3 Inductive Knowledge and Digital Taylorism

Introduction

Amazon's semiautomated warehouses are modern marvels, demonstrating both the emerging capacities and the limits of robotics and artificial intelligence (AI). In classic warehouses, goods are kept on rows of shelves, and workers roam among those shelves on foot or in a vehicle to grab goods and prepare them for shipping. For a time, Amazon sought to develop warehouse robots to directly replace such human workers, but that turned out to be impossible due to the limits of robotic hands and of humanoid robots generally.[1] So the company redesigned its warehouses and labor practices around robots' actual capabilities. Amazon worked with Kiva, a robotics firm that it later purchased, to develop robots that carry shelves to two groups of workers. One set, known as "stowers," place goods in numerically coded bins on those shelves. Amazon then uses video cameras trained on stowers' workstations, together with image recognition algorithms, to document where particular goods were placed. Kiva robots then move the shelves back to storage until needed. When an order arrives, a robot will bring the shelf to one of the second set of workers, known as "pickers," who locate ordered items on shelves, grab them, and put them into plastic bins.[2]

The Kiva robots' movement of shelves is uncanny, and even beautiful in a way. The robots and shelves are physically separated from workers by reinforced chain-link fence. This is necessary to minimize dangers, since the robots have limited abilities to sense their environment. They move quickly and constantly, guided by barcodes on the floor, often driving right toward one another and then pivoting at the last moment and moving ninety degrees in another direction. Their movements are directed and optimized by suites of algorithms that ascertain where each robot is and where it needs to go, with the overarching goal of minimizing the time to fill orders.[3] As a result, a sort of alien intelligence is immanent here. The robots function as a team or network, but each is executing instructions sent to it from a coordination node and does not consider what other robots are doing. In that sense, the robots fulfill a task once performed by humans, but in a completely new manner, and while moving and relating to each other in ways that humans never would. As this chapter explains in detail, the Kiva robots exemplify how a great deal of automation occurs today. Today's robots are rarely anthropomorphic and do not perform tasks in the same way that humans previously did. Meanwhile, their installation often leads to increased consumer demand (and therefore labor demand), and therefore does not threaten massive worker displacement.

Yet there is a dark side to this story: many human workers in such warehouses are little more than production inputs, and their work lives are anything but sublime. A picker's entire job involves grabbing goods off of shelves as quickly as possible and placing them in bins for further processing. Each individual picker works in a phone booth-sized area, where they may be unable to make eye contact with others and where they are constantly monitored by video cameras and image-recognition algorithms.[4] A large digital stopwatch in front of each picker shows how long she or he is taking to perform each task, enforcing time discipline down to the second. Once a picker has filled a bin with goods, it is sent via a network of automated conveyor belts to a third group of workers known as "packers." They spend just fifteen seconds on each order, sealing boxes with tape that is automatically dispensed at the right length and affixing a barcode to the package. The boxes then move to a sorting station, where another machine scans that barcode, determines which shipping method or company to use based on its destination and other data, and sorts the packages accordingly. Many packages are later delivered by contractors whose workers are subject to similarly pervasive surveillance and time discipline.[5]

While warehouse jobs pay relatively well—Amazon raised starting wages to $15 in 2018 following public pressure and nascent worker organizing efforts, and to $18 in 2021[6]—they are physically and mentally very demanding. As a *New York Times* reporter observed, "Unlike pickers in manual warehouses," who walk among shelves to find goods, "the pickers [at a semiautomated warehouse] have almost no relief from plucking goods off shelves, other than their breaks."[7] Since each job is exceptionally repetitive and involves virtually no teamwork, workers can be trained quickly to act as stowers, pickers, or packers. Before COVID, due to spikes and valleys in order volume, Amazon hired and trained more workers than it needed at any given time and developed a scheduling system to blast out requests for workers to take or give up shifts at all hours of the day. While many workers have a regular schedule, they also monitor shift requests so they can put in extra hours or get time off.[8] Vending machines just off the warehouse floor supply ear protection, gloves, and the Advil that many workers take regularly to alleviate inflammation from repetitive stress on joints or acute injuries from lifting heavier items. When workers do not perform rapidly enough, or they take a bathroom break without preclearance, algorithmic monitoring systems may report as much to managers, sometimes even recommending termination.[9] As noted in this book's introduction, those sorts of practices were among the issues that motivated Amazon workers in Staten Island to unionize.

Amazon's efforts to seamlessly integrate robotics, AI, pervasive surveillance, and deskilled human labor exemplify an overall sociotechnical system that this book calls "digital Taylorism."[10] The reference is to early-twentieth century Taylorism, a system of scientific management that established managerial control over the labor process. Like its forebear, digital Taylorism involves intertwined processes of automation and intensified surveillance. Where possible, companies use algorithms and robots to perform tasks once performed by line-level workers—though typically not in the anthropomorphic manner envisioned by Hollywood writers. Rather, companies aim to break jobs "into discrete, rationalized, low-skill tasks," some of which are automated and others that can be performed by workers with little specialized training.[11] Then, regardless of which sorts of tasks workers are performing, companies use new data-processing mechanisms to assign tasks, to schedule and to oversee workers, and to discipline them. Following emerging usage, the book calls that latter set of practices "algorithmic management."[12]

As the discussion in this chapter also shows, these processes fuel one another: through intensive surveillance and algorithmic management, companies can often extract, formalize, and encase employees' knowledge and know-how, sometimes even claiming intellectual property (IP) rights in that expertise. These new capabilities in turn reinforce patterns toward market concentration by leading firms, including in low-wage sectors, which chapter 5 discusses in more detail. Like its forebear, then, digital Taylorism aspires to create and enforce a division of labor in which managerial authority is centralized, while line-level workers are hired to perform tasks that require uniquely human skills such as fine motor control and situational judgment. Also like its forebear, digital Taylorism makes it much harder for workers to organize and therefore helps keep labor costs down. Indeed, that is part of its attraction to companies.

Yet digital Taylorism differs from its namesake in important respects, which this chapter also discusses. With new surveillance and data-processing devices, companies are able to oversee workers' performance in vastly more detail than before. Even when thousands of workers labor side by side, they stand at the end of surveillance spokes that extend out from a corporate nerve center. Companies are also using data-processing tools in substantively new ways, which this chapter also explores. In particular, AI today frequently operates by drawing inferences from very large data sets, spotting patterns that humans never could. This leads to what the chapter calls "inductive" knowledge, which is different in kind both from the highly formal productive knowledge sought by classical Taylorism and from the tacit or embodied knowledge that characterized craft production. These distinctions among inductive, formal, and tacit knowledge illustrate both the promise and limits of contemporary automation efforts—and how companies are using technology to reshape class relations.

This chapter says less about law than the others, in part because companies' authority to take such steps is so well established today that it no longer gets litigated. Labor law, as it evolved under neoliberalism, nevertheless operates constantly in the background. The recent expansion of employers' property rights has given them near-plenary legal authority to install monitoring devices and reshape workplace practices around them. The contractualization of employment has made it easier to terminate workers for protesting those efforts or otherwise challenging managerial authority. A helpful way to think about the role of labor law here is that it establishes a set of *background* entitlements that are so powerful that companies can often use them to erode *foreground* obligations. Digital Taylorism, therefore, puts pressure on workers' statutory labor rights, especially under wage and hour laws and collective bargaining laws.[13] The following two chapters build on this account and discuss legal issues in more detail, in part because the legal issues they address—discrimination protections, privacy rights outside the workplace, workers' rights of association, and the legal definition of employment—are more in flux.

Section 3.1 discusses the three forms of knowledge in production and how each one is a subject of class-based politics. Section 3.2 then summarizes the promise and limits of contemporary automation efforts. Section 3.3 addresses algorithmic management, showing how major companies today combine it with task automation as interlocking and complementary strategies.

3.1 The Rise of Inductive Knowledge

As discussed in chapter 1, companies can use new workplace technologies for two quite different ends: to enhance labor productivity (i.e., to enable workers to generate more output per unit of input) and to augment their power over workers and therefore limit labor costs. As also noted in chapter 1, control over workplace information has long been an important form of workplace power and a subject of pitched battles between workers and companies. For some time now, new information and communications technologies have been deepening and sharpening companies' capacity to keep tabs on what happens in the workplace and its contemporary equivalents, like the delivery driver's vehicle or the client's or the worker's home. More recently, those technologies have begun to generate new ways of seeking and knowing the world, which companies are integrating at scale into management processes. This section discusses the three forms of knowledge involved in production planning and execution today—formal, tacit, and inductive—and how companies and workers deploy or protect each to advance their interests.

Formal versus tacit knowledge Classical Taylorism and its limits were defined by the separation of conception and action. This reflected an important distinction between two sorts of knowledge, both of which have always been important to modern production. The first is what this book calls "formal" knowledge. It is highly abstract and ordered, ideally susceptible to expression in mathematical equations or detailed engineering specifications. Formal knowledge can be described and transmitted with great specificity. But it also envisions the world of production and society from the top down. The second form is "tacit" knowledge. This is perhaps best captured by Michael Polanyi's observation that "[w]e know more than we can tell."[14]

Our tacit knowledge of how the world operates and how we operate within it enables us to perform many physical day-to-day tasks—walking, cracking an egg, and so on—that we could not possibly explain or codify. This is also the realm of social norms and customs, body language, and the way that a slight change in tone of voice can signal a great deal. The sort of "situation sense" developed over time by craft workers and experts involves a combination of formal and tacit knowledge. When an experienced lawyer confronts a new legal problem and has a rough sense of how it will play out, or when a doctor examines a patient and has an intuition about what is wrong, the lawyer and doctor are drawing on both forms of knowledge.

The relationship between formal and tacit knowledge poses fascinating philosophical and scientific questions, which are directly relevant to issues of automation and the future of work. For example, in some understandings of cognition, higher-level reasoning faculties—formal knowledge—are actually dependent on and inseparable from our embodied existence in the world.[15] The tension between those forms of knowledge also helps to illustrate some challenges of classic Taylorism and automation. In essence, machines have long been very good at doing the same thing again and again in the same way. In that sense, factory machinery basically encodes formal knowledge. Humans, meanwhile, are very good at exercising situational and social judgment, performing fine motor actions, and adapting tools to new circumstances, all of which require tacit knowledge.[16] Historically, Taylorism sought to formalize some of workers' tacit knowledge so that managers could exert control over production. But there were always limits to that process, which reflected the limits of formal knowledge itself. James C. Scott, for example, argues that many social disasters of the nineteenth and twentieth centuries resulted from rulers' "high modernism," or their extreme faith in the ability of technical rationality to reorder human affairs.[17] Collective farms, tree plantations, massive modernist public housing projects, and many other efforts failed spectacularly, Scott argues, due to their disregard of citizens' social intelligence and embedded tacit knowledge. One of Scott's go-to examples is the "work-to-rule" strike, in which workers follow engineers' or managers' commands to the letter rather than using all the workarounds and flexible judgments that characterize actual social processes—and thereby bring a factory or office to a standstill.[18]

Inductive knowledge Today, there is a third emergent form of knowledge—referred to here as "inductive" knowledge—which has developed iteratively and become more important as new data-gathering and processing devices have proliferated across the lifeworld.[19] Those include cellular phones, internet-connected computers, and all manner of digital devices, including video cameras, payment processing systems, and networked appliances. Many or most of those devices generate data about users' activities and whereabouts, which is then collected and analyzed by tech companies of all stripes. As other scholars have illustrated, companies across the consumer space have sought to claim property-like rights in such data, to analyze and utilize that data to discern customers' preferences and behavior, and then over time to *shape* customer demand and other facets of social behavior.[20] For example, companies have often sought not just to gather and hold extensive user data, but also to aggregate it with data from other sources and then to draw various inferences about individuals even when the underlying data sets are spotty or anonymized.[21] An entire industry of data brokers has arisen to gather, clean up, and process such data,[22] and to sort consumers into statistically based categories like "Affluent Baby Boomer" or "Rural Everlasting," in order to market to them.[23] Chapter 4 discusses data-aggregation efforts in more detail, insofar as they intersect with contested questions of employee privacy, equal opportunity, and worker self-organization.

A related form of data analytics builds on data aggregation and statistical analysis, essentially taking it to scale and generating new ways of seeing and knowing the world. It is best reflected in the subfield of AI known as "machine learning," which differs substantially from earlier efforts to develop AI. For some time, computer scientists sought to develop AI through encoded formal knowledge, developing rules that would guide algorithms through performing certain tasks. That approach worked for tasks such as playing chess but then stalled, due in part to the enormous complexity of many fields of human endeavor.[24] For example, efforts to develop translation programs by encoding grammatical rules in a formal "if-then" format ultimately failed. Human language is too complicated, too nuanced, too situationally specific to be captured that way. As a result, researchers were constantly plugging holes in their algorithms through ad hoc patches, while struggling to obtain accurate outcomes.[25]

Machine learning works differently: it draws advanced statistical inferences from large data sets.[26] While the technique is not entirely new, several papers in the early 2010s demonstrated how machine learning could be used for purposes of image recognition in ways that had seemed impossible in the past.[27] A relatively simple machine learning algorithm can be "trained" to determine whether a particular picture is of a dog or a cat.[28] Programmers would train it by uploading thousands of pictures of dogs and cats, appropriately labeled as such—the so-called training data—into the machine. The machine would then develop statistical correlations between the pixels in images labeled "dog" or "cat" and the outcomes "dog" and "cat," and programmers would adjust the various algorithms' inferences until the overall system could recognize dogs and cats accurately. When the data sets are large enough, the results can be remarkably precise, and the applications are extensive. Banks and other financial institutions use machine learning to process loan applications and to assist in fraud prevention by spotting irregular activity.[29] Machine learning can help determine whether particular moles are cancerous and help interpret radiological scans.[30] Google has also used machine learning in its search responses and to develop language translation programs that are remarkably good.[31]

But while machine learning often replicates the *outcomes* of human judgments fairly well, it does not replicate our reasoning processes. Instead, its underlying logic can be idiosyncratic, even baffling. Google's development of a suite of algorithms to play the game Go helps to illustrate this point. Unlike chess, which has many rules, Go has very few—it simply involves placing white and black tiles on a surface. The number of possible moves over time is several orders of magnitude larger than chess, however, which made it impossible to hard-code a machine to play the game. For related reasons, there are defined national and regional styles of play, analogous to dialects, since individuals have typically learned how to play in particular communities. So Google enabled a machine called DeepMind to train itself to play Go after being given some basic parameters.[32] During one match, DeepMind made several moves that no human player would have made because they went against all conventional wisdom and did not match any known style of play. At times, these moves cost the machine in the short term but paid off in the end.[33] One observer said that watching the match was like watching "an alien civilisation inventing its own mathematics."[34] Amazon's warehouse robots evince a similar uncanny logic.

Machine learning and related techniques are therefore "inductive and atheoretical."[35] They reveal patterns in large data sets that humans could never see, and then they classify individuals or items within those data sets into either predefined or discovered groups.[36] In the hiring context—discussed in the next chapter—such techniques may provide clues to an applicant's "future health, future productivity, or likely tenure with an employer."[37] As scholars have noted in other contexts, a major use case of such technologies involves discerning whether individuals who share some observable characteristic (x) will also tend to share some unobservable characteristic (y). As Salome Viljoen has put it, "a basic purpose of data production as a commercial enterprise is to relate people to one another based on relevant shared population features."[38] In part for that reason, as also discussed in the next chapter, it is difficult to fully grasp the harms of such practices through individualized conceptions of privacy—the harms are more social by their nature, since they lead to individuals losing opportunities on the basis of membership in statistical categories.[39]

There are many use cases for inductive learning technologies in the workplace and labor markets, as discussed later in this chapter. But only rarely can they replace human workers—or even do tasks performed by human workers—on a one-to-one basis. Rather, companies tend to utilize formal and inductive knowledge in combination and to rely on human workers to supply tacit knowledge. Reflecting historical practice, moreover, companies use formal and inductive knowledge both to enhance productivity and to augment their power over workers. As a result, ultra-advanced technologies coexist with low-wage, demobilized labor across large segments of today's economy.

3.2 Automation and Its Limits

In the late 2010s, these developments in AI generated widespread fear—bordering on panic—that a looming automation wave would lead to massive unemployment or even economic collapse. Tech leaders didn't hesitate to stoke those fears by describing their products as near-magical. To take one of many examples, Sundar Pichai, the chief executive officer of Alphabet (parent company of Google) said in 2020 that AI is "more profound than fire or electricity."[40] Many major magazines ran cover stories on the supposed automation threat, with titles like "Welcoming Our New Robot

Overlords" and "Learning to Love Our Robot Co-workers."[41] A widely discussed study by two Oxford researchers predicted that "about 47% of total US employment" is at high risk of automation.[42] Labor leaders and leading labor law professors asked whether and how we need to adapt to a world with much less work.[43] Some viewed automation as a mortal threat to workers and even society, while others viewed it as our best hope for liberation from toil.[44]

Such fears are not new. In the wake of industrialization and the growth of modern factories, which had already led to a sharp decline in agricultural employment due to tractors and other implements, John Maynard Keynes famously speculated that his grandchildren would be able to work a fifteen-hour week.[45] In the 1960s, companies' incorporation of computers and other advanced information technologies into their operations—then known as the "cybernation revolution"—led again to widespread fear of a looming automation wave.[46] In those past moments of sociotechnical change, however, work did not disappear, for two interrelated reasons. First, companies passed some productivity gains on to consumers through lower prices, which bolstered consumer demand, and therefore demand for workers to produce, distribute, and sell those goods. Second, companies developed new goods and services—leisure goods, safer and more efficient appliances and automobiles, and health-care products and services—to sell to workers and consumers whose lives had become more comfortable as a result of this technological progress. Workers who wanted to be able to purchase such goods were then willing—or required—to put in the long hours necessary to afford them. Taking the long view, productivity growth, together with the expansion and deepening of global markets, led over time to shifts in the composition of the labor market: first from agriculture to manufacturing, and then from manufacturing to services.

Many wondered, however, whether this time was different. After all, the tech sector *had* revolutionized various other parts of the economy in the preceding decade. Napster and other file-sharing services, and then YouTube, iTunes, and Netflix, transformed music and video distribution. The development of smartphones completely changed photography and the market for cameras and film. Facebook and Google altered the economics of news media by aggregating stories and capturing online ad revenue. The explosive growth of Uber and Lyft then suggested that tech companies were poised to alter economic sectors that involved heavy investment in *physical* technology. This narrative may also have taken root because journalists had seen their career prospects erode due to Facebook and Google's influence, and therefore may have been inclined to believe tech companies' promises to revolutionize physical work. Some companies also stoked automation fears for self-interested reasons. When fast food workers began demanding $15 an hour, industry groups argued that if companies were required to raise wages, they would implement kiosk-based ordering more quickly.[47]

Granted, from the standpoint of Silicon Valley, perhaps it was reasonable to think that machine learning would only grow ever more powerful. Those at the center of contemporary data-gathering and -processing networks were able to "see" across much greater distances, and in much greater detail, over the course of the 2010s. As more and more facets of social behavior became legible, why wouldn't the nerve structures of the digital economy evolve into something like intelligence? And if so, then wouldn't they continue to evolve? Such thinking led to various predictions of a looming "technological singularity," the point at which AI would be able to replicate all aspects of human intelligence and then improve at an exponential rate.[48] A prominent computer scientist argued that such a machine would be "the last thing we'll ever have to invent because, once we let it loose, it will go on to invent everything else that can be invented."[49] The rate of progress in all sorts of technologies would speed up, breaking through many of the engineering challenges that bedevil robotics today, and integrate advanced intelligence into those robotics. In some versions of the story, AI would then go rogue, seeking to dominate human society or even the galaxy, or become psychotic and seek to turn all known matter into paper clips or some such.[50]

This is science fiction—it is literally the plot of the *Terminator* and *Matrix* franchises. And in reality, such a revolution in AI would likely be necessary to displace today's service workforce.

The limits of machine learning and robotics The notion that inductive knowledge could generate a world without work suffered from several basic errors. For one thing, there was no hard evidence that automation was a major threat.[51] Nobody could ever point to an army of robots standing ready to displace huge numbers of workers, especially in services, and consumers' day-to-day experiences of physical technology are often underwhelming. As the economist Martha Gimbel has said, "Any time anyone tries to claim robots are coming to take our jobs I ask them how well their printer works and that usually ends the conversation."[52] What's more, if companies were installing robotics and related technologies in large numbers, that effect would surely be visible in the data on productivity growth over time, since workers as a whole would be generating more output per hour. But productivity growth has recently been as slow as at any time since World War II,[53] and productivity growth in the manufacturing sector—where task automation has historically been easiest—has been especially tepid.[54] Meanwhile, an automation wave would lead to significant increases in unemployment, but prior to COVID-19, joblessness was relatively low in the US.[55]

Predictions of a looming automation wave also disregarded the costs of robotics and other mobile physical devices like vehicles. Bluntly, they are not cheap. The price for Amazon's Kiva system is not publicly available, but a distribution center consultant estimated the cost for a typical warehouse of 50–100 robots as being between $2 million and $4 million, and a large warehouse of 500–1000 robots as being from $15 million to $20 million. As the consultant wrote, "these are some serious figures as far as distribution center capital investments are concerned."[56] Just as important, the unit economics of robotics are dramatically different from the unit economics of algorithms. App-based tech companies have grown rapidly in part because the marginal cost of producing new versions of an app or software is effectively zero. That is why file-sharing and related technologies posed such a devastating threat to the music and entertainment industries: once music and movies can be digitized, they can be shared essentially at will, and for free, which made it impossible for artists and legacy companies to profit from those products. Robotics and other semiautonomous machines, in contrast, involve physical technology. While mass production of robots would drive down costs, those costs will never approach zero due to the simple costs of materials such as metals and plastics, as well as of intermediate physical inputs like batteries and light detection and ranging (LIDAR) or other sensing systems. Now, labor isn't cheap either, and when labor shortages or worker organization drives up labor costs, companies will have greater incentives to economize through task automation. As a result, the impetus to automate is not going away. The point is rather that progress in physical automation presents economic challenges that are very different from progress in the gig economy, file sharing, online search, social networking, and other activities in which the marginal costs of adding new units or users can approach zero.

Predictions of a looming automation wave also significantly overstated the capacities of inductive learning, simply assuming that it could substitute for tacit and formal knowledge across many spheres of human behavior and action. For better or worse, it seems increasingly clear that machine learning is not a path to imminent artificial general intelligence, and therefore not to an army of robots who would displace line-level workers.[57] The basic problem is that machine learning relies on drawing statistical inferences from massive but also discrete data sets, often in laboratory conditions.[58] Programmers and commentators often reported progress in AI with reference to particular benchmarks—for example, in image recognition or language translation—without considering whether performance on those benchmarks, in laboratory conditions, actually measured progress toward creating generalizable analytical systems.[59]

When machines or programs move into the physical and human world, however, all sorts of new and unpredictable challenges emerge that cannot be solved through statistical analysis because—almost by definition—those challenges did not appear in the training data. For example, such systems have difficulty making contextual judgments about human and social affairs since they have no innate sense about the world.[60] Minor changes in a system's input layer, therefore, can lead it to fail, sometimes catastrophically.[61] A recent language program and model by OpenAI has made progress on issues such as these, and yet it still delivers strange and absurd results at times and clearly has no awareness of the social or real-world contexts for its conversations.[62] In 2020 a programmer showed that while it was impressively accurate in answering straightforward questions, it was stumped by questions that humans would recognize as absurd, such as "How many eyes does a blade of grass have?" The program's answer: "A blade of grass has one eye."[63]

These underlying problems have stymied progress in many fields, including autonomous vehicles. After Uber grew spectacularly, companies rushed into the space, with Google, Tesla, and General Motors (GM) all promising fully autonomous cars by 2019.[64] The implications for workers would be obvious: a prominent labor reporter predicted that self-driving cars could displace five million professional drivers, including truckers, cab drivers, and other delivery drivers.[65] The first company to develop and patent fully autonomous vehicles—or the critical technology for them—would also capture massive product market rents. But for several years now, companies in the sector have been trying to lower expectations.[66] The problem is that the real world constantly presents novel situations that do not map onto the algorithms' training data. Drivers need to react to sudden changes in weather, intoxicated people running into the road, or items flying off of other cars. In those circumstances, machine learning algorithms get stumped or respond in erratic ways. That sort of event helped cause one of Uber's self-driving cars to hit and kill a pedestrian in 2018, as the image-recognition devices misidentified the pedestrian and therefore did not respond in time.[67] Similarly, former Tesla employees have accused that company of repeatedly overstating its vehicles' autonomous capabilities, and the National Highway Traffic Safety Administration has investigated the company on related grounds.[68]

In other contexts, engineers can mitigate that problem by engineering social and physical environments so that machines have to process only a limited amount of information. In the industrial and warehouse settings, for example, engineers often place production robots in cages to isolate them from humans since they are heavy and very powerful and have difficulty sensing that people are around, and therefore can injure humans quite easily.[69] Those robots are then programmed to perform particular tasks, which involves formal rather than inductive knowledge. That approach isn't plausible for vehicles, however, since it would require redesigning our entire system of roads so that pedestrians cannot access them. Alternatively, engineers could program an autonomous vehicle to stop every time it is stumped by an object in the distance, but in that case, passengers will get impatient or carsick, and other drivers will get irritated.[70] What seems more likely going forward is that the technologies that have already been developed will be deployed in environments where environmental control is easier. Autonomous long-haul trucking seems plausible, therefore, but human drivers would still need to take the trucks off highways to their final destinations within cities or suburbs.

Those challenges compound for other sorts of robots. Promotional videos by the robotics company Boston Dynamics show humanoid and dog-like robots walking unattended, dancing, and even performing backflips.[71] But the company has admitted that they are either remotely controlled by humans or programmed in minute detail—encoding formal knowledge.[72]

Even industrial automation is far more difficult than many appreciate.[73] There are basic technical limitations in the design and strength of robotic arms and hands, which lead to a trade-off between strength and safety: robotic arms that are strong enough to perform many industrial and other tasks need to be heavy, which both increases energy costs and makes it challenging to deploy them alongside humans. Similarly, Apple ended up abandoning many automation efforts after investing heavily in a secret lab in California that brought together top robotics researchers in an effort to solve production challenges.[74] Those researchers learned, for example, that it is exceptionally difficult to design a robot to insert a small screw into an iPhone chassis, given the fine motor skills and vision required for the task.[75] The Kiva example that opened this chapter gives a much clearer picture of how mobile robots are being integrated into production: they are used to perform single tasks in controlled environments that are organized around their capabilities, and they are not anthropomorphic.

Implications for service workers These technical and financial limitations suggest that the majority of automation going forward will likely be slow and iterative, especially among service-sector workers. Consider the tasks performed by delivery drivers for a company such as FedEx or Amazon Flex. Those individuals need to drive, of course. Then, once arriving at a destination, they need to park, exit their vehicles, find the packages that they are planning to deliver, and walk up to a building, often across an uneven sidewalk. Once there, they need to determine whether packages can be left safely in a particular area, or whether a particular person is responsible and capable enough to take a package—for example, not a minor, an intruder, or an individual with dementia. Similarly, baristas don't just make a cappuccino, but also answer questions from customers, modify the drink based on a customer's random requests, and hand it to the customer. Supermarket clerks don't just put boxes on shelves—a task that is still remarkably difficult for robots—but also make strategic decisions about where to store excess inventory, spot hazards, and defuse conflicts among customers if necessary.[76] Nursing home aides and home-care workers don't just dispense medication—they also move patients through irregular physical spaces and determine whether a patient's grunts signify pain versus frustration. It is impossible to automate such jobs without anthropomorphic and extremely intelligent robots, and those robots are far in the future.

In each of these cases, companies continue to automate particular tasks, of course. The FedEx driver may receive navigation instructions from an app and track packages using a barcode scanner that sends information straight to the company's servers. In a prior era, the driver would have navigated without electronic assistance and filled out paperwork to track deliveries. Both developments also enhance productivity and are in some respects desirable for workers. They also enable closer surveillance, as discussed in section 3.3. On the consumer-facing side, cafes are integrating tablet- or app-based ordering systems, which also enable more extensive worker surveillance. The list goes on, but the point should be clear enough: automation is affecting work, but through many small iterative changes rather than the sudden displacement of whole categories of workers. How these trends will play out is fundamentally unknowable, especially given the massive disruptions of COVID. By the time this book is published, engineers and computer scientists will surely have made progress on some of the challenges noted here. Conversely, by that time, other companies may have decided to cut their losses and cease trying to automate particular tasks.

As a result, lawmakers have quite a bit of space to ensure decent work in the near future, and even to shape the course of automation processes themselves. Other scholars have begun to consider how to steer workplace automation in directions that complement rather than undermine human creativity and autonomy.[77] Chapter 6 discusses and builds on those efforts and argues that worker-protective reforms are a necessary component of such efforts. By raising labor costs, such protections can actually *encourage* automation, and thus a degree of worker displacement. And yet under the right institutional conditions, including robust mechanisms for worker voice and power, companies have incentives to collaborate with workers to unlock productivity gains.

Here, some comparative evidence may help illustrate the point. As alluded to in chapter 1, German manufacturers responded to technical change and globalization in the 1980s by focusing on high-wage, high-skill, high-value-added strategies.[78] Today, there is some evidence that German companies are more likely to use "co-bots," or robots that work alongside and complement factory workers, while their counterparts in the US typically seek to replace workers entirely.[79] In 2020, moreover, several economists documented that German companies subject to a stricter form of codetermination—the German system of worker representation, including on corporate supervisory boards—had higher capital intensity than companies subject to forms of codetermination that gave workers less power. This suggests that worker voice and power can encourage companies to pursue higher-productivity strategies.[80] An important goal for labor market and industrial policy in the US going forward is to generate similar virtuous cycles of productivity and wage growth, including in service sectors when possible.

3.3 Algorithmic Management

Given the limits of robotics and machine learning discussed thus far, companies today cannot avoid labor politics entirely. But the second component of digital Taylorism—algorithmic management—is giving companies powerful new means of reducing labor costs, even while employing massive numbers of workers. The overview of Amazon's labor practices in the introduction to this chapter illustrates. The company uses robots and algorithms to perform tasks that can be engineered in precise detail through a combination of formal and inductive logic, such as moving shelves of goods, directing pickers to grab particular goods, sorting packages, and affixing mailing labels to them. Amazon then hires workers to perform tasks requiring tacit knowledge and fine motor skills, such as stowing irregularly shaped goods on shelves and grabbing them back off those shelves. The company also deploys new surveillance devices and data-processing algorithms to surveil and manage those workers at scale. That system also involves task automation, but many of the tasks involved are cognitive rather than physical, and were once carried out by managers. Line-level workers, then, are cogs in an enormously sophisticated machine whose operations are obscured from them.

Algorithmic management practices are already well established among large companies in the low-wage labor market[81] and seem very likely to become more important going forward. On the demand side, because it is difficult to generate breakthrough productivity gains in service occupations, and because so many sectors in the US are dominated by a low-wage, low-productivity model, companies frequently maintain profitability by squeezing ever-greater effort from workers. Algorithmic management techniques can help them to do so. On the supply side, algorithmic management devices and technologies are substantially cheaper than physical automation devices. Algorithmic management requires new physical sensors, including cameras, infrared and barcode scanners, listening devices, and the like. But those devices themselves do not need to move around or apply physical force to the environment, which makes deploying them cheaper and operationally simpler than deploying new robotics. Such efforts are also clearly cheaper than human-powered surveillance and management in many cases.[82]

Companies can use algorithmic management both to concentrate valuable information in fewer hands[83] and to intensify work efforts.[84] The discussion in the next subsection first surveys developments here and then assesses their effects on class relations. It focuses on production activities and relations between managers (or algorithms) and individual workers. The next chapter takes up hiring processes, discrimination, and workers' organizing, which raise distinct legal issues. As noted in the introduction of this chapter, this section doesn't say much about law because companies have broad authority to take these steps.

Algorithmic management—an overview Algorithmic management is now well established in at least three subfields of management practices: assigning tasks to workers, scheduling workers for shifts, and setting the pace of work. The most prominent examples of algorithmic tasking involve the gig economy.[85] Uber claims to match consumers and drivers more quickly and reliably than street-hail systems, drawing on data gleaned from drivers' and consumers' cell phones. There is evidence that it enhanced productivity through that method, at least for a time and in some jurisdictions.[86] But Uber has also exploited its control of tasking algorithms to erode drivers' capacity to earn a decent living.[87] As the company moved into cities (often illegally), it led customers to expect not just reliability but also quick availability, so it could build market share rapidly. Uber did so by flooding the market with drivers[88] and then tightly controlling their access to information. Early on, at least, when the app sent a ride request to a driver, the driver had only fifteen seconds to accept it—and had to do so without knowing the destination or fare. Drivers who refused too many requests, or who canceled rides after accepting them, risked deactivation. Uber also uses various behavioral "nudges" to keep drivers on the road, for example by selectively sending ride requests to drivers who seem likely to log off.[89] Similarly, on weekends when it expects a high volume of ride requests, Uber sometimes sets up incentive structures where drivers can receive bonuses if they carry out a certain number of rides and accept a certain percentage of requests.

But drivers have alleged that the company may use its panoptic knowledge about the market to prevent most drivers from reaching the bonus stage by cutting off ride requests once they get close to that point, or by sending them "phantom" ride requests that they are unable to accept.[90] In these and other ways, Uber enforces market discipline on workers, requiring them to compete—tacitly or actively—to remain on the platform and to get the best opportunities. Similar tasking strategies are used across the low-wage gig economy, among companies like Lyft, Instacart, Doordash, and Amazon. In a telling example of how precarious workers have become dependent on these companies and their tasking algorithms, in September 2020, it was reported that Amazon delivery drivers had taken to hanging smartphones from trees near distribution centers and Whole Foods stores, and then syncing their own phones to those. In that way, they hoped to get early notification of potential gigs since the company's algorithms would send pings to whichever delivery drivers were physically closest to the job site.[91]

Another transformation in management has occurred around scheduling practices, where many large companies now use algorithms to assign workers to shifts. Those algorithms claim to predict demand based on past sales, as well as factors such as weather reports, and then schedule workers accordingly in an effort to ensure that work sites are neither overstaffed nor understaffed.[92] This may involve inductive learning to determine which external and internal factors are likely to affect demand the most. Indeed, the use of such algorithms is not necessarily a negative development from the worker's perspective. If workers can specify times that they would ideally like to work, and an algorithm can figure out how to optimize the schedule for a manager, a company can reduce managerial costs and help ensure worker satisfaction.[93]

Companies have frequently used algorithmic scheduling, however, to reduce workers' time on the clock to an absolute minimum. Fast food workers have complained that their shifts are canceled at the last minute or that they are sent home in the middle of shifts without notice.[94] In other cases, companies use algorithms and apps to make schedule alterations on an on-demand basis. Amazon's scheduling practices were noted earlier in this chapter; some Amazon workers suspect that their willingness to accept shift changes on a moment's notice is factored into the company's decision of whether to offer them desirable shifts, overtime, or both.[95] Other companies have scheduled workers for shifts that make it impossible for them to fulfill caregiving responsibilities, or even to sleep. The issue came to public attention with Starbucks' practice of "clopening," where workers had to close the store one night and then open it the next morning, making it nearly impossible for them to sleep.[96] In the wake of media attention, Starbucks promised more reasonable and predictable schedules going forward.[97]

For the most part, these new scheduling practices do not even trigger scrutiny under federal working time regulations. The federal Fair Labor Standards Act (FLSA) does not actually guarantee steady hours or minimum or maximum hours.[98] All it requires regarding hours is that employers pay time-and-a-half for all hours worked over a statutory norm, which has long been forty hours per week. Worker advocates have pressed states and localities to legislate on the issue recently. The state of Oregon, as well as a number of cities including New York and San Francisco, have passed "fair workweek" laws, which require companies to provide workers more notice of schedules; some also require companies to guarantee a certain number of hours off between shifts.[99] Algorithmic scheduling practices may also facilitate "wage theft," or the failure to pay workers all they are owed under wage and hour laws. Three legal scholars who reviewed common timekeeping software programs that are often used in conjunction with algorithmic scheduling found that their default settings would often undercount hours, and that the programs enabled employers to edit down hours worked, which is a crude and obvious FLSA violation.[100]

Finally, companies are using algorithmic monitoring—and their legal powers to terminate workers at will—to require workers to perform at a rapid pace. This is again not a new development.[101] In the 1970s, as supermarkets began to introduce barcode scanners at cashier stations, the sociologist Harry Braverman warned that the technology could be used to track employee performance.[102] Since then, those systems have become far more powerful and sophisticated. The sociologist Karen Levy has documented how long-haul trucking companies use telematics-based monitoring to push their drivers to work all the hours permitted under federal law, substantially reducing the autonomy that was once a point of pride among truckers.[103] At one company, such speedups correlated with a significant increase in workplace injuries.[104] Home-care workers are increasingly required to use "electronic visit verification" apps that invade their clients' privacy, and which may lead workers to underreport their hours of work, in turn leading to lost wages.[105] Even white-collar workers now face such monitoring. Under COVID, many companies have ramped up surveillance of remote workers, using laptop cameras and facial recognition to discern, minute by minute, whether professionals were focused on assignments.[106]

Amazon has deployed similar technologies within its warehouses to determine how quickly workers are performing tasks and to push them to work faster.[107] Documents disclosed as part of a labor dispute between Amazon and a worker who alleged that he had been fired in retaliation for union organizing efforts showed that various aspects of that oversight had been automated.[108] "Amazon's system tracks the rates of each individual associate's productivity," one document said, "and automatically generates any warnings or terminations regarding quality or productivity without input from supervisors."[109] Around 300 workers in that warehouse, representing over 10 percent of the warehouse's staff, had been terminated via that process, for productivity reasons alone, in a twelve-month period.[110] Other companies are trying to observe human workers' movements very closely using networks of sensors, radar, and machine learning.[111] Still others are developing wearable devices to track workers' movements through the workplace.[112]

Effects of algorithmic management Algorithmic management seems to be maturing more rapidly than robotics, and when this book is published, various new affordances will likely be in use. Nevertheless, the field is sufficiently well developed that some of its aggregate effects on workers seem clear.

The first is that algorithmic management will often put downward pressure on wages. With near-perfect information about individual workers' performance and the capacity to run experiments in wage setting across large-scale enterprises, companies should be able to determine exactly how much they need to pay particular workers to keep them around. Put differently, these technologies are eliminating some informational asymmetries that give workers some bargaining power even if they are not in unions. As labor economists in an earlier technological era argued, when it is not easy to observe workers' output or effort levels, companies may pay above-market wages to induce employee loyalty, and therefore sound performance.[113] That could occur if individuals worked in teams, for example, or if workers had used collective action to block employer surveillance efforts. The converse also appeared true at the time: where companies can cheaply detect underperforming workers, they have less incentive to pay above-market wages.[114] In such cases, workers were subject to significant external market discipline. As should be clear from the discussion in this chapter, companies today often use data-driven workplace surveillance to detect underperforming workers. Peer-reviewed research on new tracking efforts and wages is rare, but one study of the platform Freelancer.com found that the introduction of a monitoring system that tracked keystrokes and other actions led to a substantial increase in bids, including from inexperienced workers, and lower prices for labor.[115] The authors reasoned that the monitoring system reduced the value of established workers' reputations and mitigated the risks that companies faced when hiring workers without a history on the platform.

Second, algorithmic management can erode or suppress workers' associational power. After all, workers' best means of protecting themselves against very low wages or an unsustainable pace of work in the past has been to organize.[116] But workers who are constantly supervised and physically separated from one another have little time to meet and make common cause. Chapter 4 says more about this, discussing the organizing process in detail. Similarly, workers who have multiple jobs with irregular schedules may be too exhausted to even think about planning collective action. Meanwhile, the homogenization of work tasks under digital Taylorism makes it easier for companies to plug workers into and out of jobs fairly easily, reducing companies' training costs and their incentives to invest in their workforce and making it still more difficult for workers to organize. Labor markets then resemble their neoclassical models, with workers not much different from classic commodities.

Meanwhile, the use of algorithmic management and other surveillance efforts are enabling companies to replicate and capture some of workers' tacit knowledge. This is a third aggregate effect of algorithmic management. It was also a central function of classical Taylorism, but digital technologies have extended the process into new spheres. For example, Uber has integrated global positioning system (GPS)–powered navigation into the driver side of its app, and the company may be able to improve it continuously using data from past rides.[117] But taxi drivers' specialized knowledge of how to navigate a crowded city historically gave them some labor market power. In London, cab drivers even need to pass a test showing that they know the names and locations of all streets in the area so that they can get to any location without a map.[118] In essence, Uber has captured or replicated some of taxi drivers' classical knowledge and craft skills, claimed property rights in them via IP doctrines and trade secrets, and leased that IP back to drivers. Through those efforts, the company can put downward pressure on wages since the new technology enables almost anyone with a vehicle to do the job. Similar efforts may be underway in customer service call centers as companies implement natural language-recognition and processing algorithms, and in fast food as franchisors discern facts about local management strategies or labor practices that aren't visible from the shop-floor level.

There is a final cross-cutting issue here: past a certain point, companies' efforts to reduce labor costs by squeezing workers may become self-defeating. As the industrial relations scholar Zeynep Ton has shown, retailers can get caught in a "vicious cycle" of low wages and low productivity. Overworked cashiers end up misscanning items, for example, and stock clerks end up placing excess goods in random locations because their stores are constantly short-staffed.[119] That appears to be especially common in the US for institutional and path-dependent reasons discussed previously: through the era of neoliberalism, US companies have been especially focused on labor discipline strategies. Comparative scholarship on industrial relations in retail and other service sectors such as call centers suggests that many European firms take a more collaborative approach to labor relations, with more upskilling and broader worker discretion.[120] That reflects, in part, workers' greater capacity in Europe to foreclose the zero-sum and discipline-intensive managerial practices common in the US, due to European nations' more robust collective bargaining systems and privacy protections.

Unfortunately, companies can find it quite challenging to break out of Ton's vicious cycle, especially where their competitors are mostly using the same managerial techniques. A similar dynamic appeared among delivery companies in the US after COVID. In 2021, FedEx was significantly underperforming UPS in terms of on-time package delivery and profit margins, due largely to labor shortages. As discussed in chapter 5, FedEx uses independent contractors rather than employees and keeps wages low, while UPS employs drivers directly, is unionized, and pays drivers more than its rivals. Those factors helped UPS thrive during the post-COVID labor shortage and suggest that FedEx may be unable to push as many costs onto drivers going forward.[121]

Conclusion: Institutions and Digital Taylorism

Digital Taylorism reflects a clear underlying logic: companies are using their legal and operational powers over data to automate some tasks, reorganize production accordingly, and supervise workers much more intensely. By doing so, companies can both enhance productivity and ensure that workers cannot capture a significant share of profits. As should also be clear, this process is facilitated and shaped at every stage by law. Whether and when companies can monitor workers, what they can do with the data they have gathered, and who "owns" any new products they develop using that data all involve legal questions. Through digital Taylorism, companies are using their control over technology to reconfigure workplaces and jobs, and then to monitor and manage workers in new ways, often through means that erode or suppress workers' associational power.

In that sense, companies' power over technology gives them the capacity to override workers' legal rights and to encase their own powers against workers' challenges. Those efforts have helped bring into being today's working class, which as chapter 1 discussed is heavily concentrated in service sectors. At the same time, the fact that alternative models of labor relations have persisted in other nations suggests that a different politics of workplace technology is possible. Chapter 6 returns to this question, asking how labor law reforms could contribute to such a transformative agenda.

Data and Democracy at Work

Advanced Information Technologies, Labor Law, and the New Working Class

By: Brishen Rogers

Citation:

Data and Democracy at Work: Advanced Information Technologies, Labor Law, and the New Working Class
By: Brishen Rogers
DOI: 10.7551/mitpress/11253.001.0001
ISBN (electronic): 9780262373357
Publisher: The MIT Press
Published: 2023

The open access edition of this book was made possible by generous funding and support from MIT Press Direct to Open

The MIT Press

The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Stone Serif and Stone Sans by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data

Names: Rogers, Brishen, author.
Title: Data and democracy at work : advanced information technologies, labor law, and the new working class / Brishen Rogers.
Description: Cambridge, Massachusetts : The MIT Press, [2023] | Includes bibliographical references and index.
Identifiers: LCCN 2022015330 (print) | LCCN 2022015331 (ebook) | ISBN 9780262545136 (paperback) | ISBN 9780262373364 (epub) | ISBN 9780262373357 (pdf)
Subjects: LCSH: Labor laws and legislation—United States. | Employees—Effect of technological innovations on—United States.
Classification: LCC KF3319 .R64 2023 (print) | LCC KF3319 (ebook) | DDC 344.7301—dc23/eng/20220831
LC record available at https://lccn.loc.gov/2022015330
LC ebook record available at https://lccn.loc.gov/2022015331