

Creative Destruction and the Autonomous Life

1. Introduction

At the heart of capitalism is *creative destruction*. Creative destruction is “the process by which new innovations continually emerge and render existing technologies obsolete, new firms continually arrive to compete with existing firms, and new jobs and activities arise and replace existing jobs and activities” (Aghion *et al* 2021: 1). Examples include industrial looms replacing artisan weavers, Uber and Lyft putting taxi drivers out of business, Facebook winning out over Myspace, tractor combines and other advanced agricultural tools reducing the need for farmhands, industrial robots replacing workers on assembly lines, automobiles destroying a thriving horse and buggy industry, electricity reducing the demand for kerosene and whale oil, Netflix bankrupting Blockbuster, high-quality cellphone cameras bankrupting Kodak, and much more. Joseph Schumpeter (2008: 81) introduced the term, but the idea that destruction, instability, and turmoil are part of the capitalist process is not original to him. Karl Marx and Friedrich Engels recognized this feature of capitalism when they wrote in the *Manifesto of the Communist Party* that “the bourgeoisie cannot exist without constantly revolutionizing the instruments of production, and thereby the relations of production, and with them the whole relations of society” (Marx and Engels 1972: 476). In capitalist societies, “all that is solid melts into air” (Marx and Engels 1972: 476).

There are good reasons to embrace creative destruction. Economic growth is profoundly important. It increases happiness (Stevenson and Wolfers 2013), it helps people lead objectively good lives independent its effect on their happiness (Moller 2011), it has a positive impact on our

political attitudes and institutions (Friedman 2005), and it follows from precepts of commonsense morality (Kogelmann 2022; Carroll 2023). While scholars still debate the ultimate cause of economic growth, its proximate cause is understood: innovation (Schumpeter 2017: 64; Koyama and Rubin 2022: 9). The handmaiden of innovation, however, is creative destruction (Acemoglu and Robinson 2012: 84).¹

For example, technological innovation can result in machines that automate human-performed tasks, lowering production costs and thus prices; goods and services, as a result, are more affordable for all. This is what economic growth is all about. However, these machines can also put people out of work and bankrupt incumbent firms. Innovation that induces creative destruction need not be technological, understood narrowly. For example, Amazon did not introduce (at least initially) industrial robots to automate human labor; their main innovation was in logistics. They created wealth by reducing transaction costs. But Amazon's innovation also resulted in a decline of traditional storefronts, resulting in the loss of businesses, jobs, and, in some cases, entire communities (MacGillis 2021).

While there are good reasons to embrace creative destruction, it is still destructive. What should we think about this destruction from the moral point of view? Quite surprisingly, there exists no philosophical work (as far as I can tell) on the ethics of creative destruction. Perhaps that is because the question is intimidatingly large; there is no shortage of ethical frameworks that can be adopted to analyze it. I take up only one in this paper. I examine the extent to which creative destruction undermines the ideal of autonomy.

¹ This does not imply that innovation is the *only* cause of creative destruction. Globalization—that is, an expansion of the extent of the market—can cause creative destruction when jobs are offshored to countries with lower costs of production (Autor *et al* 2013). For ease of exposition, I focus on creative destruction induced by innovation, but everything I say should apply, *mutatis mutandis*, to creative destruction induced by other causes as well.

According to many, an autonomous life is a planned life (§2). Autonomous agents form, pursue, and in many cases successfully execute life plans. Creative destruction conflicts with this conception of autonomy, I argue, because it undermines the conditions necessary for successful planning (§3). This creates a dilemma: we must either give up on the ideal of autonomy, or give up on economic growth, for creative destruction (I noted above) accompanies economic growth. Neither option is attractive.

A potential strategy to address this dilemma is to strive for a middle ground by establishing a regulatory framework to govern the process of technological change (§4). By regulating the process of technological change, the hope is that economic growth can be achieved in a manner that does not upset agents' life plans. I consider several possible regulatory strategies; they all fail, I argue. To end, I chart a different response to the dilemma: perhaps we should rethink what it means to live an autonomous life in capitalist societies (§5). I propose and defend a novel conception of autonomy that is consistent with creative destruction. I argue that those who find the traditional conception of autonomy attractive should also find my novel conception attractive as well. The upshot is that the dilemma can be avoided. We can have economic growth and the creative destruction that accompanies it without giving up on the ideal of an autonomous life.

Reflection on the ethical aspects of creative destruction could not be more topical. Recent advances in artificial intelligence (AI) such as OpenAI's ChatGPT and Google's Gemini suggest that a long stretch of creative destruction lies just over the horizon. Google CEO Sundar Pichai says AI will be a "more profound" innovation than humanity's harnessing of fire or electricity (Prakash 2023). Some have said that AI will usher in the next Industrial Revolution (Chiang 2023). Pew Research Center estimates that about a fifth of U.S. workers have jobs that are highly

exposed to AI, meaning they are likely candidates for automation (Kochhar 2023). Goldman Sachs estimates that 300 million jobs worldwide could be eliminated or degraded due to AI (Kelly 2023). If these sensational claims turn out to be even half true, then the world will experience creative destruction the likes of which it has never seen before. What should we think about this from the moral point of view? This paper takes the first step towards answering this urgent question.

2. The Ideal of Autonomy

Autonomy is a complex concept for which there are many conceptions (Buss and Westlund 2018; Christman 2020). The conception I focus on understands autonomy as self-government or self-authorship. John Stuart Mill (1978) was one of the first to state and defend this conception of autonomy, but it has been developed most fully in the work of Joseph Raz (1986) and Steven Wall (1998). Here are some statements of this conception of autonomy:

He who lets the world ... choose his plan of life for him has no need of any other faculty than the ape-like one of imitation. He who chooses his plan for himself employs all his faculties (Mill 1978: 56).

The ruling idea behind the ideal of personal autonomy is that people should make their own lives. The autonomous person is a (part) author of his own life. The ideal of personal autonomy is the vision of people controlling, to some degree, their own destiny, fashioning it through successive decisions throughout their lives (Raz 1986: 369).

It is the ideal of people charting their own course through life, fashioning their character by self-consciously choosing projects and taking up commitments from a wide range of

eligible alternatives, and making something out of their lives according to their own understanding of what is valuable and worth doing ... In short, autonomous people have a strong sense of their own identity and actively participate in the determination of their own lives (Wall 1998: 128).

Autonomous agents form, pursue, and in many cases successfully execute life plans. Though sophisticated philosophical accounts of planning exist (e.g., Bratman 1987; Rawls 1971: §63; Bratman 2007), such rigor is not needed for our purposes. We can understand life plans as consisting of a set of goals and a set of actions by which these goals are pursued (Kogelmann 2021: 100). For instance, an agent might give herself the following goals: she wants to be a lawyer, a mother, and a marathon runner. Next, the autonomous agent comes up with a set of actions that, if pursued, are likely to achieve these goals. To become a lawyer, she must first go to law school; to become a mother, she must first find a suitable partner and achieve financial stability; before she runs her first marathon, she must start training and eating right. She then performs these actions.

Why think autonomy as it has just been articulated is an attractive normative ideal? There are three reasons, and I shall go through them now. I do this because the reasons to embrace autonomy as a normative ideal will be relevant in §5 when I propose a novel conception of autonomy. To foreshadow, I shall argue that the reasons in defense of the conception of autonomy currently under consideration also support the novel conception of autonomy I develop. Thus, *if* one finds autonomy as self-authorship attractive *then* one should also embrace my new conception.

First, autonomy is an attractive normative ideal because it is bad to be overly dependent on the will of others and autonomous agents are, by definition, not overly dependent on the will

of others (Wall 1998: 146-148). They are, instead, self-authors and self-governed. Consider the following case, from Wall:

Suppose I know that you are wise and that you have an excellent understanding of what is good for me. You know my talents, temperament and vulnerabilities and you know what types of projects would best suit my nature. Further suppose I know that you are a person of goodwill who cares about my well-being. Given these facts, we can ask: Would my life go better if I let you take control of it? Would it be a better life if I always turned to you for direction as to what I should do before I took up any project or commitment? (Wall 1998: 146).

Wall's answer is "no." Many (myself included) share his judgment. What explains this intuition? It cannot be that reliance on the advisor will make the agent's life go worse because the advisor is, by hypothesis, both wise and goodwilled; if anything, following the advisor's counsel will probably make the agent's life go better. The explanation of the intuition must be that there is something bad about being too dependent on the will of another. If the agent always does as the wise and goodwilled advisor tells her, then she is overly dependent on someone else's will. If the agent takes charge of her own affairs (that is, if she is autonomous), then she is independent of anyone else's will. This seems desirable *even if* her life does not go as well as it would were the agent to constantly rely on the advisor's counsel.

Second, autonomy is an attractive normative ideal because there is something valuable about making choices (Wall 1998: 148-149). Consider a case. You face two choice scenarios. In the first scenario you must choose a career; only one option (call it X) is available, but it turns out that X is suitable for you given your talents and abilities. In the second scenario there are five career options available, of which X is one. You end up choosing X in the second scenario

because it suits you best. Most of us would rather find ourselves in the second scenario than the first. What explains this intuition? It cannot be that our life goes better in the second scenario, because by assumption we end up choosing X (which is objectively the best option) in both cases. The most plausible explanation of the intuition is that there is something valuable about making choices. The second scenario is more desirable than the first because in it choice is possible. The autonomous life is all about choice. It is about *choosing* a set of goals to pursue and then *choosing* the best means to achieve them.

Third, autonomy is an attractive normative ideal because living an autonomous life forces us to exercise our skills and capacities, and that is a good thing (Wall 1998: 150-159). This is the argument Mill offers in defense of the ideal of autonomy. In terms of the first part of the argument—the claim that living an autonomous life forces persons to exercise their skills and capacities—Mill has this to say:

He who chooses his plan for himself employs all his faculties. He must use observation to see, reasoning and judgment to foresee, activity to gather materials for decision, discrimination to decide, and when he has decided, firmness and self-control to hold to his deliberate decision (Mill 1978: 56).

That many different skills and capacities are required to form, pursue, and successfully execute life plans does not seem too controversial a claim. That it is *good* to develop and exercise different skills and capacities also does not seem too controversial. On this point, Mill (1978: 54-55) follows German philosopher Wilhelm von Humbolt, who writes that “the end of man ... is the highest and most harmonious development of his powers to a complete and consistent whole” (Humbolt 1993: 10).

There are thus three reasons to think autonomy is an attractive normative ideal: too much dependence on the will of another is bad; being confronted with and making meaningful choices is good; and forming, pursuing, and executing life plans exercises our skills and capacities, which is good.

Certain background conditions are needed for agents to form, pursue, and successfully execute life plans. According to the defenders of this ideal, it is the state's duty to ensure these conditions are in place (Wall 1998: 132; Raz 1986: 372). For instance, agents must possess certain capabilities and traits. They must have the requisite mental capacities to "form intentions of a sufficiently complex kind, and plan their execution" (Raz 1986: 327). To take charge of their own affairs, autonomous agents also need a certain kind of motivational and psychological constitution: they must have "vigor," and be free of various afflictions such as "world-weariness, emotional distress, depression, laziness and perhaps a growing sense of the meaninglessness of the world and one's place in it" (Wall 1998: 139). Moreover, autonomous agents must be independent (Raz 1986: 378). They cannot be too reliant on or subservient to the will of another, lest they give up authorship of their lives. I do not discuss these conditions any further because creative destruction does not undermine them.

Other conditions are (I shall argue below) threatened by creative destruction, in particular the requirement that agents need adequate choice options to be autonomous. What is needed for a set of choice options to be considered adequate? Three things. First, there must be sufficient variety (Raz 1986: 375). If you live in a society where only one occupation is open to you—say, working on your parent's farm—then you cannot be autonomous because there is no possibility of self-authorship. There are no choices for you to make. Wall writes: "Having access to two options that are significantly different may be better than having access to ten options that are

very much alike ... a person's option set must include a range of significantly different options" (Wall 1998: 141). Call this the *Variety Condition*.

Second, choice options must be suitable for individuals, in that "a person's option set should include options that give him the opportunity to develop his talents and capacities" (Wall 1998: 142; Raz 1986: 376). If various kinds of manual labor exhaust your career options, then perhaps there is sufficient variety in the set; however, if your gifts and interests are primarily intellectual, then this cornucopia will not provide opportunities to develop your talents, so it is inadequate. For your choice options to be adequate in this case specifically they must include some careers that develop your intellectual skills and talents. Call this the *Development Condition*.

A third component of adequate choice options is that, once an option has been selected and pursued for a sufficient amount of time, it must *remain* an option for the individual who pursues it (Raz 1986: 411). Wall (1998: 143) writes: "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." Call this the *Stability Condition*. While the Variety and Development Conditions are intuitive, the Stability Condition is not. Why is it needed for individuals to live autonomous lives?

Consider a thought experiment. You live in a country where many different career options are available that would allow you to develop your skills and talents. So, the Variety and Development Conditions are satisfied. However, there is a law in your country that says every six months you must quit your job and start a new one in an unrelated field, so the Stability Condition is violated. Every six months you face a new set of choice options; they always

include sufficient variety and always include several options that would allow you to develop your skills and talents. So, the requirement that you leave your job every six months does not violate the Variety and Development Conditions.

Even though our first two conditions are satisfied, an autonomous life is still elusive. This is because autonomy is not just about forming and pursuing plans, but also successfully *executing* them. In the case under consideration, you form plans every six months and begin pursuing them. But, because most career plans take longer than six months to achieve, you successfully execute very few (if any) of them. To ensure agents successfully execute the plans they form and pursue, the options they select and pursue must remain available for them to continue pursuing. This is why the Stability Condition is important.

Let me summarize the bits of this section that will be relevant for the arguments below. Autonomous agents form, pursue, and in many cases successfully execute life plans. There are three reasons to embrace this conception of autonomy as a normative ideal: too much dependence on the will of another is bad; being confronted with and making meaningful choices is good; and forming, pursuing, and executing life plans exercises our skills and capacities, which is good. To be autonomous, certain conditions must be in place. Particularly relevant is the necessity of adequate choice options. For choice options to be adequate the Variety, Development, and Stability Conditions must be satisfied.

3. Creative Destruction Undermines Autonomy

In this section I argue that creative destruction can undermine autonomy. It does so by undermining the conditions needed for agents to form, pursue, and successfully execute their life

plans. Creative destruction can leave agents with inadequate choice options. I do not think creative destruction undermines the Variety Condition, but I contend it undermines the Development and Stability Conditions.

Before getting to that, it is worth examining the relationship between creative destruction and the Variety Condition, for it is not implausible to think the former can undermine the latter. Creative destruction eliminates career options. Because of technological change we no longer have horse and buggy drivers, knocker-uppers, switchboard operators, whalers, and much more; we have far less people who work in occupations like manufacturing and farming, among others. If creative destruction continually eliminates jobs, then the set of available career options shrinks. The Variety Condition, however, demands a sufficiently expansive set of available career options. So, if creative destruction eliminates too many career options, then the Variety Condition will be violated.

This argument fails, because creative destruction creates new jobs as it eliminates old ones (Autor 2015: 5; Ridley 2020: 289-294; Aghion *et al* 2021: 53). It destroys jobs by introducing technology that substitutes certain kinds of labor. Horse and buggy drivers are out of work when the automobile is introduced; artisan textile producers are replaced by industrial looms. It creates jobs because new technology complements and thus increases demand for certain kinds of labor. When cars are widely used, you now need mechanics; to produce and use computers, you need software engineers and information technology professionals. The argument in the paragraph above thus errs when it infers that there will be an ever-shrinking set of career options from the fact that creative destruction eliminates jobs. Creative destruction *does* eliminate jobs, but it does not shrink the set of available career options, because it creates jobs as

well. Creative destruction changes the composition of the set of available career options, but does not shrink it.

Although creative destruction has thus far created new jobs as it destroyed old ones, some worry that things will be different with the introduction of powerful AI (Brynjolfsson and McAfee 2016: 180-181). AI robots already surpass human beings at many tasks. For instance, IBM's Deep Blue is a chess-playing robot that has beaten world champions (Goodrich 2021); DeepMind has solved protein folding problems that have eluded humans (Service 2020). If AI robots surpass humans at *all* tasks, then perhaps new jobs will not replace the old jobs that are eliminated. This puts the Variety Condition in jeopardy. The set of available career options will shrink, and no new jobs will enter as replacements. When the set of available career options gets too small, agents' choice options will be inadequate.

This argument also fails. Even if AI robots are better than humans at every economic task, there will still be work for humans so long as there is an opportunity cost to using robot labor, which will be true so long as the supply of robot labor is finite. This follows from the law of comparative advantage. As David Ricardo (2004: ch. 7) taught us, it is mutually beneficial for party *A* to trade with party *B* even if *A* is better than *B* at all tasks. This is because *A*'s labor is finite and thus subject to opportunity costs. Though *A* may be better than *B* at both tasks *x* and *y*, she will be more productive at one relative to the other. Suppose she generates more wealth doing *x* than *y*. If true, then *A* is better off if she focuses all her energy on *x* and lets *B* do *y* even though she can produce more wealth doing *y* than *B* can, because the opportunity cost of *A* doing *y* is not doing more *x* which, by hypothesis, produces more wealth for her than doing *y* does. Similarly, if robots are better than humans at both *x* and *y*, it will still be the case that one of the activities (say, *x*) is more profitable for them to engage in than the other. Robots will specialize in

x and humans in y . So long as resources are finite, eliminated jobs will be replaced by new ones, so there is no reason to think AI will lead to a shrinking set of career options. Powerful AI may cause other social problems, but it is not a unique threat to the Variety Condition.

Creative destruction does not threaten the Variety Condition because it creates new career opportunities as old ones are destroyed. But though the *size* of the set of career opportunities should remain unaffected by creative destruction, its *composition* will not. Consider an example. Suppose there are two categories of career types: manual and cognitive labor. There are three unique career options available in the manual labor category and three unique career options available in the cognitive labor category. Six unique career options, let us suppose, is sufficient for the Variety Condition to be satisfied. Suppose creative destruction eliminates the manual labor category entirely. Because creative destruction creates new jobs as it destroys old ones, there are now three new career options available in the cognitive labor category. At the end of the day, there are still six unique career options available—so if the Variety Condition was not a concern before, it should not be a concern now—but all these career options involve cognitive, and none involve manual, labor.

That creative destruction can change the composition of jobs in an economy threatens the Development Condition, which says that for agents to be autonomous they must have career options that develop their talents and capacities. To continue the example above, suppose you are gifted at and enjoy working with your hands. By contrast, you have always struggled with cognitive work and never found it rewarding. Before creative destruction there are three unique career options in the manual labor category for you to choose from; all three allow you to develop your talents and abilities, so the Development Condition is satisfied. After creative destruction, however, only career options in the cognitive labor category remain. There are still

enough options present for the Variety Condition to be satisfied, but none of these options develop your skills and talents. You are good at and like working with your hands. Because of creative destruction, no career options will allow you to refine these skills. The Development Condition is violated.

I just offered a hypothetical case to show how creative destruction can undermine the Development Condition. The concern is not merely hypothetical. It has actually happened and will continue to happen. Relevant here is what economists call *skill-biased technological change* (Goldin and Katz 2008: 90). Skilled-biased technological change “is a shift in the production technology that favors skilled over unskilled labor by increasing its relative productivity and, therefore, its relative demand” (Violante 2008: 1). An example is the impact of the computer on the economy in the late twentieth century (Autor *et al* 2006; Autor 2017: 247-250). We can divide the set of all jobs into four different categories, illustrated in Table 1 (Page 2017: 41). Jobs can either be manual (working with your hands) or cognitive (working with your head). They can also either be routine (performing the same tasks repeatedly) or non-routine (performing many different tasks that are difficult to codify). The computer substituted routine labor. This shrunk the number of career options in quadrants I and III and—because creative destruction creates new jobs as it destroys old ones—increased the number of career options in quadrants II and IV. Due to computers, there are far less clerical (quadrant I) and assembly line (quadrant III) jobs. There are far more jobs in management and software engineering (quadrant II) as well as services (quadrant IV).

	Routine Work	Non-Routine Work
Cognitive Work	I	II

Manual Work	III	IV
-------------	-----	----

Table 1

Economists are concerned with the economic impact of skill-biased technological change. It tends to raise wages among the educated and depress wages among the uneducated, leading to greater inequality (Autor 2014). I am concerned with its impact on the Development Condition. There may be people whose skills and talents naturally lend themselves to routine work, either of the cognitive or manual variety. To fully develop their capabilities, these individuals need jobs in quadrants I and III. Creative destruction has eliminated many of these opportunities. Other career options are available, for new jobs are being created in quadrants II and IV (you can be a software engineer or drive an Uber). However, the Development Condition is about matching skills to careers, and insofar as entire types of work (such as routine) disappear, that is a problem. AI may further eliminate routine work. Those who are naturally suited to this kind of work and find it rewarding will find themselves with even fewer avenues to develop their skills and talents going forward.

Creative destruction also threatens the Stability Condition. This third and final condition says that for choice options to be adequate, once an option has been selected and pursued for a sufficient amount of time, it must remain an option for the individual who pursues it. Creative destruction clearly undermines this condition because it eliminates career options. Horse and buggy drivers saw their plans upset by the automobile, those working in manufacturing plants saw their plans upset by industrial robots, and taxi drivers saw their plans upset by Uber and Lyft, to name just a few examples. New AI tools will soon upset the career plans of many. For instance, machine learning algorithms are already better than humans at diagnosing disease from

medical imaging, leading some to speculate that human radiologists will soon be obsolete (Guilford-Blake 2020). If true, this will upset the plans of all those who are currently in or who have recently graduated from radiology training programs. Properly trained AI can now score better than most human lawyers on the bar exam (Weiss 2023). This might greatly reduce the demand for lawyers. If it does, it will upset the plans of all those who are currently in or who have recently graduated from law school.

Unstable choice options induced by creative destruction can threaten autonomy in a more subtle and sinister way: it can induce *paralysis*, such that agents are unwilling to form and pursue plans in the first place because they are so uncertain about the future. In the literature on property rights and economic development, several argue that insecure property rights can lead property owners to not invest in their property at all for fear of the unknown (Acemoglu and Robinson 2012: 75). Something similar may happen when agents' choice options are in flux due to creative destruction. If new technologies continuously change the occupational landscape, it may discourage some from choosing a career at all. As one possible data point in support of this thesis, in regions that have experienced significant creative destruction, many have withdrawn from the labor force completely (Acemoglu and Restrepo 2020: 2219). That is, in regions that have experienced significant creative destruction, many working-aged people have stopped even looking for work. Why they are doing so is unknown. One possibility is that there is so much instability in choice options it becomes hard for agents to choose a path forward. Their jobs in manufacturing have been eliminated. Why expend resources to train for a new profession—say, radiology—if it too could soon be eliminated?

To sum up, creative destruction can undermine the conditions needed for individuals to live autonomous lives. By undermining the Development and Stability Conditions, creative

destruction makes it difficult for individuals to form, pursue, and execute life plans. If you think the ideal of autonomy is attractive, this is troubling.

To close this section, I want to address an inconsistency between the thesis I have just defended and the work of Raz and Wall. Neither see an inconsistency between autonomy and creative destruction. Quite the opposite, in fact. Both argue that the ideal of autonomy is perfectly suited for capitalist societies and all the churn and change they bring along with them.

Raz writes:

[Autonomy] is an ideal particularly suited to the conditions of the industrial age and its aftermath with their fast changing technologies and free movement of labour. They call for an ability to cope with changing technological, economic and social conditions, for an ability to adjust, to acquire new skills, to move from one subculture to another, to come to terms with new scientific and moral views (Raz 2009: 369-370).

Wall (1998: 166-168) makes similar remarks.² So, according to some of its main defenders, the ideal of autonomy is desirable precisely because it prepares people to live in a modern age characterized by creative destruction. I have just argued for the exact opposite conclusion. What explains the disagreement? Who is right?

Raz and Wall wrongly conflate being skilled at x and successfully achieving x . These two things typically go together, but they can come apart. In many cases, something that increases your skill at x also increases the chances that you will successfully achieve x . A training program that increases your skill as a football player also increases the chances of you making the NFL

² In a later piece, Wall (2018: 93) admits that creative destruction and autonomy may be in tension with one another: “There is a dimension of the market order that disrupts, rather than contributes to, a planning-friendly environment. The ‘creative destruction’ of the market order is necessary to its efficiency, but it also unsettles the plans and expectations of all those who are its victims.” Wall does not say anything about the relationship between autonomy and creative destruction beyond these remarks.

one day. However, in some cases, something that increases your skill at x can *reduce* the probability you achieve x . For example, living in an active warzone may increase your survival skills—you learn how to filter your own water, scavenge for food, how to cauterize wounds, and so on—but decreases your chances of actually surviving. Living through an economic recession may increase your ability to save money—because times are tough, you find little ways to save—but likely decreases your overall savings if you end up unemployed.

Similarly, though living in an environment characterized by protracted creative destruction may increase your ability to form and pursue plans—you have to come up with a new plan if your chosen occupation is eliminated!—it likely also reduces the probability of you successfully executing your plans. Successfully executing life plans, however, is what being autonomous is all about. So, while Raz and Wall may be right that creative destruction makes us *skilled* planners because it forces us to exercise our planning skills more often than we otherwise would, this does not necessarily mean it will make us *successful* planners, in that we actually achieve the plans we give ourselves. And in fact, I have just argued that creative destruction undermines successful planning.

4. Threading the Needle?

Creative destruction undermines autonomy. What ought we do in response? Perhaps we should turn our back on the ideal of autonomy so we can embrace creative destruction and the economic growth that accompanies it. A different approach says we should somehow suppress creative destruction and thus sacrifice economic growth to ensure agents can form, pursue, and successfully execute their life plans. This latter approach can be accomplished with the power of

the state. Historically it was common for states to suppress new technological innovations (Juma 2019). While the motives for doing so were no doubt complex, in some cases rulers cited the wellbeing of those who would be displaced by creative destruction. For instance, in 1589, Queen Elizabeth I denied a patent to an automated knitting machine because she had “too much love for my poor people who obtain their bread by the employment of knitting ... to forward an invention that will tend to their ruin by depriving them of employment, and thus make them beggars” (Diamond 2019: 66).

Neither option is attractive. The economic growth that accompanies creative destruction is important, and it is good for people to lead autonomous lives. Is there a way to have both? I explore that possibility in the current section. Instead of suppressing new technological innovation, perhaps the state can regulate its implementation to achieve the best of both worlds. By regulating rather than suppressing creative destruction, the hope is that economic growth can be achieved in a manner that does not upset agents’ life plans. In what follows I consider regulatory strategies the state can adopt to ensure the Stability Condition is satisfied that do not involve the wholesale suppression of creative destruction. My focus on the Stability Condition specifically does not imply that the Development Condition is unimportant. Because both conditions must be present for agents to live autonomous lives, showing that one condition (in this case, the Stability Condition) cannot be made consistent with creative destruction is enough to thwart the hopes of the regulatory middle way. That is precisely what I do in the current section.

What can the state do to ensure creative destruction does not undermine the Stability Condition? The problem, recall, is that creative destruction eliminates career options while agents pursue them, which prevents them from successfully executing their plans. I can think of

two things the state can do to stop this from happening without completely suppressing creative destruction. First, the state might try to encourage individuals to not choose careers threatened by imminent creative destruction. This is the *ex ante* strategy. Second, the state might ensure that individuals can complete their plans even after creative destruction has rendered their skills and labor obsolete by subsidizing their employment. This is the *ex post* strategy.

Turn first to the *ex ante* strategy. How can the state encourage people to avoid careers threatened by imminent creative destruction? Perhaps it can regulate the pace of technological change. Some scholars of innovation argue that it would be better for a variety of reasons if the pace of technological change was slower (Vogt 2016; Woodhouse 2016). I interpret these scholars as saying that it would be better if there was more time *between* major innovations. This would not prevent individuals from choosing careers that will soon be eliminated. People choose careers threatened by creative destruction because innovations are *unexpected*. Horse and buggy drivers whose plans were upset by the automobile chose that career because the automobile was not on their radar; whether the automobile became hegemonic in 1900, 1925, or 1950 would not have changed this. The Stability Condition cannot be rescued by altering *when* innovation happens; to save it, the state must do something to allow agents to better *anticipate* major innovations and thus build their plans around them.

Perhaps the state should adopt a regulation that says there must be a lag of x years between when a new technological innovation is announced to the public and when it can be used as a factor of production. This allows individuals to avoid career options that will soon be eliminated. If it was common knowledge that the automobile was going to replace the horse and buggy in ten years, then would-be horse and buggy drivers would do something else instead. Such knowledge could have been facilitated with a legally mandated lag between when the

automobile entered the public conscience and when it could be adopted for use. Adoption of AI tools like ChatGPT was already widespread as people learned about them. However, suppose OpenAI displayed ChatGPT's capabilities to the public, yet businesses were not allowed to use it until 2040. This would give individuals time to adjust their plans around the new technology. Perhaps some decide not to go to law school on the basis that tools like ChatGPT may radically reduce the demand for lawyers and thus eliminate that career option.

There are three problems with the mandatory lag proposal. First, it is not clear how to implement the regulation, because it is not clear what counts as a new technological innovation. Major innovations—like the automobile, airplane, AI, and more—are almost never the result of a giant leap forward. The closer you look at the history of innovation “the less likely you are to find a moment of sudden breakthrough, rather than a series of small incremental steps. There is no day when you can say: computers did not exist the day before and did the day after” (Ridley 2020: 241; Mokyr 1990: 12). While the Wright brothers were the first to take flight, many more innovations and iterations of the airplane were needed before airline travel was commercially viable and thus a threat to existing modes of transportation. Should there have been a mandatory lag between each and every iteration? Obviously not. If that were the case, then we would still be traveling by train and boat. But if you do not force the mandatory lag for every iteration, then how do you select when to enforce it? There have been thousands of tiny innovations between the Wright brothers' first flight and today's airplanes, but it would have been very difficult to know *ex ante* which innovation would be the tipping point that allowed air travel to compete commercially with existing modes of transportation. Yet the tipping-point innovation is the one you must apply the mandatory lag to if you want to successfully encourage people to avoid careers threatened by imminent creative destruction.

Second, publicly announcing new technologies long before they can be used might not prevent agents from choosing careers threatened by imminent creative destruction, because they might not pay attention to the initial public announcements. As an example, ChatGPT-1 was introduced to the public in 2018 (Marr 2023). The theory behind large language models and neural networks has been developing since at least the middle of the twentieth century (Christian 2020: ch. 1). Yet most people (myself included) only recently started paying attention to the capabilities of this technology when their students and colleagues started widely using ChatGPT-3. The technology is not as new as many think. There were warning signs, but most of us did not pay attention.

Third, knowing that a specific technology will be introduced in x years may not facilitate better planning even if people pay attention, because the actual impact of new technologies is extraordinarily difficult to predict. For example, one might reasonably have predicted that introducing ATM machines would reduce the demand for bank tellers; in fact, the number of bank tellers *increased* after ATM's were introduced (Ridley 2020: 291). What exact impact will ChatGPT and other AI tools have on the economy? The impact will be large, but it is hard to give specifics with any reasonable degree of confidence.

So much for the *ex ante* strategy. What about the *ex post* one? This strategy says the state can rescue the Stability Condition by ensuring individuals can complete their plans even after creative destruction has rendered their skills and labor obsolete by subsidizing their employment. Suppose industrial robots are invented that can substitute human labor on an assembly line. If they are adopted, many life plans will be upset. The state says to firms that might adopt them: if you continue employing human labor, we will pay the difference between the cost of human labor and the cost of robot labor for the next x years. After x years, the subsidy expires, so the

firm then adopts the industrial robots. Instead of trying to better anticipate creative destruction as the *ex ante* strategy does, the *ex post* strategy says the state should simply let creative destruction happen and respond after the fact. If your job is rendered obsolete by a new technology, then the state will pay to keep you employed until you either exit the labor force or are able to form and pursue a new plan.

There are two problems with the *ex post* strategy. The first is practical. To rescue the Stability Condition without suppressing creative destruction, the subsidy must only be deployed *temporarily*. The state must subsidize an employment sector long enough so individuals can complete their current plans or form new ones; after that, the subsidy must be removed so new technology can be adopted to increase productivity. Will the state apply the subsidy in this manner? Public choice theory tells us to consider policy interventions under the assumption that political actors (i.e., voters, politicians, bureaucrats, etc.) are knaves (Buchanan 1999; Kogelmann 2023). Under this assumption, there are reasons to think the subsidy will be deployed improperly.

Suppose a town is built around a factory that wants to substitute industrial robots for human labor. According to the *ex post* strategy, the state should subsidize employment in the factory for x years, after which it should then let the industrial robots take over. Even though the subsidy will allow current employees to complete their plans or form new ones, the town will still be devastated when the subsidy ends, because it loses its main source of employment. If no one works in the factory anymore then there will be no one to shop at local stores, eat at local restaurants, etc. Voters in the town will thus support politicians who want to extend the subsidy beyond x years, perhaps in perpetuity. Seeking election, politicians who represent the town will

support an elongated subsidy. If the subsidy never expires, however, creative destruction is delayed indefinitely; productivity gains are never captured.

If the undesirable consequences of an imperfectly implemented policy are great enough, it might be wise to give up on the policy altogether, *even if* the policy would be highly desirable assuming perfect implementation. In some cases, no policy is preferable to imperfectly implemented policy. As an example, Henry Sidgwick (2011: book 3, chap. 5) criticized the free trade arguments of Adam Smith and John Stuart Mill with his famous infant industries argument. In some cases, Sidgwick argued, protectionism makes economic sense. Yet, he did not think the state should actually pursue protectionist policies, because he knew the state would (for public choice reasons) do so imperfectly: “I do not think we can reasonably expect our actual Governments to be wise and strong enough to keep their protective interference within due limits” (Sidgwick 2011: 488). It is not unreasonable to reach a similar conclusion about the *ex post* strategy. Implemented perfectly, it might facilitate autonomy (although see the next criticism for doubts about this) in a manner that does not suppress creative destruction. Implemented imperfectly, creative destruction will be suppressed. All things considered, it is probably best if the state does not enter the business of subsidizing employment categories threatened by technological change.

The second problem with the *ex post* strategy is more philosophical. The worry is that subsidizing a career can change the nature of the career; this change can itself upset life plans. When many of us form plans, we select a career not only because it is enjoyable and we are suited to it, but also because we believe it contributes to some greater cause worthy of support. For instance, I chose to be a professor not only because I like reading, writing, and lecturing, but also because I believe in the transformative power of education. If it turns out that education has

little impact on students' lives, my career (in my eyes, at least) would be meaningless; my plan would be destroyed.

Something similar may happen when the state subsidizes careers only so people can continue working in them. Imagine someone who not only wants to work in manufacturing because they find it enjoyable and are suited to it, but also because they want to be an indispensable member of a team that produces an important good, like the steel that is used to make military ships and tanks. Yet, if the government must subsidize their employment because industrial robots outperform them, then they are no longer an indispensable member of the team that produces the important good. In fact, they are holding the team back by remaining in a job that a robot can perform more efficiently. Subsidization can steal the meaning from a job, and this loss of meaning can itself destroy life plans.

This section examined whether the state can thread the needle by regulating the process of technological change in a manner that secures both creative destruction and the conditions necessary for autonomy. I conclude that it cannot. Both the *ex ante* and *ex post* regulatory strategies are unable to ensure the Stability Condition can coexist with creative destruction. If we do not want to completely give up the ideal of autonomy or creative destruction, then we need a new approach.

5. Autonomy as Weathering Storms

Creative destruction and the ideal of autonomy conflict. Giving up either is undesirable, and there is no obvious way the state can regulate the process of technological change to render them consistent. What to do? In my view, we should rethink what it means to live an autonomous

life. In this section I propose and defend a novel conception of autonomy that is consistent with creative destruction.

Autonomy is often understood as synonymous with self-government (Wall 1998: 131; Bratman 2007: 195; Christman 2017: 5). Individuals are autonomous when they govern themselves. When autonomy is understood as synonymous with self-government, it is easy to see how its proponents are led to a conception of autonomy as forming, pursuing, and executing plans. For many, to govern is to plan; the state governs by setting goals and pursuing them through public policy (Oakeshott 1991: 26, 48; Hayek 2007: ch. 4). For instance, the state might decide it wants to increase GDP, reduce crime, or improve education. To pursue these goals, it then passes policies. To spur GDP growth the state might remove tariffs on trade, to reduce crime it might impose harsher penalties on convicted criminals, and to improve education it might increase funding for public schools. Understanding governing as an exercise in planning is most evident in Plato's (488a–489d) famous ship of state analogy. The state is like a ship with a specific destination that it must be steered toward; to be successful, it must empower those who navigate well.

There are some who want to break the tight association between governing and planning. Political philosopher Michael Oakeshott offers a variation of Plato's ship of state analogy. Here is how he understands what it means to govern:

In political activity, then, men sail a boundless and bottomless sea; there is neither harbour for shelter nor floor for anchorage, neither starting-place nor appointed destination. The enterprise is to keep afloat on an even keel; the sea is both friend and enemy; and the seamanship consists in using the resources of a traditional manner of behaviour in order to make a friend of every hostile occasion (Oakeshott 1991: 60).

According to Oakeshott, sailing the ship of state is not about steering it in a specific direction, for there is no end destination to reach. Rather, it is about keeping the ship afloat; the only goal is not to sink. On this view, governing is not an exercise in planning; it is about maintaining a satisfactory state of affairs by successfully responding to trials and tribulations. Applied to the individual, the autonomous agent is not one who forms, pursues, and executes life plans. Instead, the autonomous agent maintains a decent life for herself and her family by successfully navigating trials and tribulations with her skills, talents, and effort. To be autonomous is to weather storms by one's volition.

It will be helpful to name our two conceptions of autonomy. Call the more traditional conception that we have been working with throughout the paper thus far *Autonomy as Planning*, or AP for short. According to AP, to be autonomous agents must form, pursue, and successfully execute life plans. Call the new conception I am introducing *Autonomy as Weathering Storms*, or AWS for short. According to AWS, to be autonomous agents must maintain decent lives for themselves by successfully responding to trials and tribulations with their skills, talents, and effort. To further highlight the differences between these two conceptions of autonomy, let us consider some examples.

I already introduced an example of AP in §2 above, but it is worth repeating and adding a bit more detail. An agent, let us call her Althea, gives herself the following goals: she wants to be a lawyer, a mother, and a marathon runner. To pursue these goals, Althea comes up with a set of actions that are likely to achieve them. To become a lawyer, she must first go to law school; to become a mother, she must first find a suitable partner and achieve financial stability; before she runs her first marathon, she must start training and eating right. Althea performs the relevant

actions. After a decade of hard work, she achieves them. Because she forms, pursues, and successfully executes a life plan, Althea is autonomous according to AP.

Is Althea autonomous according to AWS? To answer that we need more detail; we need to know the severity of the challenges Althea faces in pursuit of her goals and how she responds to them. Perhaps it was not very difficult for Althea to achieve her goals. Suppose law school was easy because Althea is intellectually gifted, the marathon was a breeze because she is a natural athlete, and becoming a mother was no problem because she quickly found a suitable partner. Under these circumstances, Althea does not achieve an autonomous life according to AWS, because there were few obstacles to overcome. Of course, achieving her goals could have been very difficult for Althea. Suppose law school was a struggle that she barely made it through, her marathon training was beset by constant injury, and she spent a decade wading through failed relationships before she finally found the one. Under these circumstances, Althea is autonomous according to both AP and AWS. She sets, pursues, and achieves her goals and also must use her own volition to weather storms in the process.

Just as an agent can be autonomous according to AP, but not AWS, so too can she be autonomous according to AWS, but not AP. Let us consider a second example, this one paradigmatic of AWS. Bertha decides she wants to be a graphic artist; after getting settled in her career, she then wants to start a traditional nuclear family. She goes to art school in pursuit of her goal. While in art school, new AI tools are released that dramatically reduce the demand for graphic artists. Companies had to once hire graphic artists to design their websites and logos, but this can now be done by novices with AI image generators such as DALL-E and Midjourney. As a result, Bertha cannot find work after school. She struggles for many years. She works as an Uber driver, a waitress, and as a caretaker in a retirement home. She has a child with a man who

is not fit to be a father and soon exits her life, so she raises her daughter as a single mother while working full time. Eventually, Bertha observes that many professions in the medical field are untouched by AI. She attends nursing school in her late 30's. She graduates and gets a steady job working the nightshift at a hospital. Bertha does not love the work, but she makes enough money to send her daughter to college. In her mid-40's, Bertha meets a nice man that loves and supports her. They get married.

Things do not go according to plan for Bertha. The plan was to go to art school, become a graphic artist, then start a traditional nuclear family. None of that happened, so we should conclude that she did not live an autonomous life according to AP. Plans were set and pursued, but they fell apart. Yet, Bertha persisted through trials and tribulations and ended up creating a decent life for herself and her family with her skills, talents, and effort. She ends up with a good career, though not the one she wanted. She does get married, though far later in life than she anticipated. She raises a great kid and can support their higher education, even though she did not expect to do so alone. By her own volition, Bertha weathered the storms. She lived an autonomous life according to AWS.

We have seen that creative destruction conflicts with AP; it does not conflict with AWS. The issue with AP is that some of the conditions needed for agents to be autonomous in this sense—the Development and Stability Conditions specifically—are undermined by creative destruction. AWS will also have conditions that must be in place for agents to be autonomous in the sense it describes. Some of these conditions it will share with AP. For instance, to be autonomous in either sense will require sufficient mental capabilities, a robust psychological and motivational constitution, and so on. Some conditions will differ. To realize your autonomy in the sense described by AWS will require significant trials and tribulations; storms cannot be

weathered on placid waters. AP does not require this. Importantly, there is no reason to think autonomy in the sense described by AWS requires the Development and Stability Conditions. To navigate trials and tribulations, you do not need choice options suited to your talents and abilities. To overcome obstacles, you do not need stable choice options. (In fact, the exact opposite might be true. See the discussion of Humbolt below.) There is thus no reason to think creative destruction conflicts with AWS.

AP and AWS are two distinct conceptions of autonomy. Though there are some cases of overlap (when agents must overcome significant trials and tribulations to achieve their goals) they can also come apart (when agents achieve their goals without having to overcome significant trials and tribulations or when agents overcome significant trials and tribulations without achieving their goals). Should we prefer one conception to the other? I do not think so. I shall now argue that we ought to be indifferent between the two conceptions. They stand on equal ground. I do this by arguing that the considerations in defense of AP also support AWS. My claim: *if you think AP is attractive then you should also think AWS is attractive*. To see this, let us return to the three arguments in defense of AP covered in §2 above and inquire whether they also support AWS.

First, AP is an attractive normative ideal because it is bad to be overly dependent on the will of others, and agents who are autonomous according to this conception are not overly dependent on the will of others; they take charge of their affairs by forming plans and pursuing them. Agents who are autonomous according to AWS are also independent. According to the definition of AWS I have proposed, to be autonomous agents must successfully navigate life's trials and tribulations *with their skills, talents, and effort*. Being autonomous is not only about weathering life's storms, but weathering them *by one's own volition*. By definition, agents who

are autonomous in the sense I am proposing are not overly dependent on others. They must overcome obstacles with a significant degree of independence.

To see this, it is helpful to consider an example of someone weathering life's storms in a manner inconsistent with AWS. Suppose Bertha has very wealthy parents. After new AI tools eliminate her chosen profession, Bertha's wealthy parents tell her that, because of her bad luck, she can live off their wealth for the rest of her life, doing whatever she pleases. She never has to work again. If she accepts their offer there is a sense in which she weathers life's storms; she does not end up living a life of poverty and squalor. However, Bertha would not be autonomous according to AWS because she does not overcome the obstacles in front of her with her skills, talents, and effort. Instead of using her skills, talents, and effort, Bertha overcomes life's obstacles with her parents' money.

Note, to say Bertha must respond to trials and tribulations with her skills, talents, and effort does not imply that she cannot receive any help at all. After her career as a graphic artist does not work out, perhaps she receives some financial support from her parents, friends, or the state. That is fine. But eventually, she needs to chart a new path by her own volition. Similarly, the independence needed for AP is not *complete* independence. When forming and pursuing life plans you can ask for others' advice; you just cannot solely rely on their judgment. Your own judgment must enter the picture.

Second, AP is an attractive normative ideal because there is something valuable about making choices, and agents who are autonomous in this sense make choices: they choose what goals to pursue, they choose the means they shall use to pursue these goals, and so on. Agents who are autonomous according to AWS also make many choices. When creative destruction eliminates Bertha's preferred career, she must decide what to do next. When the father of her

future child walks out on her, she must decide whether to get an abortion, put the child up for adoption, or be a single mother. When she cannot make enough money to provide for her child as a waitress or Uber driver, she must choose a new career path.

In fact, those who are autonomous in the sense described by AWS probably (though not necessarily) make *more* meaningful choices than those who are autonomous in the sense described by AP. If you face frequent storms (i.e., are confronted with frequent challenges), then you will have to make many decisions about how to reorient the ship. If waters are placid and you are aiming at a specific port, then you may only have to make one choice about how to orient the ship. After that, you can sit back and let the winds do the rest.

Third, AP is an attractive normative ideal because living an autonomous life in this sense forces agents to develop their skills and capacities, which is a good thing. Forming and pursuing a life plan requires many different skills; these skills are refined when autonomous agents engage in these activities. As has been the case for the prior two justifications, this argument also supports AWS. To weather life's storms requires the development of many skills. After Bertha's plan of being a graphic artist is upset, she must learn several new skills: how to drive an Uber, be a waitress, and work as a caregiver in a retirement facility. When her child's father abandons her, she must learn to be a single parent. To provide a better life for her child, she must learn to be a nurse. To get married later in life, she must learn how to love again after being hurt. To navigate all the trials and tribulations life throws at her, Bertha develops many different skills and capabilities.

Like the prior argument, there is reason to think AWS satisfies the relevant normative criterion better than AP. Someone who gives herself a life plan and pursues it will surely develop many skills. Althea learns to be a lawyer, a parent, and a marathon runner. But someone who

must constantly revise and change direction in life will undoubtedly develop even more. Bertha learns to be a graphic artist, an Uber driver, a waitress, a retirement home caretaker, a single mother, a nurse, and more. In fact, Humbolt—recall, from §2 above, the German thinker who inspired Mill’s conception of autonomy—argues that to fully develop our capabilities requires constant variation. To fully develop our capabilities, Humbolt says agents need two indispensable conditions: freedom and “a variety of situations” (Humbolt 1993: 10). The latter component is needed because “even the most free and self-reliant of men is hindered in his development, when set in a monotonous situation” (Humbolt 1993: 10). Weathering storms is anything but monotonous.

I have just argued that the considerations in support of AP also support AWS. If you find the more traditional conception of autonomy attractive, then you should find my novel conception attractive as well. We ought to be indifferent between the two; both are valuable. To be autonomous in the sense described by AWS one needs rocky seas. We ought not upset placid waters just so agents can be autonomous in this sense. If you are lucky enough to sail on calm waters, then you should live an autonomous life in the sense described by AP. You should form, pursue, and execute a life plan. But if you end up living in times of creative destruction and the seas are rough, an autonomous life is still possible. You will not achieve autonomy by forming, pursuing, and executing a life plan. But you can be autonomous by weathering what Schumpeter (2008: 84) called the “gale of creative destruction.”

Works Cited

- Acemoglu, Daron and Pascual Restrepo. 2020. "Robots and Jobs: Evidence from US Labor Markets." *Journal of Political Economy* 128: 2188-2244.
- Acemoglu, Daron and James A. Robinson. 2012. *Why Nations Fail: Power, Prosperity, and Poverty*. New York: Crown Business.
- Aghion, Philippe, Céline Antonin, and Simon Bunel. 2021. *The Power of Creative Destruction: Economic Upheaval and the Wealth of Nations*. Cambridge: Harvard University Press.
- Autor, David H. 2014. "Skills, Education, and the Rise of Earnings Inequality Among the 'Other 99 Percent.'" *Science* 344: 834-851.
- Autor, David H. 2015. "Why Are There Still So Many Jobs? The History and Future of Workplace Automation." *Journal of Economic Perspectives* 29: 3-30.
- Autor, David H. 2017. "Paradox of Abundance: Automation Anxiety Returns." In *Performance and Progress: Essays on Capitalism, Business, and Society*, edited by Subramanian Rangan: 237-260. Oxford: Oxford University Press.
- Autor, David H., David Dorn, Gordon H. Hanson. 2013. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review* 103: 2121-2168.
- Autor, David H., Lawrence F. Katz, and Melissa S. Kearney. 2006. "The Polarization of the U.S. Labor Market." *American Economic Review* 96: 189-194.
- Bratman, Michael. 1987. *Intentions, Plans, and Practical Reason*. Cambridge: Harvard University Press.
- Bratman, Michael. 2007. *Structures of Agency: Essays*. Oxford: Oxford University Press.

- Brynjolfsson, Erik and Andrew McAfee. 2016. *The Second Machine Age: Work Progress, and Prosperity in a time of Brilliant Technologies*. New York: W.W. Norton & Company.
- Buchanan, James M. 1999. "Politics without Romance: A Sketch of Positive Public Choice Theory and Its Normative Implications." In *The Collected Works of James M. Buchanan, Vol. 1: The Logical Foundations of Constitutional Democracy*: 45-59. Indianapolis: Liberty Fund.
- Buss, Sarah and Andrea Westlund. 2018. "Personal Autonomy." *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. URL = <https://plato.stanford.edu/archives/spr2018/entries/personal-autonomy/>.
- Carroll, Jeffrey. 2023. "Must We Always Pursue Economic Growth?" *Utilitas* 36: 102-110.
- 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>.
- Christian, Brian. 2020. *The Alignment Problem: Machine Learning and Human Values*. New York W.W. Norton & Company.
- Christman, John. 2017. "Anti-Perfectionism and Autonomy in an Imperfect World: Comments on Joseph Raz's *The Morality of Freedom* 30 Years On." *Moral Philosophy and Politics* 4: 5-25.
- Christman, John. 2020. "Autonomy in Moral and Political Philosophy." *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. URL = <https://plato.stanford.edu/archives/fall2020/entries/autonomy-moral/>.

- Diamond, Arthur M., Jr. 2019. *Openness to Creative Destruction: Sustaining Innovative Dynamism*. Oxford: Oxford University Press.
- Friedman, Benjamin M. 2005. *The Moral Consequences of Economic Growth*. New York: Vintage Books.
- Goldin, Claudia and Lawrence Katz. 2008. *The Race Between Education and Technology*. Cambridge: Harvard University Press.
- Goodrich, Joanna. 2021. “How IBM’s Deep Blue Beat World Champion Chess Player Gary Kasparov.” *IEEE Spectrum* (25 January 2021). URL: <https://spectrum.ieee.org/how-ibms-deep-blue-beat-world-champion-chess-player-garry-kasparov>.
- 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>.
- Hayek, F.A. 2007. *The Road to Serfdom*. Chicago: University of Chicago Press.
- Humbolt, Wilhelm von. 1993. *The Limits of State Action*, edited by J.W. Burrow. Indianapolis: Liberty Fund.
- Juma, Calestous. 2019. *Innovation and its Enemies: Why People Resist New Technologies*. Oxford: 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>.

- Kochhar, Rakesh. 2023. "Which U.S. Workers Are More Exposed to AI on Their Jobs?" *Pew Research Center* (26 July 2023). URL: <https://www.pewresearch.org/social-trends/2023/07/26/which-u-s-workers-are-more-exposed-to-ai-on-their-jobs/>.
- Kogelmann, Brian. 2021. *Secret Government: The Pathologies of Publicity*. Cambridge: Cambridge University Press.
- Kogelmann, Brian. 2022. "We Must Always Pursue Economic Growth." *Utilitas* 34: 478-492.
- Kogelmann, Brian. 2023. "In Defense of Knavish Constitutions." *Public Choice* 196: 141-156.
- Koyama, Mark and Jared Rubin. 2022. *How the World Became Rich: Historical Origins of Economic Growth*. Cambridge: Polity Press.
- MacGillis, Alec. 2021. *Fulfillment: Winning and Losing in One-Click America*. New York: Farrar, Straus and Giroux.
- Marr, Bernard. 2023. "A Short History of ChatGPT: How We Got To Where We Are Today." *Forbes* (19 May 2023). URL: <https://www.forbes.com/sites/bernardmarr/2023/05/19/a-short-history-of-chatgpt-how-we-got-to-where-we-are-today/?sh=7571ada6674f>.
- Marx, Karl and Friedrich Engels. 1972. "Manifesto of the Communist Party," in *The Marx-Engels Reader*, edited by Robert C. Tucker: 469-500. New York: W.W. Norton & Company.
- Mill, John Stuart. 1978. *On Liberty*, edited by Elizabeth Rapaport. Indianapolis: Hackett Publishing.
- Mokyr, Joel. 1990. *The Lever of Riches: Technological Creativity and Economic Progress*. Oxford: Oxford University Press.

- Moller, Dan. 2011. "Wealth, Disability, and Happiness." *Philosophy & Public Affairs* 39: 177-206.
- Oakeshott, Michael. 1991. *Rationalism in Politics and Other Essays*. Indianapolis: Liberty Fund.
- Page, Scott E. 2017. *The Diversity Bonus: How Great Teams Pay Off in the Knowledge Economy*. Princeton: Princeton University Press.
- Plato. 2007. *The Republic*, translated by Desmond Lee with an introduction by Melissa Lane. New York: Penguin Classics.
- 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/>.
- Rawls, John. 1971. *A Theory of Justice*. Cambridge: Harvard University Press.
- Raz, Joseph. 1986. *The Morality of Freedom*. Oxford: Oxford University Press.
- Ricardo, David. 2004. *The Principles of Political Economy and Taxation*. Garden City: Dover Publications.
- Ridley, Matt. 2020. *How Innovation Works: And Why It Flourishes in Freedom*. New York: Harper Perennial.
- Schumpeter, Joseph A. 2008. *Capitalism, Socialism and Democracy*. New York: Harper Perennial.
- Schumpeter, Joseph A. 2017. *Essays on Entrepreneurs, Innovations, Business Cycles, and the Evolution of Capitalism*. London: Routledge.

- Service, Robert F. 2020. “‘The game has changed.’ AI trumps at protein folding.” *Science* 370: 11144-1145. URL: <https://www.science.org/doi/10.1126/science.370.6521.1144>.
- Sidgwick, Henry. 2011. *The Principles of Political Economy*. Cambridge: Cambridge University Press.
- Stevenson, Betsey and Justin Wolfers. 2013. “Subjective Well-Being and Income: Is There Any Evidence of Satiation?” *American Economic Review: Papers & Proceedings* 103: 598-604.
- Violante, Giovanni L. 2008. “Skill-Biased Technical Change.” In *The New Palgrave Dictionary of Economics*, edited by Steven N. Durlauf and Lawrence E. Blume: 1-6. London: Palgrave Macmillan. URL = https://doi.org/10.1057/978-1-349-95121-5_2388-1.
- Vogt, Thomas. 2016. “How Fast Should We Innovate?” *Journal of Responsible Innovation* 3: 255-259.
- Wall, Steven. 1998. *Liberalism, Perfectionism and Restraint*. Cambridge: Cambridge University Press.
- Wall, Steven. 2018. “Planning, Freedom, and the Rule of Law.” In *The Oxford Handbook of Freedom*, edited by David Schmidtz and Carmen Pavel: 283-299. Oxford: Oxford 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>.

Woodhouse, Edward J. 2016. "Slowing the Pace of Technological Change?" *Journal of Responsible Innovation* 3: 266-273.