

Automation

The annihilation of work

By the first decade of the twenty-first century, the average technologically augmented individual had effectively become superhuman. By and large, however, those of us who have undergone this transformation don't stop to wonder or even notice that this is the case, unless the temporary failure of our augmentive systems brings us up short. So long as everything stays up and running, we communicate instantaneously and on demand with anyone, more or less anywhere on Earth. We traverse the planet physically in a matter of hours. We enjoy near-universal access to a medium that connects virtually everyone we've ever met or ever will meet, and furnishes us beyond that with immediate access to most of the things anyone has ever spoken, written, painted, sung or committed to the screen—if not by any means all of the knowledge ever developed by the species, then certainly a very respectable fraction of it. Within the bounds imposed by our current understanding of material science, we can design and build just about anything it is possible to imagine.

Yet for all the impressive scope and reach of our power, we still face stark, seemingly fundamental limits of materiality and

mortality. Life still presents us with the broadest array of tasks we're not very good at—those we can only do for so long before tiring, those we quail at, and those that place us in physical or psychological danger. None of those tasks are getting any easier, and none of them are disappearing any time soon. And overhead arcs the distant but certain promise of mortality. In a very real sense, we are all of us pressed for time.

Some of us are willing to accept this state of affairs as a given, and possibly even as a tutor. But others are not. Out of this unwillingness, these people have set out to devise technical systems that are more capable than we are ourselves, along any axis or dimension of evaluation you might care to mention: systems that are stronger and faster than we are; that have finer perception and greater endurance; that never, ever succumb to boredom, fatigue or disgust; and that are capable of operating without human oversight or guidance, indefinitely. We are, of course, talking about using robots and automated systems to replace human labor.

The great twentieth-century economist John Maynard Keynes had foreseen much of this early on, coining the expression “**technological unemployment**” sometime around 1928.¹ He saw, with almost clairvoyant perspicacity, that societies might eventually automate away the jobs much of their labor force depended on, and his insight is borne out in **recent United States government estimates that an American worker making less than \$20 an hour now has an 83 percent chance of losing their job to automation.**² But what Keynes concluded—that the **eclipse of human labor by technical systems would necessarily compel a turn toward a full-leisure society**—has not come to pass, not even remotely.

And what neither Keynes nor any other economist reckoned with, until very recently, was the thought that the process of automation would hardly stop when it had replaced manual and clerical labor. If automation was initially brought to bear on tasks that were one or more of the “four Ds”—dull, dirty, difficult or dangerous—the advent of sophisticated machine-learning algorithms means that professional and managerial

work now comes into range. We need to be clear that automated systems might replace any one of us in our jobs, however nominally executive or “creative” those jobs might be.³

We should be careful not to overstate how fast this process might take place. Nor should we make the mistake of imagining it as some impersonal, ahistorical transformation, beyond anyone's control; like any other such process, its progress turns upon identifiable actors, acting from discernible motivations, intervening at material sites. But it's just as crucial that we not underestimate the speed, force or thoroughness with which this is taking place. We now stand at a juncture where there is no pursuit that cannot in principle be undertaken by an automated system, and we need to come to terms with what that might mean for the economy, the ways in which we organize our societies, and our own psyches. However disturbingly crude they remain, to whatever extent they are all too evidently the fruit of their creators' biases, neuroses and projections, these systems are nevertheless poised to assume responsibility for much of the work that furnishes us with a livelihood, an identity and a sense of self. And they are getting better at what they do with every passing day.

What I wish to argue is that whether they are brought together consciously or otherwise, large-scale data analysis, algorithmic management, machine-learning techniques, automation and robotics constitute a coherent set of techniques for the production of an experience I call the posthuman everyday. This is a milieu in which the rhythms we contend with, the ordinary spaces we occupy, and the material and energetic flows we support are all shaped not so much by our own needs but those of the systems that nominally serve us, and in which human perception, scale and desire are no longer the primary yardsticks of value.

This posthuman turn will profoundly reshape the environment in which we live and act. But it will also compel political change, driving realignments of belief and commitment along a set of novel axes. Our ability to think clearly about what kind of a deal we wish to pursue with these technologies—what place

we wish for them to occupy in our economies, our societies and our lives, and ultimately what kind of everyday we ourselves wish to occupy—will depend on what we know about how each of them works, and how they may come to be fused in something powerful, purposive and self-directing.

As is so often the case, the language around these subjects can be daunting, but none of this complexity need stand in the way of understanding the basic propositions on offer. Before immersing ourselves in it, though, it's worth asking a question whose answers evidently strike advocates as being so obvious and so self-evident that they hardly warrant discussion: Why do so many of us seem to want to replace ourselves so very badly?

Reflexively, most of those who are now working to develop automated systems would probably prefer to speak of supplementing, rather than replacing, human beings in all the myriad roles we now fill. They paint pictures of hybridized teams working in varying degrees of comradely harmony, toward ends that are exclusively those specified by the human partner. If asked what motivates them in their drive toward the post-human, they'd very likely say they intend no such thing at all.

Instead, there's a very good chance they would offer a selection from among the following justifications, some of which are all but self-contradictory, and others of which are far more defensible. Assuming that none of them actually nurture a suppressed craving for the twilight of the species and a will to extinction, these are the justifications people offer when they undertake the development of learning algorithms, or the fully autonomous systems based on them:

We shouldn't underestimate, firstly, the force of desire that drives not a few researchers in robotics and artificial intelligence—the magnitude of the intellectual challenge. Any ascent to human-level intelligence by a technical system would by definition rank as one of the towering achievements of human history. Its designers could be assured of being remembered and celebrated alongside the titans of science and industry, for however much longer the species subsisted. While that prospect

might not drive you or me toward this work, it's manifestly not without its power to seduce.

Nor do I think we should underestimate the degree to which some are invested in the development of autonomous systems for commercial reasons—purely, that is, as a business play, and not because they have any other personal, intellectual or ideological commitment to the idea. As long as there is money to be made in the field, there are those who will be willing to dedicate themselves to it, body and soul, whatever it might portend for the future.

We can be sure, then, that some cohort of the people working in this field will confessedly not actually care much about what automation is, what it means, or what might happen to result from its development. But others are very explicit in what they want from this set of technologies.

I don't think it's overly cynical to suggest that we might want to discount any professed concern for relieving human beings of dirty, dangerous or degrading jobs. Even in our era of acute sensitivities, employers seem to have little compunction about exposing workers to virtually any level of drudgery or peril, as long as their labor can be had cheap enough, and the consequences of catastrophe made tolerable enough.⁴ The problem, from an employer's perspective, isn't so much that their workers' dignity must somehow be protected, but that even the minimal steps necessary to do so are becoming increasingly expensive. So one clear motivating factor behind research into automation is the creation of a *cheaper* and *far less fractious* labor force—one that never demands overtime, never agitates for higher wages or better benefits, never sues on grounds of discrimination, never files for workers' compensation,⁵ and never complains about the conditions to which it is exposed.

Blue-collar wages may be stagnating globally, but automation costs are falling.⁶ A used palletizing robot in good condition, capable of going 50,000 hours between overhauls (or just over eight years of three-shift operation), may cost \$15,000 or less on the open market; while for the time being it may require an operator (at some \$36,000 a year), even that outlay compares

very favorably to the wages that would otherwise be paid out to a full-scale crew of human workers over that eight-year period. (As we shall see, that operator's job is itself in jeopardy: most of what he or she does is already within reach of rapidly developing machine vision systems and control algorithms.)

A slightly different calculus of cheapness drives much of the interest in military robotics: a *concern for the relative price of human life*, or at least those human lives that are designated as friendly. For decades, strategists have nurtured ideas about swarms of "fast, cheap, out of control" cybernetic combat systems, capable of occupying contested terrain, degrading an enemy's will to resist, and ultimately destroying that enemy's ability to wage war.⁷ The remotely piloted vehicles we think of as "drones" are a gesture in this direction; more or less disposable for all their expense, they can be shot down or blown to pieces even as an operator reposes in air-conditioned comfort 8,000 miles away. But they're merely an incremental step toward more fully autonomous combat systems that would allow a state to pursue its territorial ambitions all without exposing (its own) human combatants to risk.

There are contexts in which the appeal to *safeguarding life and limb* actually does furnish a reasonable, even an ethically compelling, case for automation. In the United States alone, more than 30,000 people are killed in motor-vehicle crashes every year;⁸ more than nine in every ten of those crashes were caused, at least in part, by driver error.⁹ Though the number of fatalities is slowly falling over the decades, as regulation tightens and safety technology improves, this remains by any standard a heartbreakingly high price to pay for our desire for unlimited mobility. Given the scale of the risks involved, not least to insurers,¹⁰ we can readily enough imagine a time when it is illegal to operate a vehicle, tool or other potentially destructive device manually, if the option exists of surrendering control to an algorithm bound to be more competent than we are ourselves.

There are, of course, other contexts where the vagaries of human discretion may produce intolerable effects. We know, for example—many of us know all too well—that our public

institutions are pervaded by pernicious, systemic bias, and that this bias not infrequently has lethal consequences for anyone on the wrong side of it. This is, in some bleak sense, an opportunity for the vendors and advocates of algorithmic systems, which are commonly promoted as being neutral, rational, objective and scientific. Often with the noblest of intentions, they aim to restore probity to processes of public administration by supplanting imperfect human discretion with a synthetic judgment that is invariably figured as *free from the contaminating influence of prejudice*.

Our inability to act without bias is not the only way in which we fall short of the standards we set for ourselves, or which are routinely expected of us at work. As they develop over time, every industry, every profession and every domain of human endeavor will demand tolerances, sustained efforts or throughput velocities that are simply unachievable by human means. This might mean the ability to sift through large evidentiary productions without losing focus or succumbing to boredom. It could entail the capacity to maintain watchful awareness at a fence or border crossing deep in the night, without becoming distracted or sleepy; or, given the pace that capital already expects from human workers, it might involve the acuity to correctly identify contaminated chickens on a production line churning by at 175 birds per minute. Algorithmic systems will be pressed into service in all of these cases, and any number of contexts beyond where the demands of performance tax human strength, endurance, sensory discernment or motor coordination, simply because they are seen as being *more reliable*.

Another way in which nominally autonomous systems are thought of as being "better" than human beings is that they are *more obedient*.

Consider the lessons of a chastening post-World War II study conducted by the US Army, which found that few soldiers bothered to aim when firing their weapons in combat, or for that matter seemed to have any particular taste for killing at all, even when faced with a clear and highly motivated enemy. Since this finding, enormous sums have been invested in research aimed at

ensuring that a military unit might actually bring kinetic force to bear on its intended target. Much of this research has focused on indoctrination and other methods of inculcating troops with the instinct to kill, and much of it has been successful.¹¹

But in recent years the thrust of this effort has moved toward removing human beings from the decision loop entirely, along with any feelings of fear, compassion or mercy that might compromise the will to act at the decisive moment. Already well advanced in the doctrine of remotely piloted aerial warfare, with its air-conditioned control trailers and rhetoric of near-bloodless “precision strike,”¹² work proceeds toward the development of fully autonomous combat vehicles in China, Israel, Pakistan, Russia, South Korea, Turkey, the United Kingdom and the United States.¹³ Though such systems are blessedly far removed from everyday experience for most of us, the long, sordid and well-documented history of “blowback,” along with the increasing material and psychic militarization of policing,¹⁴ suggests the strong possibility that techniques originally developed in and for distant theaters of war will eventually find their way into domestic life.

Finally, beyond arguments from cheapness, safety, reliability or obedience, I have no doubt that a few advocates for automation are driven by a sincere and passionate belief that automation is the surest way of achieving a *more equitable* future. Some on the left—accelerationists such as Alex Williams and Nick Srnicek and the proponents of Fully Automated Luxury Communism prominent among them—have argued that the ends of economic justice in our time are best served by maximum automation and the elimination of work.¹⁵ Thinkers of this stripe argue that the soonest possible supplantation of human labor by cybernetic means is something close to an absolute ethical imperative.

In some ways, left accelerationism is just a contemporary gloss applied to the visions of total leisure that were developed by the generations immediately preceding, in a few distinct currents. Forerunners like Constant Nieuwenhuys, the French Situationists, and the radical Italian architectural practice Superstudio explored the spatial dimensions of post-work society

between the late 1950s and mid-1970s, developing conceptions of what urban environments might look like when more broadly arranged around self-actualization and play, while it was left to second-wave feminist thought to explore the social dimensions of full automation. Shulamith Firestone organizes much of the later argument of her 1970 *The Dialectic of Sex* on “machines that ... surpass man in original thinking or problem-solving,” “the abolition of the labor force itself under a cybernetic socialism, the radical restructuring of the economy to make ‘work,’ i.e. wage labor, no longer necessary,” and ultimately the creation of a society in which “both adults and children could indulge in serious ‘play’ as much as they wanted.”¹⁶ And automated production, of course, furnishes Valerie Solanas one of the hinges of her *SCUM Manifesto*, with its ever-resonant demand to “overthrow the government, eliminate the money system, institute complete automation and destroy the male sex.”¹⁷

But for that last, these themes have recently reemerged to be picked up and developed in the contemporary accelerationist discourse, after a long interval during which utopian thought on the left seemed more captivated by the possibilities of networks than by any catalyzed by automation per se. As the accelerationists would have it, the advanced econometric modeling, algorithmic planning and networked digital fabrication we now have access to finally make it possible for someone to “do one thing today and another tomorrow” as Marx foresaw—to binge-watch in the morning, write fanfic in the afternoon, make reaction GIFs in the evening and criticize after dinner, just as she has a mind. The problematic before us then actually would become the Keynesian (or Olympian) one of learning to live “wisely and agreeably and well” under conditions of absolute and universal freedom from want.

In the end, what is it that people want from these technologies? As near as I can tell, a few want just exactly what some have always wanted from other human beings: a cheap, reliable, docile labor force. Others, though, are seeking something less tangible: sense, meaning, order, a ward against uncertainty.

They're looking for something that might help them master the combinatorial explosion of possibility on a planet where nine billion people are continually knitting their own world-lines; for just a little reassurance, in a world populated by so many conscious actors that it often feels like it's spinning out of anyone's control. These are impulses I think most of us can relate to, and intuitively react to with some sympathy. And it's this class of desires that I think we should keep in mind as we explore the mechanics of machine learning, automated pattern recognition and decision-making. For all the arrogance, the reductionism and the more than occasional wrongheadedness that crop up in this development effort, I believe it is founded in a set of responses to the world that most all of us have experienced at one time or another. If nothing else, to consider automation with any seriousness is to be presented with a long, poignant and richly elaborated index of our deepest longings and fears.

Sometime in the early months of 2015, a curious map started to circulate through my social-media networks.¹⁸ Generated by the data-graphics specialist Quoc Trung Bui for National Public Radio, it charted the most commonly held job in each of the fifty United States during the period 1978–2014.

For Nevada, this job was “retail clerk,” Washington State and Virginia “software developer,” and North Dakota “farmer.” Just about everywhere else in the United States, though, the most common occupation over this span of time was that of “truck driver.” It was in seeing that job title lettered over every corner of the map that I really began to understand what automation is going to do to the American economy—and by extension, every other national economy on Earth at a similar stage of development.

In some ways, perhaps, the NPR map was a little misleading. Due to a quirk in the way the US government classifies jobs, long-haul truckers, delivery personnel and farm-equipment operators are shoehorned into the single category “truck, delivery and tractor drivers”; Bui had chosen to apply the label

“truck driver” to this category, collapsing still further whatever distinctions may have nestled within it. But his animation of the map, and the year-by-year job tables he helpfully included, left no doubt that the phenomenon was and is real. The occupation has only become more dominant over time, displacing farmer, factory worker and machine operator as the number-one job—and, by extension, presumably the chief source of household income—in state after state.

As it happens, though the congelation of long-haul trucking, mechanized farming and urban logistics into a single job category obscures as much as it clarifies, it may not matter for our purposes. Trucking, farming and logistics are all seething sites of research into automation, and none will likely survive very long as distinctly human fields of endeavor.

“Driverless” cars like those being developed by Google and Uber may dominate the mainstream media coverage, but the spare, highway-bound performance regime of long-haul trucking is far more amenable to automation than the stop-and-start, high-complexity environment of city and suburban driving, as Tesla's July 2016 announcement of plans for an autonomous semi recognizes.¹⁹ Meanwhile, logistics has already shed most of the human labor force it once supported, with Jeff Bezos's Amazon pioneering the development of robotic warehousing and fulfillment.²⁰ Research has already moved on to attack the challenges of delivery by swarming drone as well as autonomous ground vehicle.

So when a worthy body like the Pew Research Center convenes a panel of experts²¹ and just over half of them argue that automation will create more jobs than it displaces between now and 2025, I'm forced to wonder if anyone involved has spent much time contemplating the pastel contrasts of Quoc Trung Bui's map.²² What is now the most commonly held job in twenty-nine of the fifty states will surely number among the very first to be automated. Autonomous trucking alone is going to land on the American economy (and the American worker) with devastating force.

The usual rattling off of statistics doesn't really tell us very

much, but it can at least help sketch in the outlines of what it is we face: of 702 detailed job categories in the United States, Carl Benedikt Frey and Michael A. Osborne of the University of Oxford found that 47 percent of them were vulnerable to near-term advances in machine learning and mobile robotics.²³ Among developing countries, this rises to 69 percent in India, 77 percent in China and an astonishing 85 percent in Ethiopia.²⁴ (Again, these figures refer to the percentage of job *categories* that are susceptible to replacement, not of workers in employment.) Meanwhile, against the oft-cited hope that technology would generate more jobs than it eliminated, Frey found that fewer than 0.5 percent of the US workforce have found employment in the high-technology industries that have emerged since the turn of the century. A World Economic Forum estimate that some five million jobs would be lost to automation by 2020 has to be regarded as a stark outlier, if not a gross error, especially since Bank of England Chief Economist Andy Haldane reckons that 15 million jobs would disappear over the same timeframe in the United Kingdom alone.²⁵

I'm not qualified to discuss, in any but the broadest terms, what will happen to the shape and structure of national economies in the aftermath of pervasive automation. What I can speak to, however, is what the working environment will look and feel like for those of us who do manage to remain in employment.

We often tend to visualize automation in heroic terms, as if limned in the hues of a cybernetic Socialist Realism: we imagine ranks of mighty robot tractors furrowing the grain in parallel bands ten miles across, or enormous darkened warehouses where relentless manipulators pick, pack and ship in silence. As far as industry is concerned, though—and in this instance it really is their perspective that weighs heaviest and counts most—automation also means far less elaborate technologies, like the touchscreen ordering kiosks McDonald's began introducing into its locations in the fall of 2014. In fact, automation means *anything* that reduces the need for human workers, whether it's a picking-and-packing robot, a wearable

biometric monitor, a mobile-phone app or the redesign of a business process.

Like McDonald's CEO Steve Easterbrook, some executives insist that bringing service-sector automation on line won't eliminate jobs, but merely allow them to deploy their employees in more "value-added" roles.²⁶ Former McDonald's head of US operations Ed Rensi argues just the opposite; as part of a campaign against efforts to increase the US minimum wage to \$15 an hour, he noted that "it's cheaper to buy a \$35,000 robotic arm than it is to hire an employee who's inefficient, making \$15 an hour bagging french fries."²⁷ Whichever of these two perspectives does prevail—and the pitiless logic of shareholder value strongly bolsters Rensi's position—it is clear that whatever human participants do remain in the waged labor force are in for a particularly rough ride in the years to come.

This shrunken workforce will be asked to do more, for lower wages, at a yet higher pace. Amazon is again the leading indicator here.²⁸ Its warehouse workers are hired on fixed, short-term contracts, through a deniable outsourcing agency, and precluded from raises, benefits, opportunities for advancement or the meaningful prospect of permanent employment. They work under conditions of "rationalized" oversight in the form of performance metrics that are calibrated in real time. Any degree of discretion or autonomy they might have retained is ruthlessly pared away by efficiency algorithm. The point couldn't be made much more clearly: these facilities are places that no one sane would choose to be if they had any other option at all.

And this is only the most obvious sort of technological intervention in the workplace. We barely have words for what happens when an algorithm breaks down jobs into tasks that are simple enough that they don't call for any particular expertise—just about anybody will suffice to perform them—and outsources them to a global network of individuals made precarious and therefore willing to work for very little. Naturally, this newly intensified Fordist production regime sees its workers paid by the minute, without security of tenure, a

guaranteed weekly minimum or any other form of benefit, and this too is automation.

Most of the blue-collar workers that do manage to retain employment will find themselves “below the API”—that is, subject to having their shifts scheduled by optimization algorithm, on little or no notice, for periods potentially incommensurate with their needs for sleep and restoration, their family life, or their other obligations.²⁹ (In the UK and elsewhere the practice is tolerated, the terms of such employment may be specified by so-called zero-hour contracts, which offer no guaranteed minimum of shifts.) Former US secretary of labor Robert Reich recoils from the prospect of work under these conditions: “Can you imagine if this turns into [an economy] where everyone is doing piecework at all odd hours, and no one knows when the next job will come, and how much it will pay? What kind of private lives can we possibly have, what kind of relationships, what kind of families?”

Under such circumstances, the workplace itself becomes an arena for every kind of performance monitoring and calibration. Every checkout interaction at the Target chain, for example, is rated Green, Yellow or Red by an automated system, according to whether or not the clerk hit targets for speed and accuracy, and these ratings are used to determine employee compensation. Target worker “Tessa” explained in a blog post that the company “keeps a running average of your scores for the week, month, and year. They expect over 88 percent of your transactions to make the speed cut, and your score reflects on possible raises, promotions, and sometimes even who remains as an employee.”³⁰ (An alternate—and in truth, almost equally credible—possibility is raised by a Target employee commenting at another site: “No one cares about it at the store level, it’s just to get people at Corporate to feel that your store is productive.”)³¹

Since June 2016, sales associates at every Container Store location have been required to wear “enterprise wearables” made by a startup called Theatro.³² These are devices that track employee location and provide real-time feedback regarding their interactions with shoppers. The company uses the accompanying

analytic suite to “identify top performers” (and, by implication, those at the bottom as well), and plan schedules and distribute assignments in the store accordingly.

Theatro’s devices are less elaborate than a Hitachi wearable called Business Microscope, which aims to capture, quantify and make inferences from several dimensions of employee behavior.³³ As grim as call-center work is, a Hitachi press release brags about their ability to render it more dystopian yet via the use of this tool—improving performance metrics not by reducing employees’ workload, but by compelling them to be more physically active during their allotted break periods.³⁴

Hitachi’s wearables, in turn, are less capable than the badges offered by Cambridge, MA, startup Sociometric Solutions, which are “equipped with two microphones, a location sensor and an accelerometer” and are capable of registering “tone of voice, posture and body language, as well as who spoke to whom for how long.”³⁵ As with all of these devices, the aim is to continuously monitor (and eventually regulate) employee behavior.

Whether furnished by a large and globally established enterprise like Hitachi, or a startup along the lines of Sociometric Solutions or Theatro, employers will have their choice of devices that allow them to track worker performance and attitude, along an ever-increasing number of axes and in constantly improving resolution. It’s also worth pointing out, in turn, that in every generation of product, such specialized employee monitors are notably less capable than consumer-grade biometric wearables like the Fitbit Charge or Apple’s Watch; smaller, or canner, enterprises may simply choose to capture data from the devices their employees are already voluntarily wearing. In either case, the disciplinary regime specific to the workplace becomes instead an all-but-placeless control regime, an “ambient factory” where decisions once thought to be purely personal—sleep cycles, nutritional patterns, exercise habits—become subject to employer monitoring and intervention.³⁶

Wearables work in synchrony with workstation systems to furnish bosses with a comprehensive overview of employee behavior. At KFC, Wendy’s and the RiteAid drugstore chain,

workers at the point of sale log into their workstations biometrically via the Crossmatch suite of hardware and software. Crossmatch touts this system as preventing “tardy arrivals, ‘buddy punching,’ ‘lollygagging,’ extended breaks and early departures, inventory shrink, unauthorized discounts and returns, and fraudulent gift card transactions;”³⁷ anyone who’s ever worked retail will understand this list as an almost perfect, point-for-point recitation of the tacit measures employees have always taken to compensate themselves for having to put up with abusive bosses, shitty pay and intolerable working conditions. For that matter, “buddy punching”—the act of clocking in a friend who’s late for work, possibly because they’ve had to take a child to daycare or a sick parent to the doctor—is just the kind of small act of solidarity that might save someone their job in the harsh, zero-tolerance climate of contemporary work. But these are the tactics such oversight systems are expressly designed to eliminate.

As the capacity to detect and characterize emotional states has grown, these reasonably traditional, Taylorist notions of time-and-motion efficiency have been supplemented by a concern for the worker’s affective performance.³⁸ Japan’s Keikyu Corporation, for example, began measuring the quality of its frontline employees’ smiles in 2009, scanning their “eye movements, lip curves and wrinkles,” and rating them on a 0-100 scale.³⁹

As intrusive as this may seem, smiling is at least something under an employee’s conscious control, which cannot be said for all of the measurements of “body posture, facial expressions, physiology, semantics [and] who a person talks to and when” that the management consultancy Accenture recommends to ensure employees are “exhibiting effective social behaviors.”⁴⁰ Such subconscious tells are picked up by the People Analytics suite the “emotion-aware sentiment analysis company” Kanjaya offers, which uses unstructured voice and text data to calibrate an employee’s “Attrition Risk” and “Workplace Value,” in addition to the expected “Performance.”

The concern for retention implies something that a review of similar sentiment analysis systems makes entirely explicit: the

demand that inner states be measured and used to determine the conditions of labor now applies to the white-collar workforce every bit as much as it does to checkout clerks or line workers. As well, in a theme that we’ll be taking up repeatedly, what is salient is not so much whether these tools actually perform as advertised, but whether users can be induced to believe that they do. The prejudicial findings of such “HR analytics,” i.e. that a given employee is unreliable, costly or a litigation risk, may be acted upon even if the algorithm that produced them is garbage and the data little better than noise.

Though top executives may, for the time being, manage to wriggle free from algorithmic performance evaluation, it seems highly likely that employment at all other levels will become increasingly contingent on a continuously iterated double articulation of assessment and selection that leaves no room whatsoever for the distracted, the halfway-competent, the deliberately shirking or simply the different.

Everything that wearables and workstation systems do for the shop floor, the call center and the checkout counter, analytic suites like BetterWorks do for the management echelon.⁴¹ The same anomaly detection subroutines that identify when a customer service representative’s average call length or escalation rate has climbed past the permissible value can trivially discover when someone is away from their desk too long, taking over-frequent bathroom breaks or gossiping with friends in other departments (never mind that such interdepartmental contacts are how organizations break through groupthink and actually innovate).⁴²

As employee monitoring (and self-monitoring) inexorably advances across the enterprise, the data it generates won’t simply evaporate. It will pile up in drifts, with measurements characterizing every last aspect of job performance in exquisitely high resolution rapidly accumulating in the corporate servers. It will, of course, make a tempting target for theft and exploitation—and if the past is any guide, the obligation to exert some kind of fiduciary responsibility over this cache won’t be felt nearly as urgently as it should be.⁴³

But something else happens as well, whenever a volume of data this large is gathered: feature extraction becomes easier, the development of training sets more straightforward, unsupervised learning a very real possibility. And this brings us full circle to the training of algorithmic systems, and their eventual deployment in replacement of human workers.

These disciplinary techniques will no doubt continue to be relied upon into the indefinite future, and they will continue to be responsible for their share of misery. But they are not the processes that are relentlessly and progressively reducing the scope of human labor.

Behind the wheel, the learning algorithm and the multispectral sensor do for the human driver, moving vehicles across the land more swiftly, more reliably, and more safely than the most responsible flesh-and-blood operator, whether they are charged with the transportation of freight or of passengers.

In the warehouse and the loading dock, the standardized container, the RFID tag, the autonomous pallet sled, and the development of sprawling big-box distribution centers that are very little other than enormous robotic systems in themselves—all these things sooner or later settle a quietus onto the prospect of unskilled employment in the fulfillment and logistics sector.

At retail, “seamless” point-of-sale processes and the displacement of responsibility onto the shopper themselves via self-checkout slash the number of personnel it takes to run a storefront operation, though some staff will always be required to smooth out the inevitable fiascos; perhaps a few high-end boutiques performatively, conspicuously retain a significant floor presence.

In customer service, appalling “cognitive agents” take the place of front-line staff.⁴⁴ Equipped with speech recognition and natural-language processing capabilities, with synthetic virtual faces that unhesitatingly fold in every last kind of problematic assumption about gender and ethnicity, they’re so cheap that it’s hard to imagine demanding, hard-to-train human staff holding out against them for very long.

Even in so-called high-touch fields like childcare and home-health assistance, jobs that might be done and done well by people with no other qualification, face the prospect of elimination. Xenophobia and racism subtly, if predictably, shape the possibilities here: in sector after sector, from healthcare to farming,⁴⁵ a Japan that is rapidly shrinking and aging would rather invest in developing advanced (and often specifically humanoid) robotics than admit an immigrant labor force of any significant size. There are always choices, and this is the one that Japanese society has made—but the techniques and conventions that are developed as a consequence of this choice will find purchase far beyond its shores.

In the military, in sex work, in eldercare, in one domain after another where you’d think the law or good taste or common sense might prevent someone from proposing the automation of a historically human function, the effort proceeds apace. If we can judge fairly from the statistics we’re offered, or the things that CEOs say in unguarded moments, automation is already sweeping across the economy at its foundations, taking up entry-level jobs and popping them one by one like blisters in a strip of bubble wrap. The presumption that was until very recently not merely tenable but persuasive—that even the otherwise-untrained had something unique and definitive to offer on the job, whether that thing was cognitive or affective or empathic—begins to yield before the cold equations of capital. Something new looms into view.

One of the very few points on which I am likely to agree with former US secretary of the treasury Larry Summers is his reminder that “[n]o one should speak with certainty about these matters, because there are challenges in the statistics, and there are conflicts in the data.”⁴⁶ It would be easy to overstate the impact of automation on entry-level employment, especially in the near term.

But as someone profoundly skeptical of the claims that are so often and so breathlessly made about technology—as someone who knows from personal experience full well how hard it is not so much to develop, but to deploy, integrate and make use of

technological systems—what I've seen in the course of research for this book has convinced me that automation is an existential mid-term threat to the livelihood of the most vulnerable workers. Capability is advancing more quickly than most realize, or are prepared to accept.

The dirty, dull, dangerous and demeaning jobs will be the first to be automated, and these remain, for better or worse, precisely the reservoir of opportunity for unskilled, undocumented or otherwise marginalized participants in the workforce.

What this does to the culture of work, and to labor's already imperiled ability to make demands and specify the conditions under which it produces value, deserves treatment at book length. For the present purposes, it seems safe to conclude that between algorithmic management and regulation, and the more than usually exploitative relations that we can see resulting from it,⁴⁷ hard times are coming for those who have nothing to offer the economy but their muscle, their heart or their sex. I don't doubt that those who benefit from any such state of affairs will be able to focus the rage of the permanently unemployed on immigrants and other convenient scapegoats, but eventually they too will be compelled to seek some sort of *modus vivendi*.

The tacit assumption that has held sway in the developed world for most of the past century, and which has only deepened its hold since women began to enter the workforce in great numbers around 1950, is that whatever their degree of skill or education, the vast majority of adults will eventually be able to find employment in the formal economy. And this fact structures virtually everything about everyday experience, not merely economically, but materially, psychically and socially.

It's the logic beneath the twice-daily surge we call the commute; the factor shaping the great commuter hubs of the world; the reason why there is such a thing as "rush hour" at all. It shapes the width of sidewalks, drives the content of daytime television and underwrites the viability of businesses whose only proposition is that they provide a "third place" between work

and home. It molds the way we dress, the way we eat, the skills we choose to acquire and the hobbies we undertake. It largely determines our social status. Not least, it places bounds on the way we present ourselves and the things we feel able to say out loud.

At this point, a garden-variety pop futurist would commit to the flat declarative: All of this is going to change. In the advance of automation, there will be very little that is meaningful left for anyone to do. The point will be reiterated, made again for the folks who were texting or otherwise tuned out the first time around: jobs are going away. You Better Get Ready.

The temptation, when presented with one of these would-be business oracles, is always to hiss with suppressed laughter at the shallowness of the argument and the simplicity of whatever "takeaway" they offer. But what if they're even halfway right? What if the thing that furnishes the economy with its basic structure, and day-to-day experience with its most fundamental organizing principle, is starting to erode before our eyes?

In their recent book *Inventing the Future*, Nick Srnicek and Alex Williams perceive in this set of circumstances an epochal opportunity for the left. Their argument, broadly, is that going forward, there simply won't be enough meaningful work to furnish a global labor force of five billion or more with employment capable of sustaining them—and that it is in any event perverse to defend jobs we know full well to be bullshit.⁴⁸ Instead of squandering energies in the sentimental defense of a proletarian way of life that no longer corresponds to any set of facts on the ground, they propose that there is a far more valuable effort progressive forces could dedicate themselves to at this moment in history: the struggle for a universal basic income, or UBI.

As the name suggests, most UBI plans—and the variants are many—propose that the state furnish all of its citizens with some kind of sustaining stipend, regardless of means tests or other qualifications. Most versions propose a grant at least equal to the local poverty line, in theory liberating recipients from the worst of the want and gnawing fear that might otherwise beset them in a time of mass unemployment.

The UBI is by no means strictly an argument from the utopian fringes of the left. Indeed, the terminology “basic income” itself is a market-friendly reframing of something that Fabian socialists used to think of as the “social dividend.”⁴⁹ In the United States, no guaranteed annual income initiative has ever gone further than the Nixon Administration’s 1969 experiments with the unfortunately named Family Assistance Plan (FAP).⁵⁰ More recently, the Dutch city of Utrecht and the Finnish national government have recently trialled similar measures under the market-liberal VVD party and Juha Sipilä’s center-right coalition, respectively.⁵¹ (A June 2016 Swiss UBI initiative failed, with 76.9 percent of voters opposed.)

With support across the conventional political spectrum, it may seem like some kind of UBI is far and away the most cogent response to widespread automation we have available, a cushion to buffer those who might otherwise plummet to Earth.

But the devil always nestles in the details of any such proposal. Held up to sustained inspection, the UBI can often seem like little more than a neoliberal giveaway. Its proponents on the market right clearly anticipate it as a pretext to do away with existing benefits, siphoning whatever transfers are involved in it back into the economy as fees for a wide variety of educational, healthcare and family services that are now furnished via social provision. And whichever direction it comes from, arguing for the accelerated disappearance of work is a very high-stakes gamble to make, in a world where the welfare state and its safety net are distant and receding memories and the horizontal and mutualist infrastructures that might replace them have not had time to develop.

One could, therefore, be forgiven for concluding that in practical terms, the achievement of a universal basic income will result not in anything like total leisure and unlimited self-actualization, but in the further entrenchment of desperation and precarity. When far more powerful forces are already waiting to exploit its emergence and divert its flows for their own ends, it seems unnecessarily cavalier of people who think of themselves as being on the left to “demand” a generic UBI. To

the degree that we buy the Srnicek and Williams line, what we need to insist on is the implementation of income guarantees in a context that protects our ability to spend that windfall as we see fit.

Other, less central aspects of the visions we’re presented of a world without work might trouble us as well. In his book *The Zero Marginal Cost Society*, Jeremy Rifkin fetishizes fully automated logistics, without considering how often logistics workers specifically have constituted the most radical faction of industrial labor. Without for a moment romanticizing the circumstances that gave rise to their militancy, we might want to remember how frequently in the past it’s been workers toiling in the most oppressive industries who have offered themselves as the insurgent brake on unfettered capital accumulation. On our way to a world of total automation, we may often have time to contemplate what a society winds up looking like when its most mutinous voices have fallen silent.

It can be briefly amusing to stand alongside the accelerationists as they proclaim, “Workers of every nation unite! You have nothing to lose but your jobs!” But the euphoria soon fades, swept away by the swiftly sobering recognition that there is terribly little chance for a soft landing in any of this, for any one of us.

Though we may debate the degree to which choice and conscious authorship are involved in it, it seems important to note that automation is a directional process whose initial stages we’ve already entered. In this respect David Graeber’s empty, signifier-shuffling “bullshit jobs” are a signal from the future. They’re not so much a return as an anticipation of the repressed: the surfacing in the present, and pricing into contemporary ways of doing and being, of the recognition that there simply won’t be enough meaningful work for anyone to do following the eclipse of human judgment.

The UBI, as a rearguard action aimed at ensuring the continuity of business as usual, seems to be predicated on the survival of at least some higher-order executive jobs. But what its advocates neglect to consider is how total the process of supplantation will

be. If we allow anything like this to happen, it will be a wave that sweeps any notion of consequential work away, for the wealth managers and creative directors and project managers as much as the truck drivers and pipefitters. What we will discover, I think, is that we urgently need to reinvent (particularly, but not just) a left politics whose every fundamental term has been transformed: a politics of far-reaching solidarity, capable of sustaining and lending nobility to all the members of a near-universal *unnecessariat*.

But we will also discover something else. We needed work, though not in any hackneyed dignity-of-labor sense. We certainly didn't need to rouse ourselves to Stakhanovite exertions on behalf of uncaring employers; we didn't benefit from being forced to simulate team-spiritedness, amidst all the profoundly dispiriting banality of our fluorescent-lit cubicles; and it goes without saying we didn't need to suffer the insults of cretinous customers, working out their neuroses and class frustrations across the counter. All too often work cost us our health, our dreams, our lives. But it also offered us a context in which we might organize our skills and talents, it gave us some measure of common cause with others who labored under similar conditions, across all bounds of space and time, and if nothing else it filled the hours of our days on Earth. Though these goods came at far too high a price, I don't know that we are wise to consider living entirely without them, or are practically prepared to do so.

Even under the very best of foreseeable circumstances, circumstances in which we are successfully able to organize the infrastructures of solidarity we will need to call upon for our sustenance once work has gone away, I don't know that we're psychically equipped to withstand total freedom from obligation.⁵² As feeling beings, we become habituated to any emotion long and persistently experienced, and eventually desensitized to its appearance. Some degree of variation, effort and friction—in other words, some measure of its absence—is necessary to the experience of pleasure.⁵³ In a world without work, as Keynes suggested almost a century ago, we dispense with most all of the

friction and effort bound up with the struggle to survive, and therefore finally and implacably arrive at the question of what life is for. At a time when the alienation and anomie attendant upon the withdrawal of work is literally killing (some of) us⁵⁴, perhaps it's worth considering what the answers might cost us.

Some researchers insist that in the aftermath of total automation, we will learn to value people not for the advantage we're able to generate through work, but for who we are inherently, as unique and irreplaceable individuals. I leave that thought here without further comment.

In the end, it's difficult to stand back far enough to weigh the meaning of the changes rippling across this thing we call the economy, something so total in our lives that its outer bounds are virtually impossible to perceive. What we can see is that for all the sorriest, most predictable and all-too-human of reasons, human discretion is progressively becoming decentered within it.

As algorithms develop the ability to plan optimal courses of action through mind-bendingly complex multivariate decision spaces, and blockchain-based distributed autonomous organizations the ability to capture and organize wealth, it seems as though human intention is sure to follow. Such decenterings may not be particularly upsetting to anyone who's ever sat with a set of questions we inherit from feminist, post-structuralist and ecological thought: questions about the death of "Man," the agency of nonhuman actors, the consequences of our decisions for the other life we share a planet with, and the extreme unwisdom of trying to articulate some kind of binary distinction about the boundary between humanity and its others in the first place.

I believe questions like these come from the best that's in us, and I cannot imagine how impoverished we would be if nobody had ever thought to press them. But it remains difficult for me, at least, to conceive what an economy might be for, if not the generation and apportionment of wealth as humans experience it.