Philosophy & Technology (2023) 36:36
https://doi.org/10.1007/s13347-023-00631-w

RESEARCH ARTICLE

The Future of Work: Augmentation or Stunting?

Markus Furendal[1] · **Karim Jebari**[2]

Received: 13 December 2022 / Accepted: 3 April 2023 / Published online: 4 May 2023
© The Author(s) 2023

Abstract

The last decade has seen significant improvements in artificial intelligence (AI) technologies, including robotics, machine vision, speech recognition, and text generation. Increasing automation will undoubtedly affect the future of work, and discussions on how the development of AI in the workplace will impact labor markets often include two scenarios: (1) labor replacement and (2) labor enabling. The former involves replacing workers with machines, while the latter assumes that human–machine cooperation can significantly improve worker productivity. In this context, it is often argued that (1) could lead to mass unemployment and that (2) therefore would be more desirable. We argue, however, that the labor-enabling scenario conflates two distinct possibilities. On the one hand, technology can increase productivity while also promoting "the goods of work," such as the opportunity to pursue excellence, experience a sense of community, and contribute to society (human augmentation). On the other hand, higher productivity can also be achieved in a way that reduces opportunities for the "goods of work" and/or increases "the bads of work," such as injury, reduced physical and mental health, reduction of autonomy, privacy, and human dignity (human stunting). We outline the differences of these outcomes and discuss the implications for the labor market in the context of contemporaneous discussions on the value of work and human wellbeing.

Keywords The future of work · Automation · Human-AI cooperation · The goods of work · Work at Amazon · The ethics of work · The ethics of technology · Autonomy · Surveillance capitalism

✉ Markus Furendal
markus.furendal@statsvet.su.se

Karim Jebari
karim.jebari@iffs.se

[1] Department of Political Science, Stockholm University, Stockholm, Sweden

[2] Institute for Futures Studies, Stockholm, Sweden

1 Introduction

The last decade has seen a rapid improvement of a set of technologies referred to as "artificial intelligence" (AI). The increasing capacities of AI technology have, in turn, led to a renewed discourse about the future of work and automation (Acemoglu & Restrepo, 2020; Frey & Osborne, 2017; Paus, 2018).[1] While many have worried that the implementation of this technology will bring about mass unemployment (Harari, 2017), others have pointed out that if societies ensured access to a social minimum to everyone, the productivity increases that AI enables could open the door to an attractive, post-work existence (Danaher, 2019; Floridi, 2014). Yet, there are also reasons to be skeptical of the likelihood of technological unemployment, since AI technology may not only replace but also boost human productive capacities (Willcocks, 2020). Although computers now beat even the best humans at chess, human–computer teams known as "centaurs" could—until recently—beat solo-playing computers (Kasparov & Greengard, 2017). By analogy, we could expect that the combined efforts of human-AI teams can combine the capacities of each and create a new kind of highly productive worker able to compete with AI tools (Willcocks, 2020). The main thesis of this article, however, is that this simple view overlooks the fact that human-AI cooperation can work in two distinct ways: either by *augmenting* human capabilities, or by *stunting* human capabilities. While both kinds of cooperation lead to higher productivity, the latter, we will argue, also prevents workers from attaining the *goods of work* while also bringing about significant *bads of work*. Illustrating this with a case study of how AI technology is implemented in the warehouses of the pioneering company Amazon, we suggest that, while it is likely that AI technology will be able to increase the productivity of many workers, it might also impair their working conditions in the process.

The next Section (2) introduces the notions of labor replacing and labor enabling AI technologies and summarizes the debate on AI-driven automation. Section 3 then presents a framework for conceptualizing what the goods and bads of work are in modern societies. Section 4 illustrates how labor enabling AI could improve people's experience of work and the goods they can access through their productive activity. Section 5 then considers how AI is already implemented in Amazon warehouses in a manner that rather stunts human capabilities. Four objections to our analysis are considered in Section 6, before Section 7 concludes.

[1] Here, by "work" we refer to free wage labor, i.e., not subsistence labor, domestic labor (if unpaid) or slave labor. This excludes other forms of work, like unpaid care work, which both create value and carry burdens and benefits. Yet, while other forms of labor are important to understand and analyze in the context of AI-driven stunting, the relational and power dynamics in these cases are different from that of wage labor, and would require a distinct analytical framework. For example, it could be argued that slave labor can never contribute to the goods of life, and that such labor is inherently stunting. For the purposes of this paper, we thus focus narrowly on free wage labor. See also Smids et al. (2020, p. 506).

2 Technological Innovation and Work

By "automation" we mean, in this context, a process by which machines increase labor productivity, i.e., the value of produced goods and services per work hour, by performing some or all tasks that were previously done by humans. In other words, automation happens when the set of tasks that can be produced with capital rather than labor expands (Acemoglu & Restrepo, 2018, p. 5). For this reason, we will not discuss the process by which machines merely allow a task to be performed by a worker with less pay, for example by outsourcing or by using child labor, unless that process also makes the worker more productive. While such a process may increase the value of goods relative to wages, it does not increase the value of goods relative to *working hours*. We also do not refer to innovations that merely shift the execution of a task from a worker to a third party—like when customers use digital check-in counters at airports—since the task is not performed by capital but merely a different (unpaid) person. Although this frees up time for paid airline employees, it does not necessarily increase overall productivity.[2] While innovations that shift the workload from employees to customers can increase the profits of a particular company, these do not contribute to a society's general long-term economic growth, as real productivity enhancing technology does. This is why productivity enhancing technology is of special importance.

Innovations have enabled increases in labor productivity for as long as there has been work. A recurring worry, however, has been that when machines perform the work previously done by people, demand for labor will drop and unemployment will rise. A common rebuttal to this worry is that it is a mistake to believe there is simply a fixed amount of work that needs to be done in a given economy (the "lump of labor fallacy"). Optimists instead note that even though automation may render some jobs superfluous, the increased productivity will help the economy grow, thereby increasing demand for labor overall (Acemoglu & Restrepo, 2018). Regarding our current situation, and the future of automation, various leading theorists envisage two broad possibilities for how this could happen (Frey, 2019, p. 13). The *labor replacing scenario* implies that we primarily substitute machines for human labor, rendering human workers redundant, as, for example, elevator operators became redundant by improvements in elevator safety. This would be undesirable, it is argued, as technological unemployment might cause popular discontent. In turn, it might also lead to political and social obstructions of technological progress (Gallego & Kurer, 2022).[3] By contrast, the *labor enabling scenario* implies that machines primarily make human workers more productive, by creating the kind of centaurs mentioned above. For example, office workers have become more productive due to word processing software (at least

[2] While machines have historically also replaced the work of non-human animals, for example when the steam engine replaced horse mills, this is not the focus of our discussion.

[3] As we have noted, however, some authors have also argued that labor replacement might be a good outcome, both because many current jobs are detrimental to the kinds of goods we discuss here, and because it could make society less work-centric, and allow people to pursue the goods outside of paid employment. (Cf. Danaher, 2019) We discuss this point in Section 6, below.

with regards to the quantity of written output) and stand to become even more productive in the near future as generative AI technologies, such as chatbots, are integrated into the same software.[4] Similarly, since machines are well suited to perform dull and repetitive tasks, AI automation could make both manual and office workers more productive and potentially more satisfied.

Thus, the labor enabling scenario is often described as a positive outcome (Brynjolfsson & McAfee, 2014), where increased productivity could lead to either higher wages or cheaper/better products and services (Agrawal et al., 2018). The economic historian Carl Benedikt Frey suggests that, unlike labor-replacing technologies, labor-enabling technology has historically been accepted by workers and allowed for productivity increases without political instability (Frey, 2019, pp. 131–137). While some authors have raised the possibility that AI-technology will aggravate economic inequality, this possibility has been seen as having a relatively straightforward solution: progressive taxation could both reduce inequality and make a strong social safety net and a robust welfare state possible (Field, 2019). Consequently, Frey and other theorists argue, policy makers should enact reforms to encourage the development of technologies that lead us down the path of labor enabling rather than the path of labor replacement.[5] This scenario is also popular among AI-developing companies. It is commonplace for large tech companies to talk of AI ethics and notions like "responsible AI," and many companies have internal research and policy development around how such an ideal could be achieved (de Laat, 2021). Large tech companies, including Amazon, DeepMind, and Microsoft also founded the international non-profit "Partnership on AI" in 2016. By its own description, this organization aims to support policymaking around AI, including seeking "shared answers and recommendations for actionable steps that need to be taken to ensure AI supports an inclusive economic future." (PAI, 2023). What ambitions like these entail is not always clear, and such initiatives are sometimes accused of being a kind of "ethics washing," staving off regulation and marginalizing issues that do not fit the corporate agenda (Bietti, 2020). Early results from studies of such initiatives suggest that they can nevertheless have beneficial effects, but that it is often too early to tell whether this is actually the case (de Laat, 2021).

[4] It should be noted that Frey and others are not entirely clear about the distinction. In many cases, labor enabling technologies allow fewer workers to perform tasks that require many workers. Consequently, this leads to some degree of "replacement.".

[5] The distinction may seem somewhat ambiguous since some technologies both replace workers and create new jobs. It is important to consider the level of analysis, however: a piece of technology such as the automobile may be labor-replacing in one sector—such as the horse-based transportation sector in the early twentieth century—but labor-enabling at the general level of the economy since it allows for entirely new sectors to emerge. For the purposes of this paper, we focus on how AI can be used in labor-enabling ways in particular workplaces, and will not address the general issue of labor-replacement effects at the societal level. Now, as we go on to discuss, a piece of labor-enabling technology could lead an employer to reduce the number of work-hours per unit of output, which could lead to layoffs and/or fewer recruits in the future. However, since our main aim is to analyze the distinction between human augmentation and human stunting introduced below, this is not a major issue in the present context.

A problem with analyzing the impact of AI on the future of work in the aggregate, however, is that economic research suggests that technical innovations impact different kinds of work differently. In the 1970s, Harry Braverman proposed that mechanization and the increased division of labor led to a routinization of many tasks. Instead of technology increasing demand for skilled workers, machines could be operated by a less skilled workforce, whose jobs consequently were less rewarding than before (Braverman, 1998). Similarly, many Western countries have since experienced a kind of job polarization, where the share of workers employed in the highest-skilled and the lowest-skilled occupations grow, while middle-skilled jobs such as manufacturing decline (Goos et al., 2009). When considering the effects of automation, it is hence important to note, not only how machines affect the aggregated output of an economy, but also how they impact the quality and character of individual jobs. While the macroeconomic perspective undoubtedly tells us something important, a political-philosophical analysis will often stress the way in which the burdens and benefits are distributed within and across groups (see also Santoni de Sio et al., 2021).

In light of these insights, we will argue that the discussion on the future of work must recognize not only the distinction between labor-replacing and labor-enabling scenarios. A further distinction is needed, between two ways in which labor-enabling technologies can increase productivity: *human augmentation* technologies which promote workers' pursuit of a variety of objective goods primarily accessible through one's work (see Scanlon, 1998 for a discussion on objective goods); and *human stunting* technologies, which conversely increase productivity in ways that make the job worse for the individual workers, either by harming workers in some way or by preventing the pursuit of the work-related objective goods.

Developing this distinction allows us to better understand AI-driven disruption of work, but it also constitutes a contribution to the rapidly expanding multidisciplinary research effort to understand and study moral and political aspects of the future of work. In the last few years, researchers have addressed issues like the globalization of work, power relations between employer and employee, the gig economy, and precarious work.[6] Others have focused specifically on developing theories about, or empirically studying, how digitalization, robotization, and AI may impact the quality of work, or workers' chance to experience meaningful work. Smids et al. (2020), for instance, identify a number of aspects connected to meaningful work and discuss how the implementation of robots (broadly defined) may impact people's chance to achieve this. Similarly, both Bankins and Formosa (2023) and Hosseini et al. (2023) discuss technology and meaningful work in light of the ethical principles of beneficence, non-maleficence, autonomy, justice, and either explicability or responsibility. Berkers et al. (2022) offer an exploratory comparative study of the role of robotization in work design, and what the motives are for adopting new technology. A striking fact is that just like our paper, much of this research focuses on logistics warehouses. The main reason, we suspect, is that the warehouse is a confined and

[6] See, for instance, Bieber and Moggia (2021), and Vredenburgh (2022), Smids et al. (2020), and Hsieh (2008).

controlled space in which it is much easier to adopt special automation technologies than in other, more chaotic, environments. This allows automation to progress at a higher pace in this sector, and thus makes it a testing ground which is useful to study. It also indicates that there may be benefits from widening this field and studying other kinds of use-cases and sectors in future studies.

Although there are similarities between our approach and articles such as Smids et al. (2020), we believe that the concept of centaurs, and the analysis of the stunting-augmenting distinction that we offer provide some new theoretical insights. First, as we discuss below, the rate at which AI is currently developing suggests that many more human workers soon could be classified as centaurs, forming teams with generative AI systems that can create text, images, video, or audio. This makes it crucial to understand the issues raised when employers expect human workers to engage in such cooperation, not least because their adoption is more likely to be guided by efficiency concerns rather than ideas about how to promote the goods of work.[7] We believe that the notion of centaurs is apt to describe this particular kind of human–machine interaction, and argue below that policy makers should be aware of the risks that labor-enabling policy could inadvertently promote human stunting if not explicitly designed to avoid this outcome. Second, although more attention is being paid to different ways in which new technologies can improve or frustrate workers' opportunities to attain the goods of work, we believe that the specific way in which we conceptualize what is at stake complements the principle-based suggestions offered in similar accounts. There is a value, we suggest, in our complementary analysis of what meaningful work is and how technology affects it. Relatedly, although we do not offer a full analysis of what explains technology adoption, we believe that one benefit of our analysis is that it is sensitive to the fact that whether AI is stunting or augmenting depends not only on the available technology, but also on power relations in the workplace. We believe future studies could help illuminate this issue further. Third, and perhaps most importantly, the final section on objections to our account addresses concerns one may have about many of the recent analyses of automation and the future of work. To our knowledge, most other contributions to this emerging literature have tended to stop short of acknowledging and engaging with such critical views, and the section hence fills an important role in countering objections that could be raised against the need to worry about human stunting and related concepts.

3 Work Matters—on the Goods and Bads of Work

The distinction between human augmentation and human stunting presupposes that work can be more or less "good" or "meaningful." This section explains what we mean by this. First, jobs that are dangerous and can result in bodily injury, mental

[7] As one warehouse manager noted regarding such decisions, "The focus is on its potential savings and not the consequences for people […] I do not think that we would reject an option that would save us hundred thousand euros if it would potentially make our people less happy. We would be the first company to do so." Cited in Berkers et al. (2022).

illness, or decreased welfare of those who perform them, are quite clearly bad for these specific reasons. Beyond this, a number of suggestions have been made as to what the value of work is.

Karl Marx, perhaps the most influential thinker on the value of work, argues that the "species-essence" of humans is to be productive and that the right kind of work is good for us as individuals. Specifically, work under capitalist forms of ownership and production creates alienation, by not involving the development and deployment of essential human powers (Marx, 1999, chap. 2; Wolff & Leopold, 2021).

In an early essay in the contemporary academic debate about meaningful work, Adina Schwartz similarly argues that individuals who work in monotonous and routine jobs will be prevented from rationally forming and acting on autonomous decisions about their lives (Schwartz, 1982, p. 639). In the same vein, Ruth Yeoman has argued that meaningful work is a fundamental human need, necessary to realize not only the value of autonomy but freedom and dignity as well (Yeoman, 2014).

To make our central argument compatible with as wide a set of theoretical assumptions as possible, our analysis will not be tied to a specific conception of flourishing or autonomy. Instead, we will draw on work by Anca Gheaus and Lisa Herzog, who have argued that the goods (and, conversely, bads) of work "…have distinctive value because they are very widely desired by individuals as non-substitutable elements of their conception of the good life and […] they cannot be obtained with money." (Gheaus & Herzog, 2016, p. 79).[8] Importantly, we will follow Gheaus and Herzog in taking certain aspects of how societies are normally arranged as a given: As it stands, wage labor is for an overwhelming majority of persons the only source of a sufficient income, and most people consequently lack the opportunity to refuse to work altogether. This means that the bads of work are difficult or impossible to avoid in our society for most citizens that do not belong to the wealthy. Yet the bads of work are not an intrinsic part of any kind of work. Conversely, the goods achieved through work in our current society could potentially be pursued outside of work in a less work-centric society. However, given that most people are required to spend so much of their waking time at work, they currently lack the time to pursue the goods of work outside of the realm of work. This is especially true for some forms of precarious work, when employers can on short notice decide work schedules and dictate work hours. Moreover, the ability to pursue goods outside of work may also be more difficult for workers whose jobs also fail to provide these goods. For example, if a job causes injury, pain, or frequent exposure to disease, then this can also make pursuing excellence and other goods more difficult outside of work. The possibility to pursue goods usually attained through work outside of work is also reduced significantly as the amount of time that one has to work increases. It is not rare that individuals need to work more than 40 hours a week to provide for themselves.

[8] There are a number of theories exploring what constitutes good work that we cannot discuss here, including Arnold (2012), Veltman (2016), Roessler (2012), and Gomberg (2016).

Gheaus and Herzog identify four distinct kinds of goods. First, work could offer a chance to *pursue excellence* by developing one's skills. Think of a skilled craftsperson, for instance, honing his or her talents and utilizing them to create a refined product. Second, given the many institutions and norms that induce people to work, there are several values that are achievable mainly in one's work but more difficult to pursue if one is unemployed. Gheaus and Herzog cite sociological evidence that many people value the chance to *make a social contribution* through their work. Think, for instance, of the jobs that during the COVID pandemic were called "essential," like nursing, teaching, or waste collection.[9] Third, and relatedly, Gheaus and Herzog argue that the chance of receiving *social recognition* is intimately tied to performing paid work, and they suggest that both the kind of work one does and one's place in the job hierarchy (and, often, one's income), have a large impact on the kind of recognition one receives in society more generally. Finally, Gheaus and Herzog suggest that people working together can seek a kind of *community*, in the sense of "…the experience of doing things together with people with whom they stand in relatively free and equal relationships." (Gheaus & Herzog, 2016, p. 76). Certain ways in which work is organized are more conducive to realizing this value, and the authors suggest that non-hierarchical organizations like worker cooperatives might be preferable to more traditional forms of workplaces.

Gheaus and Herzog also note that jobs might create "bads" when they are organized in ways that prevent the pursuit of the goods just identified, or when they undermine or fail to protect other goods, "such as the health of employees, their discretionary time and, in cases in which jobs are structured by oppressive hierarchies, their freedom from non-domination." (Gheaus & Herzog, 2016, p. 73). Boring and repetitive jobs that require little skill prevent workers from pursuing excellence, for instance, and workers are unlikely to receive social recognition from others for engaging in it.

In addition to what Gheaus and Herzog lists, we suggest that there are a few and fairly straightforward ways in which jobs can be bad: (1) when they involve a substantial risk of injury, disease, or mental illness (for example, PTSD); (2) when workers are subject to intrusive surveillance; (3) when workers have to unflinchingly obey arbitrary rules (for example, not being allowed to yawn) (Vincent, 2021); and (4) when work tasks violate shared norms of human dignity (Sainato, 2019) (for example, not being allowed to visit the bathroom for extended periods of time, forcing workers to soil themselves).

In sum, this leads us to the view that work is bad when, and because, it exposes workers to the risk of being physically injured or negatively impacted psychologically, and when it is organized in ways that prevent workers from pursuing the four types of goods of work mentioned above. Our analysis will consequently proceed under the assumption that technology is used in ways that are human augmenting when it increases productivity in ways that promote, or at least do not prevent, the

[9] Although the pandemic catalyzed a greater appreciation and outpouring of public support for people with "essential jobs," this has not really changed the fact that they are seldom the most highly remunerated. On the limitations of relying on markets to decide the "social value" of one's contribution, see Furendal (2019).

pursuit of the goods of work, and it does so without subjecting workers to physical or psychological injury. Conversely, productivity-enhancing technology leads to human stunting when it involves reorganizing work or transforming tasks in ways that prevent the pursuit of the goods of work, or bring about more of the bads of work.

4 Human Augmentation

The former world champion in chess, Garry Kasparov, introduced a form of chess to a wider audience that has been referred to as "centaur chess," where a human and a machine compete as a team ("centaur" as an analogy of the mythological creature that is half-human and half-horse) against another human–machine team (Kasparov & Greengard, 2017). The idea is that by combining human and machine capabilities, centaur chess can increase the level of play by producing blunder-free games with the qualities and the beauty of both perfect tactical play and highly meaningful strategic plans.

Centaur chess represents an example of human augmentation. Human augmentation will often allow humans to specialize in focusing on the bigger picture, creating a space for human task solving that is more creative, interesting, and rewarding. These kinds of human–machine teams have been held up as one reason to be optimistic about the relevance of humans in a future where machines are able to perform most tasks better than any particular human (Brynjolfsson & McAfee, 2012; Cowen, 2013). Ideally, we would leave all dirty, dangerous, and dull jobs for the machines, and let them do the tedious and unsafe tasks of other jobs, allowing humans to focus on the tasks that require dexterity, manual skill, creativity, the ability to plan, to empathize, etc. At a general level, human augmentation could reduce the extent to which people are exposed to the bads of work, such as the risk of injury and repetitiveness. Human-augmenting technology could also enable promotion of the goods of work. First, as machines carry out dull and repetitive tasks, humans can concentrate on tasks that require skill and develop those skills. Second, as machines improve labor productivity, it allows human workers to make a greater contribution to society. Thus, human-augmenting technology at work allows individuals greater opportunities to express their skills and their humanity in a way that can have a substantial effect on developing capabilities.

Consider, for instance, how the invention of the electronic spreadsheet transformed the job of being an accountant. With the introduction of programs like VisiCalc and Excel, firms no longer needed to employ people to recalculate whole spreadsheets with pen, paper, and a pocket calculator every time the value in a cell changed. At the same time, the software also made the job accessible to others, beyond highly trained professionals, and enabled creative work processes that were essentially inconceivable before. One amazed accountant described it like this in 1984:

> Before, you would suggest a change to a client, get a staff member to calculate it, send it to the typist, to the proofreader, and recalculate it to make sure there

> weren't any errors. Now you have a machine right there with the client. Want to see what happens with a different return on investment? Sheltering? Interest rates changing by half of a percent? It's done in a minute. Before you'd be tempted to say, 'Let's leave it the way it was.' The whole mental attitude toward preparing projections has changed. (cited in Levy, 2014).

So, while the electronic spreadsheet made certain jobs in accounting firms obsolete, it also made qualified accounting services affordable for more organizations who could thereby employ more accountants. The job also became accessible to more workers, and more stimulating overall.

We expect that recent advances in generative AI technology, including large language models like GPT-4, will have similar effects, since a significant share of the tasks that office workers tend to do can be readily automated. These models can already work with the kinds of text that white-collar workers spend much of their time processing, including summarizing information, drafting messages, and checking for clarity and grammatical errors in existing text. Yet, since AI systems are far from infallible, it is likely that there will be need for a human in the loop to make sure that what the AI produces does not contain factual errors or violate norms. Indeed, such human-AI centaur teams are likely to execute white-collar tasks at a much higher rate and quality than either humans or AI systems on their own. There are reports of people using ChatGPT, for instance, as a form of creative partner to create first drafts that the human then edits and improves, or vice versa. Instead of merely cutting and pasting, an increasingly important skill among human workers is likely to be prompting, collecting, and combining AI-produced text (cf. Floridi & Chiriatti, 2020).

In other words, human augmentation technologies can improve work, not only by eliminating some of the bads of work, such as hazardous tasks, but also by increasing the goods of work, such as the opportunity to develop and pursue excellence in one's skills. We suspect that enthusiasm about the labor-enabling scenario often implicitly relies on assumptions that AI technology will primarily be labor-augmenting. It is perfectly possible, however, that the ways in which AI systems enable productivity increases will stunt human capabilities. As we will suggest presently, this is in fact how much of the already existing AI technologies function.

5 Human Stunting at Amazon's Fulfillment Centers

We suggested above that technology is used in ways that are stunting when it improves the productivity of workers while it simultaneously reduces the goods of work or brings about more of the bads of work for workers. Unlike human-augmenting AI-cooperation, human stunting hence increases productivity in ways that systematically either introduces bads of work or that reduces the goods of work. To illustrate how AI technology can be used for the purposes of human stunting, we will here describe how this kind of technology has been deployed at Amazon, one of the most innovative AI-developing companies. Amazon is known for its disruption of well-established industries through technological innovation and its move

toward automation is well publicized (Heater, 2022). It is the second-largest private employer in the USA and one of the world's most valuable companies, passing Walmart as the most valuable retailer in the USA by market capitalization in 2015 (Cheng, 2016; Streitfeld & Kantor, 2015). Although this case might not be representative of how AI-driven automation is currently implemented in the economy as a whole, looking closer at Amazon is motivated by the assumption that practices and technological applications used in the most innovative companies might spread and become commonplace down the line.[10]

The warehouses where Amazon stores the products that are delivered to people's homes are known as fulfillment centers (FCs). FCs are organized as multi floor towers lined with shelves divided into cells stocked with products. Here, workers called "stowers" shelve products, while "pickers" retrieve them. Cells and products are identified through barcodes and inventory software. Notably, the company employs "chaotic storage," since stowing items more or less randomly allows for more efficient picking and use of storage space than sorting them by categories. Warehouse workers can thus be described as forming "centaurs" with their handheld barcode scanners, without which it would be impossible to know the location of any particular item. Scanners also monitor worker performance, and nudge them into keeping up the "Amazon pace," not quite running but walking as fast as possible between different areas of the FC (Delfanti, 2019). In more newly constructed FCs, however, workers occupy fixed workstations, and mobile shelves are brought to them. This suggests that Amazon is trying to improve productivity by reducing the time spent walking between different cells (Delfanti & Frey, 2021, p. 656).[11]

5.1 The Bads of Work

During the last few years, FCs have seen a substantial increase in workplace injuries per work hour, and Amazon warehouses have significantly higher rates of serious injury than warehouses run by other companies (Greene & Alcantara, 2021). An analysis presented by the British trade union GMB suggests that ambulance calls to warehouses spike around big shopping dates such as Black Friday and Amazon's "Prime day." (GMB Union, 2021). Whereas workers in FCs with fixed shelves easily become injured by walking, often up to 20 km per shift, warehouses with mobile shelves and stationary workers bring other kinds of problems, and the increase of injuries has been the highest in the most automated warehouses (Jamieson, 2015). Robotic tools have enabled a work environment where workers stand in cage-like

[10] Many of the issues we discuss in relation to Amazon have also been reported to occur in other sectors of the economy, including among the micro workers that perform the tedious task of preparing data on which to train AI models, or to complement AI systems when they cannot be trusted. See (Jones, 2021).

[11] The company recently revealed a fully autonomous warehouse robot that can lift and move so-called "GoCarts," i.e., wheeled bins containing packages. Amazon's rationale explicitly refers to the human-augmenting scenario: "Our vision is to automate GoCart handling throughout the network, which will help reduce the need for people to manually move heavy objects through our facility and instead let them focus on more rewarding work." (Heater, 2022).

structures surrounding work stations and focus on a single task where they are superior to machines: manual dexterity. This has led to jobs with a more limited set of tasks, which increases repetitive movements and hazardous lifting. In other words, increased automation has likely directly increased injury rates at FCs. This example illustrates a case that conflicts with a common narrative in the literature: when AI is used to perform simple tasks, this enables workers to focus on more meaningful, creative, or otherwise rewarding tasks (Bankins & Formosa, 2023). In this case, robots in FCs carry out part of what used to be a worker's task, walking to a shelf and picking up packages. The result is that pickers now perform only a subset of the tasks they used to have, and that their job has become more repetitive and more dangerous.

The handheld scanner and other tools that track worker performance in real-time make high intensity workloads possible, and in some FCs pickers are expected to scan 400 items per hour. Workers report that missing such targets are detected by automated surveillance and result in write-ups, which in turn could lead to being fired (Sainato, 2020). Working under such stress has serious long-term health effects, and the turnover rate in FCs is consequently high. An anonymous worker describes that despite managers telling workers that good performance will enable them to access less taxing roles higher up in the hierarchy, it is in fact impossible to keep up the work pace: "If you are faster, never call in sick with a backache, after a while your back is gone, you have carpal tunnel, psoriasis caused by stress... and those [employees] are the first they set aside." (cited in Delfanti, 2019, p. 50). Similarly, pregnant workers have had miscarriages they ascribe to the highly stressful work environment (Abril & Harwell, 2021; Barbaro, 2018).[12]

Intrusive surveillance has a significant impact on behavior and the ability to exercise autonomy. A person that perceives that they are being watched becomes self-aware, cautious, and may find it difficult to engage in cognitively challenging tasks. Moreover, being under constant surveillance makes it impossible for us to control how we present ourselves to the world, which fundamentally undermines autonomy (Macnish, 2022). Schoeman (1984) argues that privacy provides a way to control intimate information about oneself and is crucial for the development of one's personality and inner self. At FCs, digital surveillance serves various productivity-enhancing purposes, including the enforcement of intrusive rules (e.g., how much water workers can drink, rules for social media posting during off-work hours) and social isolation (during COVID, workers were not allowed to come close to each other during breaks, a policy enforced by special purpose smart "proximity cameras"). Privacy violations can be distinguished into two types: (1) passive monitoring and (2) active surveillance. Whereas the former consists in passive collection of information, the latter combines this with active and constant feedback. FCs have both systems. For example, workers are assigned a certain amount of time to collect and register packages. When a worker fails to register the package in the allotted time window, instant feedback is produced on a screen in their workstation. Other forms of monitoring are more discrete. Tracking systems record worker productivity

[12] Note that in certain jurisdictions safety standards could mean that the employer is legally liable for what is described in some of these examples.

silently and are connected to an opaque score system that determines who gets promoted and who gets fired. A worker can be immediately terminated if their score falls under a certain target, the level of which is unknown to them (Delfanti, 2019, p. 50 f).

These close monitoring and opaque productivity measures raise two additional concerns, apart from the privacy issue. First, digital tools that make work more hazardous by being more repetitive also reduce autonomy as workers are limited to a small number of movements and are not expected to exercise planning, take initiatives, or learn new skills. Secondly, by strongly encouraging competition, it risks undermining other goods of work, such as the chance to cooperate and seek community with one's colleagues. Despite close surveillance, the relentless pace required encourages workers to cut corners when possible, and to avoid certain heavy and slow tasks that may reduce the number of items they can scan per hour.[13]

These aspects can also be understood in light of the concept of "surveillance capitalism," which Shoshana Zuboff suggests is a system centered around the collection of personal data for the purposes of control and profit making (Zuboff, 2015). While her analysis primarily focuses on how companies mine data from consumers, it arguably also helps illuminate how employers such as Amazon gather data about their workers. Surveillance capitalism, Zuboff argues, relies on opaque instruments to extract data, and continuously experiment with human behavior in a way that alienates humans while finessing techniques of behavioral prediction and modification. By being the subject of continuous behavioral experimentation and social manipulation, enabled by new technology, employees at Amazon are arguably commodified and dehumanized in ways that are more wide-ranging and visceral than consumers of Amazon's products, who are the main subject of Zuboff's analysis (Zuboff, 2019).

Although Amazon is a pioneering AI company and thus a special case, we believe lessons about human stunting at FCs may generalize to other forms of work. Our discussion above illustrates that AI systems are often not only technically opaque, but also make workplaces more opaque to workers, by enabling work that isolates workers (at FC workstations, or in the gig economy and microwork) or undermines their sense of control as they have to submit to being automatically evaluated according to a standard they cannot understand (Vredenburgh, 2022). It is easy to see how this may spread to other forms of workplaces, including white-collar jobs (Kantor & Sundaram, 2022).

5.2 Pursuing the Goods of Work

Apart from subjecting workers to the bads of work, the automation of Amazon FCs also prevents workers from pursuing the four goods of work identified by Gheaus and Herzog (2016). Clearly, there is little or no room for workers to *pursue excellence* when they are primarily performing repetitive and menial tasks. As one worker

[13] Thanks to an anonymous reviewer for the latter observation.

describes it, "You just need to follow the scanner, which tells you go here, go there, pick this, and pick that. You don't need to do anything else, don't need to think. Eight hours can last 24 hours because you are in a limbo." (Delfanti, 2019, p. 47). At the same time, workers are discouraged or prevented from using work time to develop their skills or pursue their interests by, for instance, listening to audio books or podcasts in headphones (Stuart, 2021). It is more difficult to determine whether automation prevents warehouse workers from having a sense that they can *make a social contribution* or *receive recognition* from others for the work they do. On the one hand, the point of the job is to expedite orders as quickly as possible. Managers hence regularly point out that working fast is valuable, "…thanks to you, many children will smile, you have brought joy in the homes of thousands of families." (Delfanti, 2019, p. 49). On the other hand, the standardization and simplification of tasks enables the interchangeability of each individual worker associated with Taylorism: Newly employed workers only need a few hours of training to become proficient pickers. Knowing that you do not bring anything to your workplace that could not be provided by practically anyone else is likely to undermine your sense of making a valuable contribution. The high turnover rate also impacts the possibility of achieving a meaningful sense of *community* with co-workers, and the monitoring technologies described above discourage employees from interacting and building relationships. Finally, the goods of work are distinct in that they cannot be obtained by money, which means that workers cannot be fully compensated by a higher pay. Indeed, the long shifts and (in some jurisdictions) lack of job security for FC workers is likely to leave workers too exhausted or with too little time to engage in non-work activities that provide an opportunity to pursue these goods.

6 Objections

Skeptics might agree that work of the kind we have just described stunts human capabilities by exposing workers to the bads identified above, or preventing them from pursuing the goods of work. Yet, these characteristics, they might continue, are not new or endemic to the kinds of AI-driven productivity augmentations of jobs we focus on, but rather an unavoidable consequence of the division of labor and industrial production. One of Marx's complaints about capitalism, illustrated by the classic Charlie Chaplin movie "Modern Times," was that the worker "…becomes an appendage of the machine, and it is only the most simple, most monotonous, and most easily acquired knack, that is required of him." (Marx & Engels, 1848). Furthermore, the sense of powerlessness among Amazon workers could be due to the hierarchical organization of the workplace rather than being subservient to their scanner. Yet, such hierarchies are characteristic of all firms and not only those with advanced automation technology (Anderson, 2017).

In response, we recognize that human stunting could happen even when there is no automation, simply due to the division of labor and the hierarchical workplace. AI automation nevertheless represents a shift since the technology has the potential to significantly boost human stunting practices. Whereas employers have always

been able to monitor workers, for instance, AI technology allows constant and intrusive surveillance of every employee at virtually no cost, and this is likely to significantly increase its prevalence and reach.[14] We agree, then, that the diffusion of AI technologies to workplaces that are already organized to prevent workers from pursuing the goods of work does not create entirely new risks but exacerbates existing ones. And this analysis seems to be new in the context of how AI may disrupt the future of work. In this context, the "labor enabling" scenario has almost universally been described as benign. While there have been extensive discussions on how technology could change the power relation between workers and employers (for example, in the gig economy), the possibility of productivity improvements that also stunt humans in various ways have not come under the same scrutiny.

More generally, AI technology can be used in ways that are positive and negative from the perspective of workers. But technology is not "neutral" in the sense that it never guides or nudges social dynamics in a certain direction (Gunkel & Bryson, 2014). Rather, technology can be thought of as providing us certain affordances, making some actions possible, creating opportunities and situations that make certain outcomes more likely. For example, while it is true that "guns don't kill people, people do," guns certainly make it easier to kill. For example, the widespread ownership of guns in the USA means that suicides are much more likely to be successful than in most other high-income countries (Miller et al., 2013). In other words, while a gun cannot act by itself, it makes certain acts possible and other acts more likely. The prevalence of technology can also have more subtle social effects. For example, it has been hypothesized that widespread gun ownership in the USA makes police work more dangerous, and thus makes police officers more likely to shoot suspects (Nagin, 2020).

Likewise, although worker surveillance is already commonplace in industrial production, AI technology drastically cuts the costs and increases the opportunities for employers to monitor and control their workers, allowing for the spread of these functions to new sectors of the workplace. As Gunkel and Bryson argue, technology is not a "mere tool," and certain kinds of technology facilitate certain kinds of social outcomes (Gunkel & Bryson, 2014). This does not imply technological determinism. Technology is a social process, and can be influenced by other social forces. For example, political institutions can direct how AI technology is deployed in the workplace.

A second objection, however, could admit the undesirability of a labor-enabling scenario where the productivity increase comes from human stunting rather than human augmentation. Nevertheless, these problems can be expected to disappear sooner or later, as automation progresses. Amazon workers are right to complain about being treated as robots, but this is just because Amazon is in the process of changing their FCs in order to fully replace workers with robots. With this line of reasoning, automation is not only the source of the problem but also its solution.

[14] Reports suggest that the automatic surveillance of remote white-collar workers which was implemented during the COVID pandemic is becoming more widespread (Kantor and Sundaram, 2022).

In response, it does not seem like FCs will be fully automated in the near future. Indeed, a recent study of patents owned by Amazon suggests that the company is planning for people to be around for the foreseeable future. While not all of these patents will be realized or implemented, many of them seem to be further innovations in the field of technology that would result in human stunting, including the requirement for workers to wear devices for data collection and "augmented reality" visors. For example, the company seems concerned that workers think too much while performing tasks, which can lead to "agent confusion" (i.e., workers not reacting quickly to instructions). To address this, a variety of visual and tactile aids are suggested to guide workers in real time as they pick products including lights pointed at the commodity to be retrieved, vibrations on bracelets, or arrows that are layered onto a worker's visual field through augmented reality visors (Delfanti & Frey, 2021, p. 662).

These tools would not only convey information to workers, however, but also monitor workers' bodies in unprecedented detail. The visors would collect information about what workers see and how they move their field of view in response to stimuli. Bracelets would allow the inventory management system to know where workers' arms are at any time, and give appropriate feedback if those appendages are not moving in the right way.

Two additional points should be made. The first is that the same reason that makes Amazon a suitable illustrative case—its economic ability to research or acquire pioneering automation technology—suggests that many smaller companies are likely to adopt less sophisticated automation technologies. It could be the case that while workers are fully replaced in large corporations, other sectors of the economy will be only partially automated and begin employing people in human-stunting centaur teams. Second, even if future technological innovations will replace humans in the long run, it is a cold comfort to workers whose capabilities are stunted in the short- and medium term. And it would entail that we no longer see a labor enabling scenario but rather a labor replacing one, with all the downsides usually ascribed to it.

A *third objection* says that it would require unacceptable paternalism to prevent workers and employers from entering particular kinds of work contracts, even when those contracts involve some degree of stunting. On this view, employees sell their work to employers and if the contract specifies that the employers may surveil workers or control them through AI-assisted augmented reality, then it is legitimate. Indeed, even if some work is generally regarded as unattractive, this is insufficient motivation for preventing adults from consenting to perform it in exchange for a salary. Not only libertarians but also most liberal egalitarian theories of justice argue that society should not be arranged after assumptions about what is objectively good or meaningful to people. State interventions to promote more meaningful work could hence be seen as perfectionist: A preference for such work is just one out of many preferences people who seek employment have, and workers who do not mind performing non-meaningful work would be treated

unfairly if such an intervention was enacted (Arneson, 1987; see also Kymlicka, 2002, chap. 5; and Parr, 2022).

In response, note that we have not in fact argued in favor of any potentially paternalistic policy intervention but rather aimed to defend the much weaker thesis that human stunting is a serious moral concern which ought to receive more attention in the academic, policy, and public discussion about the future of work. The minimalist interpretation of our argument, then, is that when surveillance or other AI-assisted technologies introduce the kind of bads we have described, then this can be considered to be human stunting. Whether human stunting can be permissible all-things considered is another discussion. Our analysis draws on Gheaus & Herzog's account of the goods of work, in order to avoid making unnecessarily strong assumptions about what people value (Gheaus & Herzog, 2016, pp. 72; 82 f). We have hence not argued that the goods of work are the only goods that matter, or that reasonable people will always agree on how to trade off one against the other, and so forth. It is possible, for instance, that individuals will be willing to take up stunting jobs if they pay more, or that labor interests as aggregated in trade unions coalesce around embracing automating technologies of both the augmenting and the stunting kind. Furthermore, we do not suggest that the problems we have identified show that governments should step in to regulate or ban technological development. On the other hand, the claim that state intervention would always be unacceptably paternalist could perhaps be questioned by considering the possibility that individuals have formed their preferences in the wrong kind of environment. The fact that many workers are currently unable to pursue the goods of work suggests, for instance, that they might mistakenly ascribe it lower value relative to, say, a higher income (cf. Arneson, 2009). Similarly, if workers face a collective action problem where each party is forced to accept human stunting automation in order to avoid unemployment, then state intervention could perhaps help create a new equilibrium which everyone prefers, all things considered.

This response could provoke a final objection, however, which says that our focus is misguided. By accepting the fact that modern societies are organized in ways that make it practically impossible to pursue the goods of work outside of paid labor, we fail to recognize the possibility of seeking better alternatives. If AI really is as transformative as we have assumed, the proper goal should be to develop AI technologies that reduce the amount of work that humans need to do, and allow people to pursue a flourishing life outside of work. Additionally, given the ecological costs of increased production and consumption, it might be essential for societies to abandon the model of pursuing flourishing lives in paid work or through consumer goods. While the first three objections could be said to accuse our analysis of being radically utopian, this final objection instead says that we are myopic and not radical enough.[15]

In response, we agree that ours is not the most radically utopian analysis of the future of work. The main reason is that we have focused on a short- to medium time frame, and we deem it highly unlikely that technology will make human labor

[15] Thanks to an anonymous reviewer for pressing us on this point.

redundant in the foreseeable future. First, despite the rapid progress in AI software, it is still incredibly difficult to automate many of the tasks performed in chaotic and non-standardized environments that humans routinely perform with little effort. Consider, for instance, the extent to which expectations about autonomous vehicles have been frustrated in the last decade. It would be too quick to assume that significant fractions of the labor that is currently necessary to reproduce society and provide basic goods and services can simply be eliminated (Cf. Gourevitch, 2022). Second, even when there is technology that allows tasks to be automated, it does not mean that it will necessarily happen, or that automation will spread throughout the economy at an accelerating or even steady pace (*contra* Altman, 2021). As we discussed above, technology may allow for certain scenarios to happen, but technology adoption is complex and depends as much on economic and political incentives and institutions as it does on technological capacities. Judging from history, it is highly unlikely that increased productivity will lead to drastic reductions of working time. And even if this happens over the long run, many workers may suffer stunting effects in the short- to medium term (cf. Frey, 2019).

So, although utopian analyses of what will happen over the long term are no doubt important and interesting, we nevertheless assume that the largely overlooked observations we have made about the short- to medium term are crucial and pressing issues. Ideally, however, the broad project of analyzing the future of work should proceed at both levels of analysis. Even though we are ultimately agnostic about whether certain goods of work would be lost in a world without wage labor, our argument is fully consistent with more radical proposals to transform society.

7 Conclusions

This paper has offered a political-philosophical analysis of the augmenting-stunting distinction, which should not be overlooked in the debate on the future of work. Our analysis shows why we should be more cautious about this future than some of those who hope for a labor-enabling scenario suggest. Although teaming up with autonomous systems could enable humans to have more stimulating work lives, it is also concomitant with a substantial risk that AI technologies will stunt human capabilities in ways detrimental to large groups of workers. The pioneering ways in which Amazon fulfillment centers have achieved high productivity using human–machine interaction suggest where the rest of the economy might be heading. It is of course possible that coming waves of AI-driven automation primarily will enable human augmentation. It seems plausible, for instance, that automation will affect different kinds of labor markets differently, depending on the skill level of workers, the institutions in place, and the presence or absence of policies regarding reskilling. Our analysis complements and contributes to the emerging literature on the impact of increasingly complex technology on the quality of or access to meaningful work. Having developed and illustrated the concepts of centaurs and human-augmenting and -stunting technology, we join others in calling for additional empirical work (Smids et al, 2020), to gain insight into what actually happens in workplaces when humans start cooperating with AI systems. In addition, in light of ours and similar

theoretical accounts of what is at stake, we believe more attention needs to be paid to the role of political institutions in promoting a particular scenario.

Acknowledgements We are grateful for comments and feedback from the editor and two anonymous referees for Philosophy & Technology. The paper has also benefited from comments at the conference of Arbetarrörelsens forskarnätverk in Bommersvik in November 2021, from the audience at the Paris Institute for Advanced Study in April 2022, as well as Jelena Belic and other participants at the workshop on AI and work at TU Delft in September 2022.

Author Contribution Both authors contributed equally to all parts of the paper.

Funding Open access funding provided by Stockholm University. Furendal's work was supported by Marianne and Marcus Wallenberg Foundation, grant number MMW 2020.0044.

Data Availability Not applicable.

Declarations

Ethics Approval and Consent to Participate Not applicable.

Consent for Publication Not applicable.

Competing Interests The authors declare no competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

References

Abril, D., Harwell, D. (2021). Keystroke tracking, screenshots, and facial recognition: The boss may be watching long after the pandemic ends. Washington Post. https://www.washingtonpost.com/technology/2021/09/24/remote-work-from-home-surveillance/. Accessed 24 Apr 2023)

Acemoglu, D., & Restrepo, P. (2020). The wrong kind of AI? Artificial intelligence and the future of labour demand. *Cambridge Journal of Regions, Economy and Society, 13*, 25–35. https://doi.org/10.1093/cjres/rsz022

Acemoglu, D., Restrepo, P. (2018). Artificial intelligence, automation and work. Working Paper 24196, National Bureau of Economic Research, Cambridge, MA. https://doi.org/10.3386/w24196

Agrawal, A., Gans, J., & Goldfarb, A. (2018). *Prediction machines: The simple economics of artificial intelligence*. Harvard Business Press.

Altman, S.A. (2021). Moore's Law for Everything [WWW Document]. https://moores.samaltman.com/. Accessed 22 Feb 2023

Anderson, E. (2017). *Private government: How employers rule our lives (and why we don't talk about it)*. Princeton University Press.

Arneson, R. J. (1987). Meaningful work and market socialism. *Ethics, 97*, 517–545. https://doi.org/10.1086/292864

Arneson, R. J. (2009). Meaningful work and market socialism revisited. *Analyse & Kritik, 31*(3), 139–51.

Arnold, S. (2012). The difference principle at work. *Journal of Political Philosophy, 20*, 94–118. https://doi.org/10.1111/j.1467-9760.2010.00393.x

Bankins, S., Formosa, P. (2023). The ethical implications of artificial intelligence (AI) for meaningful work. J Bus Ethics. https://doi.org/10.1007/s10551-023-05339-7

Barbaro, M. (2018). The human toll of instant delivery. The Daily, Nov. 11, 2018. https://www.nytimes.com/2018/11/26/podcasts/the-daily/warehouse-workers-instant-delivery.html. Accessed 21 Mar 2022

Bieber, F., & Moggia, J. (2021). Risk shifts in the gig economy: The normative case for an insurance scheme against the effects of precarious work*. *Journal of Political Philosophy, 29*, 281–304. https://doi.org/10.1111/jopp.12233

Bietti, E. (2020). F*rom ethics washing to ethics bashing. Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 210–219. New York, NY, USA https://dl.acm.org/doi/10.1145/3351095.3372860

Berkers, H.A., Rispens, S., Le Blanc, P. M., 2022. The role of robotization in work design: A comparative case study among logistic warehouses. *The International Journal of Human Resource Management*. https://doi.org/10.1080/09585192.2022.2043925

Braverman, H. (1998). *Labor and monopoly capital: The degradation of work in the twentieth century*. NYU Press.

Brynjolfsson, E., McAfee, A. (2014). The second machine age: work, progress, and prosperity in a time of brilliant technologies. New York. Norton & Company.

Brynjolfsson, E., & McAfee, A. (2012). *Race against the machine: How the digital revolution is accelerating innovation, driving productivity, and irreversibly transforming employment and the economy*. Lexington, Mass. Digital Frontier Press.

Cheng, E. (2016). Amazon climbs into list of top five largest US stocks by market cap. CNBC. URL https://www.cnbc.com/2016/09/23/amazon-climbs-into-list-of-top-five-largest-us-stocks-by-market-cap.html (accessed 3.21.22).

Cowen, T. (2013). *Average is over: Powering America beyond the age of the great stagnation*. Dutton.

Danaher, J. (2019). *Automation and Utopia — Human flourishing in a world without work*. Cambridge. Harvard University Press.

de Laat, P. B. (2021). Companies committed to responsible AI: From principles towards implementation and regulation? *Philosophy and Technology, 34*, 1135–1193. https://doi.org/10.1007/s13347-021-00474-3

Delfanti, A. (2019). Machinic dispossession and augmented despotism: Digital work in an Amazon warehouse. *New Media & Society, 23*, 39–55. https://doi.org/10.1177/1461444819891613

Delfanti, A., & Frey, B. (2021). Humanly extended automation or the future of work seen through Amazon patents. *Science, Technology, and Human Values, 46*, 655–682. https://doi.org/10.1177/0162243920943665

Field, A. J. (2019). Review: The technology trap by Carl Benedikt Frey. https://eh.net/book_reviews/the-technology-trap-capital-labor-and-power-in-the-age-of-automation/. Accessed 24 Apr 2023

Floridi, L. (2014). Technological unemployment, leisure occupation, and the human project. *Philosophy and Technology, 27*, 143–150. https://doi.org/10.1007/s13347-014-0166-7

Floridi, L., & Chiriatti, M. (2020). GPT-3: Its nature, scope, limits, and consequences. *Minds and Machines, 30*, 681–694. https://doi.org/10.1007/s11023-020-09548-1

Frey, C. B. (2019). *The technology trap: Capital, labor, and power in the age of automation*. Princeton University Press.

Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change, 114*, 254–280. https://doi.org/10.1016/j.techfore.2016.08.019

Furendal, M. (2019). Defining the duty to contribute: Against the market solution. *European Journal of Political Theory, 18*, 469–488. https://doi.org/10.1177/1474885117693401

Gallego, A., & Kurer, T. (2022). Automation, digitalization, and artificial intelligence in the workplace: Implications for political behavior. *Annual Review of Political Science, 25*, 463–484. https://doi.org/10.1146/annurev-polisci-051120-104535

Gheaus, A., & Herzog, L. (2016). The goods of work (other than money!). *Journal of Social Philosophy, 47*, 70–89. https://doi.org/10.1111/josp.12140
GMB Union. (2021). Amazon worker injuries spike in build up to key shopping dates. GMB. URL https://www.gmb.org.uk/news/amazon-worker-injuries-spike-build-key-shopping-dates (accessed 3.21.22).
Gomberg, P. (2016). Why distributive justice is impossible but contributive justice would work. *Science and Society, 80*, 31–55. https://doi.org/10.1521/siso.2016.80.1.31
Goos, M., Manning, A., & Salomons, A. (2009). Job polarization in Europe. *The American Economic Review, 99*, 58–63. https://doi.org/10.1257/aer.99.2.58
Gourevitch, A. (2022). Post-work socialism? *Catalyst, 6*(2), 8–48.
Greene, J., Alcantara, C. (2021). Amazon warehouse workers suffer serious injuries at higher rates than other firms. Washington Post. https://www.washingtonpost.com/technology/2021/06/01/amazon-osha-injury-rate/. Accessed 24 Apr 2023
Gunkel, D. J., & Bryson, J. (2014). Introduction to the special issue on machine morality: The machine as moral agent and patient. *Philosophy and Technology, 27*, 5–8. https://doi.org/10.1007/s13347-014-0151-1
Harari, Y. N. (2017). Reboot for the AI revolution. *Nature, 550*, 324–327. https://doi.org/10.1038/550324a
Heater, B. (2022). Amazon debuts a fully autonomous warehouse robot. TechCrunch. URL https://techcrunch.com/2022/06/22/amazon-debuts-a-fully-autonomous-warehouse-robot/ (accessed 10.7.22).
Hosseini, Z. et al. (2023). Assessing the artificially intelligent workplace: An ethical framework for evaluating experimental technologies in workplace settings. *AI and Ethics*. https://link.springer.com/10.1007/s43681-023-00265-w
Hsieh, N. (2008). Survey article: Justice in production. *Journal of Political Philosophy, 16*, 72–100.
Jamieson, D. (2015). The life and death of an Amazon warehouse temp. Huffington Post.
Jones, P. (2021). *Work without the worker: Labour in the age of platform capitalism*. Verso Books.
Kantor, J., & Sundaram, A. (2022). *The rise of the worker productivity score*. N. Y. Times.
Kasparov, G., & Greengard, M. (2017). *Deep thinking: Where machine intelligence ends and human creativity begins* (1st ed.). PublicAffairs.
Kymlicka, W. (2002). *Contemporary political philosophy: An introduction*. Oxford University Press.
Levy, S. (2014). A spreadsheet way of knowledge. WIRED. URL https://www.wired.com/2014/10/a-spreadsheet-way-of-knowledge/ (accessed 3.21.22).
Macnish, K. (2022). *Surveillance ethics*. Internet Encyclopedia of Philosophy. https://iep.utm.edu/surv-eth/
Marx, K. (1847). Wage labour and capital. Marx/Engels internet Archive. https://www.marxists.org/archive/marx/works/1847/wage-labour/. Accessed 24 Apr 23
Marx, K., & Engels, F. (1848). *Communist manifesto (Chapter 2). Marx/Engels Internet Archive*. https://www.marxists.org/archive/marx/works/1848/communist-manifesto/ch01.htm. Accessed 24 Mar 2022
Miller, M., Barber, C., White, R. A., & Azrael, D. (2013). Firearms and suicide in the United States: Is risk independent of underlying suicidal behavior? *American Journal of Epidemiology, 178*, 946–955.
Nagin, D. S. (2020). Firearm availability and fatal police shootings. *The Annals of the American Academy of Political and Social Science, 687*, 49–57. https://doi.org/10.1177/0002716219896259
PAI. (2023). AI, Labor, and the Economy. https://partnershiponai.org/program/ai-labor-and-the-economy/ (March 24, 2023).
Parr, T. (2022). In cash we trust? *Journal of Applied Philosophy*. EarlyView. https://doi.org/10.1111/japp.12583
Paus, E. (Ed.). (2018). *Confronting dystopia: The new technological revolution and the future of work*. Cornell University Press. https://doi.org/10.7591/9781501719868
Roessler, B. (2012). Meaningful work: Arguments from autonomy*. *Journal of Political Philosophy, 20*, 71–93. https://doi.org/10.1111/j.1467-9760.2011.00408.x
Sainato, M. (2020). *'I'm not a robot': Amazon workers condemn unsafe, grueling conditions at warehouse*. The Guardian.
Sainato, M. (2019). *"Go back to work": outcry over deaths on Amazon's warehouse floor*. The Guardian.
Santoni de Sio, F., Almeida, T., van den Hoven, J. (2021). The future of work: freedom, justice and capital in the age of artificial intelligence. *Critical Review of International Social Political Philosophy 0*, 1–25. https://doi.org/10.1080/13698230.2021.2008204

Scanlon, T. (1998). *What we owe to each other*. Harvard University Press.
Schoeman, F. (1984). Privacy: Philosophical dimensions. *American Philosophical Quarterly, 21*, 199–213.
Schwartz, A. (1982). Meaningful work. *Ethics, 92*, 634–646. https://doi.org/10.1086/292380
Smids, J., Nyholm, S., & Berkers, H. (2020). Robots in the workplace: A threat to—or opportunity for—meaningful work? *Philosophy & Technology, 33*, 503–522.
Streitfeld, D., & Kantor, J. (2015). *Jeff Bezos and Amazon employees join debate over its culture*. New York Times.
Stuart, F., 2021. My life as an Amazon warehouse worker. Jacobin.
Veltman, A. (2016). *Meaningful work*. Oxford University Press.
Vincent, J. (2021). *Amazon delivery drivers have to consent to AI surveillance in their vans or lose their jobs*. The Verge.
Vredenburgh, K. (2022). Freedom at work: Understanding, alienation, and the AI-driven workplace. *Canadian Journal of Philosophy, 52*, 78–92. https://doi.org/10.1017/can.2021.39
Willcocks, L. (2020). Robo-apocalypse cancelled? Reframing the automation and future of work debate. *Journal of Information Technology, 35*, 286–302. https://doi.org/10.1177/0268396220925830
Wolff, J., Leopold, D. (2021). Karl Marx, *The Stanford Encyclopedia of Philosophy* (Spring 2021 Edition), Edward N. Zalta (ed.), https://plato.stanford.edu/archives/spr2021/entries/marx
Yeoman, R. (2014). Conceptualising meaningful work as a fundamental human need. *Journal of Business Ethics, 125*, 235–251. https://doi.org/10.1007/s10551-013-1894-9
Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power* (1st ed.). PublicAffairs.
Zuboff, S. (2015). Big other: Surveillance capitalism and the prospects of an information civilization. *Journal of Information Technology, 30*, 75–89. https://doi.org/10.1057/jit.2015.5

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.