The Ethics of AI Automation: The Importance of Treating Like Cases Alike

1. Introduction

Recent advances in artificial intelligence (AI) have led many to speculate that a significant number of jobs may be automated soon. Automation refers to the substitution of human labor for machine labor. In what is perhaps the most famous prediction, Carl B. Frey and Michael A. Osborne (2017: 268) conclude that "around 47% of total US employment is in the high risk category. We refer to these as jobs at risk—i.e. jobs we expect could be automated relatively soon, perhaps over the next decade or two." Other predictions are similarly startling. For instance, Goldman Sachs estimates that 300 million jobs worldwide could be automated due to AI (Kelly 2023).

Such predictions have led to a surge in research that examines the ethics of forthcoming AI automation. Some of this work examines the moral desirability of a world in which machine labor is so powerful it completely substitutes all human labor (e.g., Danaher 2019; Bostrom 2024; Knell and Rüther 2024). While a fascinating possibility, many believe this is unlikely, for a few reasons. First, historically machines not only substituted (and hence reduced the demand for) human labor, but they complemented (and thus increased the demand for) it as well (Autor et al. 2023: 13). Second, AI may never acquire all human capabilities. Some have recently argued that "AI systems are not capable of categorical rationality, don't possess (moral) concepts and understanding, don't have affective experience and moral intuition and cannot cultivate virtues and develop practical wisdom" (Kyriacou 2024: 10).

Third, *even if* AI robots do develop capabilities that surpass humans at all tasks, it still does not follow that humans will be without work. The law of comparative advantage tells us that it is in the interest of Britain to trade with Portugal even if Britain can produce everything more efficiently; British workers' labor is scarce, so it's smart to use it only for the most valuable tasks and import other goods from Portugal (Ricardo 2004: ch. 7). Similarly, even if AI robots are better than humans at all tasks, it may still make sense for humans to work, so long as AI robots are sufficiently expensive to use (Kogelmann 2024). For many reasons, we doubt the strong view that AI will fully automate all human jobs, creating a post-work utopia (or perhaps dystopia).

But even if the strong view is false and forthcoming AI automation does not render work obsolete, it may nonetheless lead to significant short-term unemployment. Moreover, AI automation will undoubtedly change the character of work, and could have other impacts on politics, culture, and society broadly construed. Reflecting on these more realistic possibilities, many researchers have found cause for concern. For a variety of reasons, forthcoming AI automation, they argue, is morally problematic.[1] We push back on these arguments in the current paper. We do not offer a positive argument for why forthcoming AI automation is good. Our goal is to show those voicing moral objections to the rise of AI automation are incorrect in their assessment. Though the arguments we canvass are diverse, they all fall prey to the same flaw: they fail to treat like cases alike.

To demonstrate this, we show that criticisms of AI automation apply to past forms of automation that most think were overwhelmingly positive. For of course, AI automation is not all that new; it is the continuation of a long and slow process of using technology to increase

[1] Not everyone agrees with this assessment. See, for instance, Spencer (2024).

productivity by both substituting and complementing human labor. Ever since the first human harnessed the first plow to the first oxen, humans have been automating their labor. Much of this has been desirable. To treat like cases alike, critics of forthcoming AI automation must be willing to extend their proposed policies to all forms of automation that share AI's alleged moral shortcomings. Most would hesitate to do this, however.

Here is how the paper proceeds. In the next section we overview the existing literature that criticizes forthcoming AI automation from the moral point of view (§2). After this we argue that these criticisms also apply to the automation of agriculture, a form of automation many deem overwhelmingly positive (§3). Inspired by this symmetry, we then present and defend a novel ethical principle, which says that if a policy seems appropriate for addressing challenges related to AI automation, then it must also seem appropriate if applied to the automation of agriculture in the past (§4). Like cases must be treated alike. Most policies proposed by the new AI skeptics fail to clear this hurdle.

2. Moral Concerns with Forthcoming AI Automation

What should we think about forthcoming AI automation from the moral point of view? We've conducted what we believe is an exhaustive reading of all academic papers exploring the ethics of AI automation released in the last decade. Reviewing these papers, we believe ethical concerns about AI automation can be grouped into three different camps: concerns about how AI automation will lower the *quality* of work, concerns about how it will disrupt the *stability* of work, and concerns about how it will cause broader social, political, and cultural *ills*. We explicate all three camps in this section.

2.1 *The Quality of Work.*

The first camp worries AI automation will lower the quality of work. Some jobs are good, and some are bad. Many people want to be entrepreneurs, software developers, tenured professors, artists, and project managers, but few want to spend their workday on an assembly line performing the same task repeatedly. That there are good and bad jobs is uncontroversial, but what *makes* a job or good bad is less clear. At the most abstract level, we can say good jobs have certain features that make them good (Gheaus and Herzog 2016: 70; Veltman 2016: 117; Celentano 2024: 33). We call these *good-making features*. Candidate good-making features include (but are not limited to) the opportunity to achieve excellence, the opportunity to make an important social contribution, a good fit with other aspects of a worker's life, the opportunity to develop virtue, self-direction, the achievement of dignity, and more (Gheaus and Herzog 2016: 74-79; Veltman 2016: 117; Celentano 2024: 33). Whatever the list of good-making features ends up being (we take no stand on its exact content), we can say that job j_1 is of a higher quality than job j_2 if and only if j_1 has more good-making features than j_2. To say AI automation will result in lower quality jobs is thus to say it will produce jobs with less good-making features when compared to existing jobs.

For instance, John Danaher and Sven Nyholm (2021) argue that *achievement* is a good-making feature of work that AI automation is liable to eliminate or at the very least make more difficult to realize. Several components are necessary for achievement: you must produce something valuable, you must make a significant causal contribution to its production, you must exert effort, all of which must be voluntary (Danaher and Nyholm 2021: 231). AI automation upsets some of these components. For instance, AI automation reduces the causal impact humans

have in the production process (Danaher and Nyholm 2021: 234). Where individuals would once themselves write up a report, they now do so with a heavy lift from ChatGPT. This undermines achievement. As another example, AI automation reduces required effort: "automation, almost by necessity, reduces the cost of the human commitment to producing a workplace output" (Danaher and Nyholm 2021: 233). Because an agent can now finish her work in four hours instead of eight with the help of AI, she achieves less.

Another representative of the first camp is Denise Celentano (2019; 2024). She highlights several good-making features of work, but we want to focus on two: *self-direction* and *self-development* (Celentano 2024: 33). Self-direction "refers to having room for conceiving the tasks to perform besides merely executing them" (Celentano 2024: 33). This good-making feature of work is present when individuals not only perform various tasks, but also decide which tasks to perform. A second good-making feature of work is self-development. Because "monotonous and repetitive work with no opportunities for learning leaves little room for the development of one's capacities," a good job will afford agents opportunities to develop their skills and talents (Celentano 2024: 34).

Celentano worries forthcoming AI automation will result in jobs that do not have these good-making features. Algorithmic management can thwart self-direction, for instance (Celentano 2024: 37). Consider the case of a delivery driver whose entire job is governed by an algorithm. In the name of efficiency, the algorithm dictates the order of deliveries, expected arrival times, suggested routes between houses, suggested driving speeds, when and where he can stop for breaks, and so on. This leaves the driver little control to make decisions based on his knowledge, experiences, and preferences. Without the algorithm, the delivery driver has far more opportunities for self-direction.

Ghost work is “invisible human labor operating behind the scenes of AI” (Celentano 2024: 36). This includes things like “spending long hours labeling images … and cleansing the internet of inappropriate content” (Celentano 2024: 36). Increased use of AI will lead to more ghost work; this is bad for self-development. Ghost workers engage in narrow, low-skill tasks like data labeling, content moderation, or image tagging that offer little room for development of their skills and talents. The isolated nature of these tasks also prevents workers from building valuable professional networks or gaining comprehensive industry knowledge. Ghost workers find themselves trapped in roles that offer few opportunities for skill expansion or career advancement. Their self-development is hindered.

2.2 *The Stability of Work*.

The second concern with forthcoming AI automation is that it will create job instability. Jelena Belic (2023) voices this concern most forcefully, but she is not alone (e.g., Celentano 2024: 33; Kogelmann 2024). To better understand it, we must introduce some terminology. Many philosophers uphold *autonomy* as an important moral ideal (e.g., Mill 1978: 56; Raz 1986: 369; Wall 1998: 128). Autonomy is the “ability to develop, revise and pursue one’s plans and commitments” (Belic 2023: 4). The autonomous agent sets goals; for instance, to become a lawyer, a mother, and a marathon runner. She develops a plan to achieve these goals. For her legal career, she recognizes the need to attend law school. To start a family, she prioritizes finding a compatible partner. To complete a marathon, she commits to a rigorous training regimen. With her roadmap in place, the autonomous agent pursues these actions to achieve her goals (Kogelmann 2021: 100).

To facilitate autonomy, several background conditions must be in place. For instance, agents must have certain cognitive and emotional capacities (Raz 1986: 327; Wall 1998: 139). They must have access to a sufficient variety of choice options (Raz 1986: 375; Wall 1998: 141). Most relevant for our purposes, autonomy demands a certain degree of stability (Raz 1986: 411). As one defender of the ideal of autonomy puts it: "In order for people to have access to a sufficiently wide range of options they must not only have access to options that would allow them to develop their capacities and talents, but also to particular options that have become indispensable to this development" (Wall 1998: 143). It's easy to see why stability is important. If you are pursuing a certain career option as a part of your life plan, and all of the sudden it is eliminated, then your plan has been irretrievably upset. Unable to execute your plan, an autonomous life is fleeting.

AI automation can induce instability and thus thwart life plans in many ways. First, AI automation can render goals *obsolete* while agents pursue them. This is likely to happen if individuals set very specific goals concerning the kinds of careers they want to have. A prime example is in the field of medical imaging, where machine learning algorithms have demonstrated superior diagnostic capabilities compared to human practitioners. This development has led to speculation about the future relevance of human radiologists (Guilford-Blake 2020). AI diagnostic tools could significantly impact the career plans of current radiology students and recent graduates, effectively nullifying their long-term professional aspirations as they pursue them.

Even those without very specific career goals can see their plans upset by forthcoming AI automation. This brings us to the second way AI automation may upset life plans: by making chosen means *ineffective* for achieving desired ends. Consider an individual with the general aim

of attaining a comfortable lifestyle to support a large family and enjoy certain luxuries. This person might choose a legal career as a means to this end. However, the advent of AI tools like ChatGPT may significantly alter the legal landscape (Weiss 2023). These technologies have the potential to reduce the demand for human lawyers, leading to decreased wages. Consequently, an individual who went to law school as a means to achieve her end of attaining a certain lifestyle may find her plan undermined.

The third way AI automation might disrupt life plans is by inducing *decision paralysis*. As AI technologies rapidly evolve and reshape the labor market, individuals may become hesitant to form any plans at all. This reluctance stems from heightened uncertainty about which skills and professions will remain viable options in the future. This phenomenon bears similarities to arguments in development economics, where insecure property rights are thought to discourage asset owners from making investments due to uncertainty about future returns (Acemoglu and Robinson 2012: 75). In the context of AI automation, the constant flux in occupational landscapes may deter people from making decisive career choices. For example, a high school student passionate about graphic design might hesitate to pursue a degree in this field, fearing that AI image generation tools could significantly reduce job prospects for human designers by the time they graduate. This uncertainty may lead to indecision, or the formation of no plan at all.

2.3 *Broader Social, Political, and Cultural Ills*.

The third and final set of concerns about forthcoming AI automation is that it will lead to broader social, political, and cultural ills. For example, Jake Burley and Nir Eisikovits (2023) worry that AI automation may result in extremist political movements grounded in feelings of

replacement. A group experiences feelings of replacement when they lose valued status that they believe they deserve, see this status transferred to a group they deem undeserving, and yearn to reclaim their former position (Burley and Eisikovits 2023: 1363). Feelings of replacement "are predictive of political action in the best case and political violence in the worst case" (Burley and Eisikovits 2023: 1363). White nationalist movements, for example, are driven by feelings of replacement.

Burley and Eisikovits worry AI automation may lead to the kinds of feelings of replacement that inspire extremist political movements. Suppose within the span of a few years all truckers lose their jobs to AI. Truckers lose status to another group (robots) they deem unworthy. They will surely desire to regain their lost status. The truckers may also believe their status was something owed to them: they went through arduous training to receive a CDL license and put in years of service, so they are entitled to the job. Finally, the status that was lost will be remembered fondly among truckers. Truck-driving is lucrative; it is unlikely many will be able to make as good of money elsewhere. With all the conditions in place, feelings of replacement are likely. The "results may well become socially and politically explosive" (Burley and Eisikovits 2023: 1368).

Let us summarize this section. Based on our reading of the literature, three broad concerns have been raised about forthcoming AI automation. The first worries AI automation will produce jobs with less good-making features when compared to existing jobs. The second worries AI automation will upset life plans and thus inhibit autonomy. The third worries AI automation will lead to broader social, political, and cultural ills outside the workplace, such as political extremism.

3. The Automation of Agriculture

In the last section we highlighted three concerns with forthcoming AI automation: it may lower the quality of work, it may disrupt the stability of work, and it may cause broader social, political, and cultural ills. We dispute none of these concerns. Our central claim is that these drawbacks are not sufficient to demonstrate that forthcoming AI automation should be stopped, or its implementation significantly regulated. The current section takes the first step toward establishing this conclusion. We argue other forms of past automation were subject to the *exact same* criticisms and concerns. More specifically, automation in agriculture—which includes plows, mechanized tractors, combine harvesters, seeders, planters, milking machines, produce pickers, hydroponic systems, spraying pesticides via aircraft, and more—caused the exact same problems many worry forthcoming AI automation will cause. The normative significance of this is left for the next section. This section simply highlights the facts.

Automation in agriculture has been of seismic importance. The great economic pessimist Thomas Malthus (2015) argued that sustained improvement in living standards was impossible, because population growth always outstrips humankind's ability to produce food. Though innovation might allow farmers to extract more food per acre of land, this ultimately does not matter because population growth will eventually eat up (literally) these productivity gains. This is known as the *Malthusian trap*. Malthus was right until he was wrong (Clark 2007: 193). For most of human history, increased agricultural productivity was offset by population growth; the Malthusian trap was real. But during the Industrial Revolution, something happened: productivity growth in agriculture outstripped population growth. Population kept growing but it did not matter, because the ability to produce food grew *even faster*. Developing countries escaped the Malthusian trap. There are far more people today than there were when Malthus

wrote *An Essay on the Principle of Population*. Contrary to his prediction, their standards of living are beyond his imagination.

Most importantly, famines (which were once common) have almost entirely disappeared. When famines now do exist, the issue, as Amartya Sen (1983) shows, is almost never a lack of food, but rather how food is *distributed*. The disappearance of famines is due to the fact that world agricultural production from 1800-2000 increased more quickly (by tenfold) than world population (six- to seven-fold) (Federico 2005: 1). Yet amazingly, the labor it takes to produce this food has dropped dramatically. In traditional agrarian societies, agriculture employed more than 75 percent of the workforce and, "as late as 1950, about two-thirds throughout the world" (Federico 2005: 1). But by 2005, only 2.5 percent of the total workforce in advanced countries worked in agriculture (Federico 2005: 1).

This rapid increase in agricultural productivity is owed to many factors (e.g., better fertilizers, genetically modified crops, etc.), but the most important factor is technology that automated farm work (Federico 2005: ch. 6). Indeed, the mechanized tractor *alone* is estimated to have resulted in social savings equivalent to eight percent of GNP (Steckel and White 2012: 2). Given all this, it is hard not to look at automation in agriculture with anything but immense gratitude. Our lives of prosperity and abundance are unimaginable without John Deere. And yet, automation in agriculture was not all good. We can think of three downsides: agriculture automation resulted in lower quality jobs, upset life plans, and it catalyzed broader social, political, and cultural ills.

3.1 *The Quality of Work.*

Let's begin with the claim that automation in agriculture led to jobs with less good-making features than the prior status quo. One good-making feature of work, according to Celentano, is self-direction. Mechanized agriculture reduced self-direction on the farm in a few different ways. First, self-direction is easy when you're the boss, not so much when you're the worker. Mechanized farm tools resulted in less bosses and more workers, because the tools created larger farms (Schmidtz and Moss 2015). Optimal farm size increases as technology enhances productivity (Foster and Rosenzweig 2022: 641). Economies of scale explain why. Tools that automate farm labor have high fixed costs but facilitate tremendous amounts of work; the tools are thus effectively cheaper the more land that can be farmed with them. For example, in 1964 before the widespread adoption of mechanized tomato harvesters in California there were 1,072 farmers growing tomatoes on farms that averaged 132 acres in size. After adoption of the technology in 1975 there were 845 farmers growing tomatoes on farms that averaged 354 acres in size (Schmitz and Moss 2015: 283n6). If being your own boss is the pinnacle of self-direction, then 227 California tomato farmers lost this opportunity due to automation in the span of just ten years.

Second, agriculture mechanization not only reduced the number of bosses, but it also limited their choices, further constraining self-direction. Tractors and other mechanized farm tools were expensive. Most farmers lacked the capital to buy one outright, so they were financed through debt. To service the debt, farmers were forced to pursue only the most lucrative farming opportunities (Ankli 1980; Ankli and Olmstead 1981). This meant producing certain types of crops at large enough scales. By implication, many kinds of farming were no longer options. Third, another way the adoption of the mechanized tractor reduced self-direction was by reducing the domain of nuanced decision-making. Many farmers complained that tractors

"lacked flexibility" because the horsepower in the engine was fixed; by contrast, horses allowed farmers to "adjust the scale of horsepower to suit the task that was being performed" (Juma 2016: 130). Horses were modular; tractors were not. The room for decision-making by farmers, thereby, shrank with the mechanized tractor.

Mechanized agriculture plausibly reduced other good-making features of work on the farm. Danaher and Nyholm, recall, emphasize achievement. According to them, to achieve you must produce something valuable, you must make a significant causal contribution to its production, you must exert effort, all of which must be voluntary. Mechanized agriculture clearly undermines some of these conditions. When you pick fruit by hand you exert effort and make a significant causal contribution to the final product. If you merely attend to a mechanized picker, you do less of each, resulting in less achievement.

Mechanized agriculture not only reduced good-making features of work on the farm, but off the farm as well. As we shall soon see in detail, tools like the tractor displaced many workers. Automation on the farm meant less need for farmhands. Where did they go? Many went to work in manufacturing and factories. In 1870, 45.9 percent of the workforce in the U.S.A worked in agriculture; by 1940, only 17.4 percent did. By contrast, during that same period blue collar work—which includes craft workers, operatives, and laborers—rose from 33.5 percent to 39 percent of the workforce (Gordon 2016: 255).

Was a factory job better than a farm job? That depends. Working in a factory was physically easier; workers were also shielded from the elements. Yet, factory work was repetitive (Gordon 2016: 257). Thus, "repetitive jobs" rose from 28.9 percent to 53.7 percent from the 1870 to 1940 period (Gordon 2016: 255). Repetitive jobs are bad from the perspective of self-development, which, recall, is a good-making feature of work according to Celentano. Karl Marx

(1990: 386) explains why: if factory workers were "obliged to toil the year round … in a repetition of the same work, might it not blunt their ingenuity, and render them stupid instead of alert and dexterous; and might not our workmen lose their reputation instead of maintaining by such external slavery?"[2] Because it was so repetitive, factory work (again, where many displaced farm workers went) hindered self-development.

3.2 *The Stability of Work*.

Let's move on to the second criticism of agriculture automation: it resulted in job instability and upset plans. Alan L. Olmstead and Paul W. Rhode (1994) demonstrate the extent of this instability by looking at Census data. They provide evidence both that automation was occurring and that it significantly impacted employment categories. To illustrate that automation was occurring, they compare the number of horses and tractors on farms. In 1910, there were 23,934,000 horses and 1,000 tractors on American farms. But by 1960, there were 3,089,000 horses and 4,685,000 tractors (Olmstead and Rhode 1994: Table 1). How this impacted employment is easy to see.

In 1900 there were 99,061 people employed as stable owners and laborers and there were zero garage owners and laborers. By 1930, only 10,354 stable owners and laborers remained, while the number of garage owners and laborers grew to 136,658 (Olmstead and Rhode 1994: Table 2). Looking at the impact of mechanized agriculture on employment through a different lens, in 1870, 45.9 percent of the workforce worked on farms; by 2009, it was only 1.1 percent (Gordon 2016: 255). Farmers, farmhands, stable owners, horse breeders, and more all had to

[2] These are not Marx's original words; here he approvingly cites (as he often does throughout *Capital*) economist Malachy Postlethwayt to help make his point.

adopt completely new careers as a result of the tractor. Automated agriculture resulted in many upset life plans.

3.3 *Broader Social, Political, and Cultural Ills*.

Finally, mechanized agriculture led to significant social, political, and cultural ills. New technologies are almost always resisted because they produce winners and losers (Mokyr 2002: 232; Taylor 2016: ch. 7). Automation in agriculture was no exception. This resistance manifested in at least three different ways. First, automation led to protest and strike activity among farm laborers. The United Farm Workers union cited farm mechanization and its impact on employment and wages as a primary concern (Majka 1981: 543). Iconic United Farm Workers leader César Chávez said "They thought they could go ahead and build machines and displace jobs. There's no way they can get away with it. There has to be a parallel program for works" (Valdés 1994: 221).

Second, automation led to intense lobbying. The Horse Association of America (HAA) was established in 1919 (Olmstead and Rhode 1994: 39; Juma 2016: 127). Its official mission was to "champion the cause of livestock dealers, saddle manufacturers, farriers, wagon and carriage makers, hay and grain dealers, teamsters, farmers, breeders, and other business interests that had a financial or emotional interest in horses and mules" (Juma 2016: 128). It ran a sweeping campaign that tried to convince legislators and farmers to stick with horses. Third, automation led to court battles. For example, the University of California system was a leader in innovating new ways to automate farm tasks. Lawsuits were filed on behalf of farmers, claiming that the university's research "displaced farmworkers, eliminated small farms, harmed

consumers, impaired the quality of rural life, and impeded collective bargaining" (Juma 2016: 135). Farm automation was politically contentious.

There were far subtler and more insidious social and cultural ills that resulted from automation in agriculture. Horse- and oxen-drawn plows were a form of automation, where animals rather than machines substituted human labor. There is overwhelming evidence that the adoption of the plow resulted in inegalitarian gender norms (Alesina et al. 2013; Alesina et al. 2021; Hopster et al. 2022). This is because working a plow in conjunction with horse or oxen requires significant upper body strength, so it was only men who could effectively use them. Women were forced out of the field and into the home. This division of gendered labor "generated norms about the appropriate role of women in society. Societies characterized by plough agriculture, and the resulting gender-based division of labor, developed the belief that the natural place for women is within the home" (Alesina et al. 2013: 471).

Remarkably, the negative impact the plow had on gender norms persists *even today*. Researchers find that "individuals, ethnicities, and countries whose ancestors engaged in plough agriculture have beliefs that exhibit greater gender inequality today and have less female participation in non-domestic activities, such as market employment, firm ownership, and politics" (Alesina et al. 2013: 527). Moreover, "domestic violence is higher in societies traditionally based on plough agriculture, where women participated less in agriculture production" (Alesina et al. 2021: 70-71). This clearly qualifies as a negative social, political, and cultural ill that resulted from agriculture automation.

4. Treating Like Cases Alike

In the last section we showed that the *exact same* criticisms leveled against forthcoming AI automation applied to the past automation of agriculture. And yet, most people view automation in agriculture as overwhelmingly positive. The lives we live today are unimaginable without it. It is literally the reason humanity escaped the Malthusian trap. In this section we explain the moral significance of these facts.

4.1 *The LCA Principle.*

Skeptics of forthcoming AI automation cite three concerns: it will lower the quality of work, it will disrupt the stability of work, and it will cause broader social, political, and cultural ills. In the prior section we demonstrated these same concerns applied to the automation of agriculture, such as the mechanized tractor. This symmetry leads to the following moral principle:

> *Like Cases Alike (LCA)*: Policies proposed to mitigate the downsides of AI automation should be evaluated against their hypothetical application to past technological transitions, such as automated agriculture.[3]

To put it colloquially, LCA says: if you wouldn't use the policy at the dawn of the automated agriculture age, then you shouldn't use it at the dawn of the AI automation age. The rationale for LCA lies in the basic demand for logical consistency.[4] The push for policy interventions in

[3] One may notice a structural similarity with the LCA principle and the meta-ethical principle of moral supervenience. Thanks to an anonymous reviewer for flagging this similarity. For more on supervenience in ethics, see McPherson (2022).

[4] To anticipate potential confusion, this principle is one of *logical* consistency, not one concerning symmetry in *belief formation.* Hence, it is distinct from David Bloor's (1991: 7-9) "symmetry principle" that maintains that all beliefs, be they true or false, should be symmetrically explained by causes of one kind. Thanks to an anonymous reviewer for suggesting this clarification.

response to forthcoming AI automation stems from three worries. These concerns mirror those present at the onset of mechanized farming. To maintain consistency, one must acknowledge that if these concerns are compelling enough to warrant a specific policy for AI automation, then they should carry equal weight when considering the same policy in the historical context of agricultural mechanization. Like cases must be treated alike.

One might object to LCA by arguing that the two cases are not, in fact, alike: though *concerns* about AI and agriculture automation are similar, their *benefits* differ vastly. As we noted in the prior section, the benefits of automating agriculture were extraordinary. Because the benefits of mechanized farming were so great, the concerns about its downsides were worth ignoring. The benefits of AI automation, so the objection goes, will not be as significant. The concerns about its downsides are thus worth taking seriously. Because the expected benefits of automation are different between the two cases, one can, without pain of inconsistency, believe a policy should be adopted to mitigate the downsides of AI automation but not agriculture automation. This objection does not deny that like cases should be treated alike; it asserts the cases are not alike.

This objection misunderstands the historical case of farm automation. The benefits of automating agriculture were never certain. Early tractors were not very good because they could not reliably navigate all terrains. It was said that "a mule is the only fool proof tractor ever built" and, in the early days of the technology, that was true (Juma 2016: 127). When asking whether a policy should have been applied at the dawn of the farm mechanization age, we must keep in mind that the benefits of automating agriculture were uncertain. And yet, the promise of

potentially miraculous results was, in our view, sufficient reason to ignore the potential downsides and plow forward.

The same circumstances apply to AI automation. We simply do not know the magnitude of its benefits. Some are very optimistic. Google CEO Sundar Pichai says AI will be a "more profound" innovation than humanity's harnessing of fire or electricity (Prakash 2023). OpenAI CEO Sam Altman (2021) believes it will usher in an age of material abundance. Some have said that AI will usher in the next Industrial Revolution (Chiang 2023). If these predictions turn out to be even half true, then AI's benefits *will* rival the benefits of mechanized agriculture. So, it's not fair to say the two cases are unalike because their benefits differ. In fact, the two cases are exactly alike: in the early days of the respective technologies, the benefits are uncertain but potentially extraordinary.

4.2 *Evaluating Policy with LCA.*

With LCA established, we can now examine policies proposed to mitigate the downsides of AI automation. It goes without saying that any kind of ban on forthcoming AI automation will run afoul of LCA. Anyone who advocates such a policy must, by LCA, also hold that a complete ban on mechanized agriculture in the historical counterfactual would have been a good idea. No reasonable person can think this. As we have noted, agriculture automation massively increased humankind's ability to produce food. It is how humanity escaped the Malthusian trap. Forgoing such gains is unthinkable. As it was unthinkable then, so LCA says it is unthinkable now. Thankfully, few people actually think we should ban AI automation outright. Doing so would

prove ineffective, as clandestine forms of AI automation would still be explored. Moreover, a complete ban might allow rival countries to gain an advantage in AI technology, with serious repercussions.

Many of the new AI automation skeptics do not propose complete bans on forthcoming AI automation but rather have more nuanced policy proposals. Danaher and Nyholm (2021: 234-235) believe there are four ways to respond to the fact that AI automation reduces the quality of work by making achievement elusive: (1) embrace other good-making features of work besides achievement not threatened by AI; (2) change the character of AI automated work to bring human achievement back into the fold; (3) deemphasize individual achievement and start emphasizing team achievement; (4) find avenues for achievement outside the workplace. Only (2) constitutes a *policy* response to AI automation; (1), (3), and (4) are what we might categorize as *personal* responses to it. Since LCA evaluates policies, we examine only proposal (2) in detail. This is not to say (1), (3), and (4) are flawed responses, only that they are about apples, and we are interested in oranges.

Proposal (2) says we should change the character of AI automated work to bring human achievement back into the picture. What would this look like in practice? The authors offer an example:

> … we could try to find a way for the 'human touch' to be retained in the automated workplace. For example, imagine a furniture factory that relies heavily on machines to manufacture tables and chairs. The machines produce safer and cheaper furniture than human workers ever could. Suppose though, at the very end of the manufacturing line, human workers add distinctive carvings or markings to the furniture. These workers are

> given a high degree of freedom and autonomy in doing this and hence are able to exercise their creativity when it comes to marking up the furniture (Danaher and Nyholm 2021: 235).

To test this proposal against LCA we must ask, what would this look like applied in the context of agriculture? It would require adding superfluous human elements to the automation process. For instance, though automation can plow fields, plant corn, harvest it, shuck it, pack it, etc., we might only allow robots to plow and harvest. Humans can still plant the seeds. They can shuck and then package the corn in individually made decorative parcels. We believe most would reject such a proposal. Adding superfluous human labor to corn production raises the price of corn by raising production costs. Corn is now more expensive. More expensive corn means people will go hungry, which is bad. If deployed at the dawn of mechanized farming, such a policy would keep humanity in the Malthusian trap longer than it otherwise needed to be. LCA thus tells us to reject the policy.

Let's turn to Celentano's proposal. Ideally, we automate away all tasks that do not contain good-making features of work, so only work with good-making features remains. Celentano recognizes this is not possible. AI automation that facilitates desirable work—she offers robots that aid nurses as an example—also requires undesirable work, like ghost work. Thus, her realistic proposal is to let automation happen, but simply *share* the residual undesirable work it produces. She writes: "the residue of AI-supporting tasks should be minimized, for example by sharing it through rotation or other organizational arrangements among human workers" (Celentano 2024: 39). For example, instead of having some people engage exclusively in the ghost work required to run AI systems everyone should spend time performing it. Perhaps

we all have jobs that, in general, contain good-making features, but we only perform them 30 hours a week instead of 40; we then all spend ten hours a week doing ghost work, which lacks good-making features.

Would it have been a good idea to impose this policy at the dawn of the age of farm automation? To answer this, we need to know what the policy would look like in that context. The tractor, we argued in the prior section, resulted in more undesirable work both on and off the farm. Let's consider off the farm: displaced workers went to work in factories, which was repetitive, hindering self-development. Celentano's proposal says all workers should be required to share this work. So, instead of having people work solely in factories, everyone—including lawyers, doctors, managers, entrepreneurs, etc.—spend one day a week in the factory to share the undesirable work.

This proposal would have resulted in a massive drag on economic productivity. Since Adam Smith's *Wealth of Nations*, we know that the division of labor is one of the most important contributing factors to economic prosperity. The division of labor makes us more productive because it increases workers' skill, results in less wasted time and, most importantly, results in more innovation (Smith 1982: ch. 1). To this John Stuart Mill (2006: 129) adds that the division of labor efficiently matches skills to tasks. To say everyone must do some factory (or other undesirable) work is to deny these benefits. This would result in less economic growth than we otherwise would have had if people were allowed to specialize. Growth is associated with all the good things in life, like happiness, health, education, tolerance, and so on (Kogelmann 2022: §2). Sacrificing growth just to share the undesirable work that resulted from mechanized agriculture would have been a bad idea. If that policy had been adopted, we would not be as happy as we are

today, we would not be as well-educated, we would not live as long, etc. Since sharing the bad work that resulted from the tractor would have been unacceptable, so too is sharing the bad work that results from AI automation.

Belic (2023: 17) sees two ways to regulate forthcoming AI automation: a *proactive* and a *reactive* approach. She favors the proactive. The proactive approach tries "to decrease the risk of the expectation of stability being frustrated" (Belic 2023: 17). The main policy proposal associated with this approach is *automation taxes*, or taxes on capital goods that substitute human labor. How do automation taxes result in less upset plans? If the taxes are sufficiently high, then they prevent all automation. This, we have seen, violates LCA: no reasonable person would stop agriculture automation, full stop. A more moderate version says the tax should not be so high as to stop all automation, but it should be high enough to slow it down. Instead of embracing the new automation technology tomorrow, the tax makes firms think twice. They will adopt the new technology eventually, just not right now. This gives individuals time to adjust their plans accordingly.

Even this moderate version of the proactive approach violates LCA. Suppose a hefty tax was levied on mechanized tractors that was not sufficiently high to prevent them from ever being adopted, but high enough to ensure their adoption was delayed by 30 years. As we have seen, mechanized farm tools like the tractor were a major factor in reducing famine. To say we should have pushed the adoption of the tractor back 30 years to allow farm workers to adjust their plans is to say we are okay with an extra 30 years of famine. This is unacceptable. If its benefits are

significant enough, pushing back the adoption of technology is a grave moral error. By LCA, this reasoning also applies to AI automation.[5]

Let's turn to Belic's reactive approach, which she is less sanguine about. This proposal says that "we could also see that the frustration of legitimate expectations amounts to a distinctive form of harm and as such, entitles those harmed to some form of compensation, all else being equal" (Belic 2023: 16). In other words, we simply write a check to those whose life plans are upset by automation to compensate them. In our view, this is the only policy proposal for addressing AI automation consistent with LCA. In the counterfactual, we see nothing objectionable about the government compensating farmers who lose their jobs to tractors. This is really just a special case of unemployment insurance, which was developed around the time of the tractor's adoption anyways (during the New Deal). Such a policy is not objectionable in the historical counterfactual because it does not in any way stop, delay, or mitigate the benefits of farm mechanization. The benefits of innovation are secured and those harmed by it are offered minimal financial security.

4.3 *LCA and Worst-Case Scenarios.*

[5] For separate criticism of automation taxes, see Parr (2022a; 2022b). One might argue that automation taxes should not be imposed to stop or slow down the adoption of automation (which LCA forbids), but rather to compensate workers. This, however, would not be an example of the *proactive* approach but rather an example of the *reactive* approach. The approaches are not defined by their policy instruments, but their goals. Taxing automation in a manner that does not stop or slow it down but is done to compensate workers after the fact is an example of the reactive approach, which we examine in the next paragraph. As it turns out, we believe this is the only policy proposal consistent with LCA.

LCA crucially depends on symmetry between AI and agricultural automation. Because the moral concerns with AI automation are similar to those present at the dawn of the automated agriculture age, whatever policy is proposed to address these concerns must also seem appropriate if applied to the historical case of farm mechanization. Like cases must be treated alike. What if, however, the moral bads of AI automation are significantly worse than the moral bads of automated agriculture? What if AI automation resulted in significant long-term unemployment? What if it resulted in an energy crisis due to the high energy demands of AI systems?[6]

At the beginning of this paper, we stipulated that we didn't want to focus on extreme and unlikely scenarios (e.g., the obsolescence of human work), but rather more moderate and realistic ones. Nonetheless, it is worth pondering what our analysis has to say, if anything, about these more extreme possibilities. To start off, LCA depends on symmetry between the two cases. If AI automation results in significant long-term unemployment or other novel bads that were not present during the automation of agriculture, then LCA falls silent. The principle tells us to treat like cases alike, so if cases are not alike (due to worst-case scenarios about AI automation coming true), then the principle should not be applied.

But even if LCA is no longer applicable, it does not follow that the general methodological approach of looking back to earlier forms of automation for guidance is without value. In fact, it might be helpful to consider a *counterfactual history* in which the doom-and-gloom concerns about AI automation were present during the automation of agriculture. From here, we can ask whether automating agriculture would still have been worth it and, if so, what

[6] The International Energy Agency estimates that data center energy usage by 2026 could increase to between 620 and 1,050 TWH—roughly equivalent to the energy demands of Sweden or Germany (Vincent 2024).

the appropriate policy response would have been. This might help us determine how to respond to doom-and-gloom AI scenarios.

Consider an example. Suppose AI automation will result in lower quality jobs, upset life plans, broader social, political, and cultural ills, as well as permanent 25 percent unemployment. Because automated agriculture did not result in permanent 25 percent unemployment, like cases are no longer alike, so LCA cannot be applied. But now let's consider an alternative history. Suppose automated agriculture resulted in lower quality jobs, upset life plans, broader social, political, and cultural ills, as well as permanent 25 percent unemployment. Would it have been worth it? Should the tractor have been banned?

This is a difficult question; careful analysis requires more space than we have. In our opinion, automated agriculture would have still been worth it in this counterfactual history *so long as* it was coupled with appropriate policy responses. Tools like the mechanized tractor are the reason there are no longer famines in developed countries; they are how humanity escaped the Malthusian trap. This is worth permanent 25 percent unemployment, so long as policies are in place to ensure the permanently unemployed live decent lives (e.g., a basic income guarantee). From this counterfactual history one might infer a similar response would be reasonable if AI had the impact posited: it would still be worth it, assuming an appropriate policy background to alleviate structural unemployment. Of course, this is only one example. AI automation can do far worse than 25 percent permanent unemployment; this would then require considering other counterfactual histories.

So, while doom-and-gloom scenarios mean LCA is no longer useful for policy analysis, our more general suggestion to consider historical cases of automation when thinking about how to address AI automation still has merit.

Works Cited

Acemoglu, Daron and James A. Robinson. 2012. *Why Nations Fail*. Crown Business.

Alesina, Alberto, Benedetta Brioschi, and Eliana la Ferrara. 2021. "Violence Against Women: A Cross-Cultural Analysis for Africa." *Economica* 88: 70-104.

Alesina, Alberto, Paola Giuliano, and Nathan Nunn. 2013. "On the Origins of Gender Roles: Women and the Plough." *Quarterly Journal of Economics* 128: 469-530.

Altman, Sam. 2021. "Moore's Law for Everything." URL: https://moores.samaltman.com/.

Ankli, Robert E. 1980. "Horses vs. Tractors on the Corn Belt." *Agricultural History* 54: 134-148.

Ankli, Robert E. and Alan L. Olmstead. 1981. "The Adoption of the Gasoline Tractor in California." *Agricultural History* 55: 213-230.

Autor, David, David A. Mindell, and Elisabeth B. Reynolds. 2023. *The Work of the Future*. MIT Press.

Belic, Jelena. 2023. "Institutions, Automation, and Legitimate Expectations." *Journal of Ethics* (Online First).

Bloor, David. 1991. *Knowledge and Social Imagery*. University of Chicago Press.

Bostrom, Nick. 2024. *Deep Utopia*. Ideapress.

Burley, Jake and Nir Eisikovits. 2023. "Workplace Automation and Political Replacement: a Valid Analogy?" *AI and Ethics* 3: 1361-1370.

Celentano, Denise. 2019. "Automation, Labour Justice, and Equality." *Ethics and Social Welfare* 13: 33-50.

Celentano, Denise. 2024. “Labor Automation for Fair Cooperation: Why and How Machines Should Provide Meaningful Work for All.” *Journal of Social Philosophy* 55: 25-43.

Chiang, Sheila. 2023. “A.I. is not all hype. It’s the ‘fourth industrial revolution playing out,’ says Wedbush’s Dan Ives.” *CNBC* (28 June 2023). URL: https://www.cnbc.com/2023/06/28/ai-is-not-all-hype-its-the-fourth-industrial-revolution-dan-ives.html.

Clark, Gregory. 2007. *A Farewell to Alms*. Princeton University Press.

Danaher, John. 2019. *Automation and Utopia*. Harvard University Press.

Danaher, John and Sven Nyholm. 2021. “Automation, Work, and the Achievement Gap.” *AI and Ethics* 1: 227-237.

Federico, Giovanni. 2005. *Feeding the World*. Princeton University Press.

Foster, Andrew D. and Mark R. Rosenzweig. 2022. “Are There Too Many Farms in the World? Labor Market Transaction Costs, Machine Capacities, and Optimal Farm Size.” *Journal of Political Economy* 130: 636-680.

Frey, Carl B. and Michael A. Osborne. 2017. “The Future of Employment: How Susceptible are Jobs to Computerization?” *Technological Forecasting and Social Change* 114: 254-280.

Friedland, William H. 1991. “Engineering Social Change in Agriculture.” *University of Dayton Review* 21: 25-42.

Gordon, Robert J. 2016. *The Rise and Fall of American Growth*. Princeton University Press.

Guilford-Blake, Roxanna. 2020. “Wait. Will AI Replace Radiologists After All?” *Radiology Business* (18 February 2020). URL: https://radiologybusiness.com/topics/artificial-intelligence/wait-will-ai-replace-radiologists-after-all.

Hopster, J.K.G., C. Arora, C. Blunden, C. Eriksen, L.E. Frank, J.S. Hermann, M.B.O.T. Klenk, E.R.H. O’Neill, and S. Steiner. 2022. “Pistols, Pills, Pork, and Ploughs: The Structure of Technomoral Revolutions.” *Inquiry* (Online First).

Juma, Calestous. 2016. *Innovation and Its Enemies*. Oxford University Press.

Kelly, Jack. 2023. “Goldman Sachs Predicts 300 Million Jobs Will Be Lost Or Degraded By Artificial Intelligence.” *Forbes* (31 March 2023). URL: https://www.forbes.com/sites/jackkelly/2023/03/31/goldman-sachs-predicts-300-million-jobs-will-be-lost-or-degraded-by-artificial-intelligence/?sh=7eb44c5782b4.

Knell, Sebastian and Markus Rüther. 2024. “Artificial Intelligence, Superefficiency and the End of Work: a Humanistic Perspective on the Meaning of Life.” *AI and Ethics* 4: 363-373.

Kogelmann, Brian. 2021. *Secret Government: The Pathologies of Publicity*. Cambridge University Press.

Kogelmann, Brian. 2022. “We Must Always Pursue Economic Growth.” *Utilitas* 34: 478-492.

Kogelmann, Brian. 2024. “Creative Destruction and the Autonomous Life.” *Journal of Business Ethics* (Online First).

Kyriacou, Christos. 2024. “Artificial Moral Intelligence and Computability: an Aristotelian Perspective.” *AI and Ethics* (Online First).

Majka, Linda C. 1981. “Labor Militancy among Farm Workers and the Strategy of Protest: 1900-1979.” *Social Problems* 28: 533-547.

Malthus, Thomas. 2015. *An Essay on the Principles of Population and Other Writings*. Penguin Classics.

Marx, Karl. 1990. *Capital, Volume 1*. Penguin Classics.

McPherson, Tristram. 2022. "Supervenience in Ethics." *The Stanford Encyclopedia of Philosophy* (Summer 2022 Edition), ed. Edward N. Zalta, URL = <https://plato.stanford.edu/archives/sum2022/entries/supervenience-ethics/>.

Mill, John Stuart. 1978. *On Liberty*, edited by Elizabeth Rapaport. Hackett Publishing.

Mill, John Stuart. 2006. *Collected Works of John Stuart Mill, Vol. 1: Principles of Political Economy*. Liberty Fund.

Mokyr, Joel. 2002. *The Gifts of Athena*. Princeton University Press.

Olmstead, Alan L. and Paul W. Rhode. 1994. "The Agricultural Mechanization Controversy of the Interwar Years." *Agricultural History* 68: 35-53.

Parr, Tom. 2022a. "Automation, Unemployment, and Insurance." *Ethics and Information Technology* 24: 1-11.

Parr, Tom. 2022b. "Automation, Unemployment, and Taxation." *Social Theory and Practice* 48: 357-378.

Prakash, Prarthana. 2023. "Alphabet CEO Sundar Pichai says that A.I. could be 'more profound' than both fire and electricity—but he's been saying the same thing for years." *Fortune* (17 April 2023). URL: https://fortune.com/2023/04/17/sundar-pichai-a-i-more-profound-than-fire-electricity/.

Raz, Jospeh. 1986. *The Morality of Freedom*. Oxford University Press.

Ricardo, David. 2004. *The Principles of Political Economy and Taxation*. Dover Publications.

Sen, Amartya. 1981. *Poverty and Famines*. Oxford University Press.

Schmitz, Andrew and Charles B. Moss. 2015. "Mechanized Agriculture: Machine Adoption, Farm Size, and Labor Displacement." *AgBioForum* 18: 278-296.

Smith, Adam. 1982. *An Inquiry into the Nature and Causes of the Wealth of Nations, Vol. 1*. Liberty Fund.

Spencer, David A. 2024. "AI, Automation, and the Lightening of Work." *AI and Society* (Online First).

Steckel, Richard H. and William J. White. 2012. "Engines of Growth: Farm Tractors and Twentieth Century U.S. Economic Welfare." NBER Working Paper 17879.

Taylor, Mark Zachary. 2016. *The Politics of Innovation*. Oxford University Press.

Valdés, Dennis Nodín. 1994. "Machine Politics in California Agriculture, 1945-1990s." *Pacific Historical Review* 63: 203-224.

Veltman, Andrea. 2016. *Meaningful Work*. Oxford University Press.

Vincent, James. 2024. "How much electricity does AI consume?" *The Verge* (16 February 2024). URL: https://www.theverge.com/24066646/ai-electricity-energy-watts-generative-consumption#.

Wall, Steven. 1998. *Liberalism, Perfectionism and Restraint*. Cambridge University Press.

Weiss, Debra Cassens. 2023. "Latest version of ChatGPT aces bar exam with score nearing 90th percentile." *ABA Journal* (16 March 2023). URL =

https://www.abajournal.com/web/article/latest-version-of-chatgpt-aces-the-bar-exam-with-score-in-90th-percentile.