WorldCat](#)

Balwit, Avital, & Korinek, Anton. (2023). Aligned with whom? Direct and social goals for AI systems. In J. Bullock (Ed.), *The Oxford handbook of AI governance*. Oxford University Press.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Bergemann, Dirk, Bonatti, Alessandro, & Gan, Tan. (2021). *Markets for information*. Yale mimeo.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Boix, Carles. (2023). AI and the economic and informational foundations of democracy. In J. Bullock (Ed.), *The Oxford handbook of AI governance*. Oxford University Press.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Bostrom, Nick. (2014). *Superintelligence*. Danod.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Brynjolfsson, Erik, & McAfee, Andrew. (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. W. W. Norton & Company.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Budd, Christopher, Harris, Christopher, & Vickers, John. (1993). A model of the evolution of duopoly: Does the asymmetry between firms tend to increase or decrease? *Review of Economic Studies* 60(3), 543–573.

[Google Scholar](#) [WorldCat](#)

Choi, J. P., Jeon, D.-S., & Kim, B.-C. (2019). Privacy and personal data collection with information externalities. *Journal of Public Economics* 173, 113–124.

[Google Scholar](#) [WorldCat](#)

Christian, Brian. (2019). *The alignment problem: Machine learning and human values*. W. W. Norton & Company.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Farboodi, Maryam, Mihet, R., Philippon, Thomas, & Veldkamp, Laura. (2019). Big data and firm dynamics. *American Economic Review: Papers and Proceedings* 109, 38–42.

[Google Scholar](#) [WorldCat](#)

Ford, Martin. (2015). *Rise of robots*. Basic Books.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Gurie, Sergei, Henry, Emeric, & Zhuravskaya, Ekaterina. (2022). Checking and sharing alt-facts. *American Economic Journal: Policy* 14(3), 55–86.

[Google Scholar](#) [WorldCat](#)

Hanson, Jon D., & Kysar, Douglas A. (1999). Taking behavioralism seriously: Some evidence of market manipulation. *New York University Law Review*, 74, 630.

[Google Scholar](#) [WorldCat](#)

Huntington, Samuel P. (1991). *The third wave: Democratization in the late twentieth century*. University of Oklahoma Press.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Jones, Charles I., & Tonetti, Christopher. (2020). Non-rivalry in the economics of data. *American Economic Review* 110(9), 2819–2858.

[Google Scholar](#) [WorldCat](#)

Judis, John B. (2016). *The populist explosion: How the great recession transformed American and European politics*. Columbia Global Reports.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Kleinberg, J., Lakkaraju, H., Leskovec, J., Ludwig, J., & Mullainathan, S. (2018). Human decisions and machine predictions. *Quarterly Journal of Economics* 133(1), 237–293.

[Google Scholar](#) [WorldCat](#)

Klemperer, Paul. (1995). Competition when consumers have switching costs: An overview with applications to industrial organization, macroeconomics and international trade. *Review of Economic Studies* 62(4), 515–539.

[Google Scholar](#) [WorldCat](#)

Klinova, Katya. (2023). Governing AI to advance shared prosperity. In J. Bullock (Ed.), *The Oxford handbook of AI governance*. Oxford University Press.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Lanier, Jaron. (2018). *Ten arguments for deleting your social media accounts right now*. Hoffmann and Co.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Lee, David S. (1999). Wage inequality in the United States during the 1980s: Rising dispersion or falling minimum wage? *Quarterly Journal of Economics* 114(3), 977–1023.

[Google Scholar](#) [WorldCat](#)

Lee, Kai-Fu. (2018). *AI superpowers: China, Silicon Valley, and the new world order*. Houghton Mifflin Harcourt.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Levitsky, Steven, & Ziblatt, Daniel. (2018). *How democracies die*. Crown.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Levy, Ro'ee. (2021). Social media, news consumption, and polarization: Evidence from a field experiment. *American Economic Review* 111(3), 831–870.

[WorldCat](#)

MacCarthy, M. (2011). New directions and privacy: Disclosure, unfairness and externalities. *I/S* 425(2011).

<https://ssrn.com/abstract=3093301>.

[Google Scholar](#) [WorldCat](#)

Marantz, Andrew. (2020). *Antisocial: Online extremists, techno-utopians and the hi-jacking of the American conversation*. Penguin.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Markoff, John. (1996). *Waves of democracy: Social movements and political change*. Pine Forge Press.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Mishra, P. (2017). *Age of anger: A history of the present*. Macmillan.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Mosleh, Mohsen, Martel, Cameron, Eckles, Dean, & Rand, David G. (2021). Shared partisanship dramatically increases the social tie formation in a Twitter field experiment. *Proceedings of the National Academy of Sciences* 118(7), 1–3.

[Google Scholar](#) [WorldCat](#)

Neapolitan, Richard E., & Jiang, Xia. (2018). *Artificial intelligence: With an introduction to machine learning*. 2nd ed. Chapman and Hall/CRC.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Pasquale, Frank. (2015). *The black box society: The secret algorithms that control money and information*. Harvard University Press.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Posner, Eric A. (2002). *Law and social norms*. Harvard University Press.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Posner, Eric A., & Weyl, E. Glen. (2019). *Radical markets*. Princeton University Press.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Russell, Stuart J. (2019). *Human compatible: Artificial intelligence and the problem of control*. Penguin Press.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Russell, Stuart J., & Norvig, Peter. (2009). *Artificial intelligence: A modern approach*. 3rd ed. Prentice Hall.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Sandel, Michael J. (2020). *The tyranny of merit: What's become of the common good?* Penguin Press.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Shapiro, Carl, & Stiglitz, Joseph E. (1984). Equilibrium unemployment as worker discipline device. *American Economic Review* 74(3), 433–444.

[Google Scholar](#) [WorldCat](#)

Simonite, Tom. (2020). Algorithms were supposed to fix the bail system. They haven't. *Wired*.

<https://www.wired.com/story/algorithms-supposed-fix-bail-system-they-havent>.

Snyder, Timothy. (2017). *On tyranny: Twenty lessons from the twentieth century*. Tim Duggan Books.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Sunstein, Cass. (2001). *Republic.com*. Princeton University Press.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Thompson, Derek. (2019). Should we be afraid of AI in the criminal-justice system? *The Atlantic*.

<https://www.theatlantic.com/ideas/archive/2019/06/should-we-be-afraid-of-ai-in-the-criminal-justice-system/592084/>.

Tirole, Jean. (1989). *Industrial organization*. MIT Press.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Tirole, Jean. (2021). Digital dystopia. *American Economic Review* 111(6), 2007–2048.

[Google Scholar](#) [WorldCat](#)

Varian, Hal R. (2009). Economic aspects of personal privacy. In W. Lehr, and L. Pupillo (Eds.), *Internet Policy and Economics* (pp. 101–109). Springer. https://doi.org/10.1007/b104899_7.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Vosoughi, Soroush, Roy, Deb, & Aral, Sinan. (2018). The spread of true and false news online. *Science* 359, 1146–1151.

[Google Scholar](#) [WorldCat](#)

West, Darrell M. (2018). *The future of work: Robots, AI and automation*. Brookings Institution Press.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Zuboff, Shoshana. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. Profile Books.

[Google Scholar](#) [Google Preview](#) [WorldCat](#) [COPAC](#)

Notes

- * I am grateful to many co-authors who have contributed to my thinking on these topics and on whose work I have heavily relied in this essay. They include David Autor, Jonathon Hazell, Simon Johnson, Jon Kleinberg, Anton Korinek, Ali Makhdoumi, Azarakhsh Malekian, Andrea Manera, Sendhil Mullainathan, Andrew Newman, Asu Ozdaglar, Pascual Restrepo and James Siderius. I am grateful to David Autor, Lauren Fahey, Austin Lentsch, Vincent Rollet, James Siderius and Glen Weyl for comments.

- 1 The field of “AI” today is dominated by the suite of current artificial intelligence technologies and approaches, mostly based on statistical pattern recognition, machine learning, and big data methods. The potential harms of AI I discuss in this paper are relevant for and motivated by these approaches. Nevertheless, I will also emphasize that “AI” should be thought of as a broad technological platform, precisely because the general aspiration to produce “machine intelligence” includes efforts to improve machines in order to complement humans, create new tasks and services, and generate novel communication and collaboration possibilities.
- 2 More generally, the fact that AI-intensive firms are using data from and customizing products to their existing customers introduces intertemporal linkages, which could create lock-in effects and rich-get-richer dynamics, as in the switching cost and dynamic oligopoly literatures, such as in Klemperer (1995) and Budd, Harris and Vickers (1993).
- 3 Firm 0 could offer a lower price and steal some customers from firm 1, but it can be verified that this would lead to lower profits.
- 4 If we had allowed this firm to also market a standardized product, it would additionally compete for consumers $i < 1/2$, about whom it has no data. Our assumption rules out this possibility.
- 5 There are some market structures and pricing schemes that may prevent the adoption of AI technologies when they are inefficient in this case. For example, if workers can take a very low or even negative wage in order to work in task 1 (so as to increase their productivity in task 2), this may outweigh the cost advantage of new firms that enter and specialize in using AI for task 1. The issues are similar to the ones that arise in the context of firm-sponsored general training, and as in that case, labor and credit market imperfections would typically preclude the possibility that workers fully pay for all of the benefits that they receive by taking wage cuts (see Acemoglu and Pischke, 1999).

Notes

- * My belief in this precautionary regulatory principle is also the reason why I became one of the signatories of the call to halt the training of large language models for six months, circulated here <https://futureoflife.org/open-letter/pause-giant-ai-experiments/>.