

Navigating AI's Impact on Labor: Challenges, Scenarios, and Policy Pathways*

Tillman Schenk[†]

Arturs Kanepajs[‡]

David Conrad[§]

November 25, 2024

Abstract

The rapid development of Artificial Intelligence (AI) is set to transform labor markets, presenting both significant opportunities and challenges. However, there is considerable disagreement on precisely how AI will impact labor demand. This paper presents a structured framework to navigate this debate by categorizing potential outcomes along two key dimensions: AI's net task displacement effect and its aggregate productivity effect. Based on these dimensions, we identify four scenarios for AI's impact on labor demand, ranging from minimal disruption to transformative changes. To address these uncertainties, we propose a two-pronged policy approach: steering policies to guide technological development toward favorable labor market outcomes, and adaptation policies to mitigate AI's adverse effects.

Keywords: Artificial Intelligence, Labor Markets, Policy Responses, Productivity, Automation

*This research was conducted as part of the AI Safety Camp, an initiative dedicated to fostering collaboration and advancing research in the field of AI safety. The authors are deeply grateful for the coordinative and administrative support provided by the AI Safety Camp, without which this project would not have been possible. Arturs Kanepajs also gratefully acknowledges the financial stipend he received from the AI Safety Camp.

[†]Independent, tillmanschenk@gmail.com

[‡]Independent, akanepajs@gmail.com

[§]Independent, davidconrad93@gmail.com

Executive Summary

The rapid development of Artificial Intelligence (AI) is set to transform labor markets, presenting both significant opportunities and challenges. However, analysts broadly disagree on precisely *how* AI will affect workers, underscoring the uncertainty of its impacts and complicating the design of effective policies. This paper provides policymakers with an accessible overview of the key unknowns, explores scenarios for AI's potential effects on the workforce, and outlines appropriate policy responses.

Echoing insights from modern economic research, the starting point of our analysis is that the impact of AI on labor demand depends on two crucial factors. Focusing on these dimensions provides a structured framework to navigate the broader debate on AI's labor market effects. Firstly, the *net task displacement effect* refers to the extent to which AI automates tasks previously performed by human labor, reducing overall labor demand. However, this concept also accounts for the new tasks AI generates, which may require human workers. Thus, the *net effect* captures the combined impact of tasks displaced by automation and those newly created by AI. Secondly, the *aggregate productivity effect* captures the efficiency gains brought about by AI. These generally lower prices, increase consumer wealth, and, in turn, elevate labor demand.

We subsequently argue that these two dimensions give rise to four potential scenarios for how AI might impact labor demand. At the same time, they provide a framework for organizing the diverse positions analysts hold regarding AI's effects on the workforce:

1. In a **low-effect AI** scenario, AI automates only a limited number of tasks and fails to significantly boost productivity growth. Consequently, its overall impact on labor demand remains minimal. Economists who consider this scenario plausible often highlight the recent slowdown in productivity growth despite increased investments in R&D.
2. **So-so automation** describes a scenario characterized by significant labor displacement paired with only modest productivity gains, leading to substantial losses for workers. Analysts who view this scenario as likely frequently argue that AI is likely to extend patterns observed since the onset of the computer revolution.
3. Under a **productive synergy** scenario, few workers are displaced while efficiency gains are substantial, leading to a positive impact on overall labor demand. Economists who deem this scenario credible commonly draw analogies to earlier general-purpose technologies, which significantly boosted productivity without causing widespread unemployment.
4. **Technological supremacy** describes a scenario characterized by substantial productivity

gains and significant displacement effects, resulting in a transformative and largely unpredictable impact on labor markets. The so-called 'Singularity' represents an extreme version of this scenario. Scholars who regard this scenario as plausible typically highlight advanced AI's presumed capability for self-improvement.

Figure 1 illustrates the framework by mapping the four scenarios onto a two-dimensional space defined by AI's net task displacement and aggregate productivity effects.

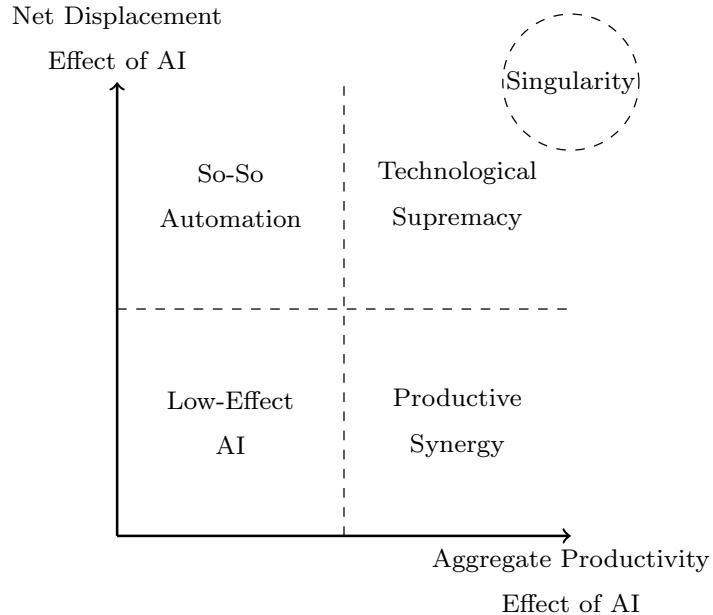


Figure 1: A framework illustrating the four potential scenarios of AI's impact on labor demand, categorized by the net displacement effect and aggregate productivity effect of AI.

Motivated by the wide range of potential scenarios for how AI might impact labor demand—and the uncertainty surrounding its trajectory—we argue that policymakers should adopt two complementary sets of policies to ensure that AI's adverse effects on labor demand remain minimal:

- **Steering policies** aim to guide the direction of technological progress, ensuring that AI developments lead to favorable labor market outcomes. The trajectory of technological development and diffusion is shaped by the actions of public and private actors, who respond to incentives, adhere to standards, and comply with regulations. Steering policies include both efforts to influence what AI systems are developed—such as government-funded R&D or tax policies that incentivize the use of labor in production—and strategies to guide how these systems are integrated into economies, including tools like regulatory sandboxes.

While steering policies can play a critical role, they are often relatively blunt instruments with implications that reach beyond the domain of AI. Moreover, critics argue that the downstream

applications of specific innovations are inherently difficult to predict. For these reasons, we contend that steering policies should be complemented by a second set of policies:

- **Adaptation policies** aim to address the adverse effects of AI on employment and economic stability without directly shaping the trajectory of technological progress. Crucially, the relevance and design of specific adaptation measures depend on which scenario ultimately materializes. For example, while a universal basic income may be well-suited to a technological supremacy scenario, it would be far less necessary under a low-effect AI scenario.

Ultimately, only time will reveal how AI will shape the workforce. While existing empirical evidence remains limited and must be interpreted with caution, this does not mean policymakers can afford to remain idle. Although some scenarios and policies may appear futuristic or premature, the rapid pace of technological advancement makes it imperative to start preparing for the changes ahead.

1 Introduction

The advent of transformative Artificial Intelligence (AI) technologies is poised to significantly disrupt labor markets. This paper seeks to examine these disruptions and propose strategies to mitigate potential adverse effects. Central to our analysis is a framework that categorizes and differentiates predictions regarding AI’s impact on labor demand. Drawing on the work of Acemoglu and Restrepo (2018a, 2018b, 2019), we argue that AI’s effects on labor demand are shaped by two critical factors: (i) the net extent to which AI technologies automate or replace human tasks, and (ii) the aggregate productivity effects of AI deployment. Using this two-dimensional framework, we identify and explore four distinct scenarios for how AI could influence labor demand. For each scenario, we detail its defining characteristics and evaluate the economic arguments supporting its plausibility. To effectively address the broader implications of AI on the workforce, we then argue that policymakers should adopt a two-pronged policy approach. The first approach focuses on steering policies that aim to guide technological progress toward sustaining robust labor demand. This includes encouraging labor-augmenting innovations over automation and maximizing productivity gains. In contrast, the second approach comprises adaptation policies which are not meant to directly influence the trajectory of technological development but aim to buffer the workforce against potential adverse impacts.

The structure of this paper is as follows: First we outline the theoretical underpinnings of our framework. We then discuss each scenario as well as the key arguments brought forward for each scenario. In the second part of our paper, we then turn towards a discussion of the implications of our findings for policymakers.

2 Framework

Building on the canonical model offered by Acemoglu and Restrepo, we argue that AI’s impact on labor demand may depend on two factors: Firstly, there is, what we call the net task displacement effect of AI. This refers to the adverse effect AI has on labor demand by allowing for the substitution of workers with machines, less the positive effect AI has on labor demand by creating new labor-intensive tasks. Secondly, there is (ii) the aggregate productivity effect of AI. By making the production process more efficient, AI allows firms to reduce the price they charge for their products, which makes consumers richer and stimulates the demand for labor as an input to production.

To understand why labor displacement negatively affects labor demand, imagine production as performing a set of well-defined tasks. Acemoglu and Restrepo (2018a) give the example of textile production requiring the production of fiber, the production of yarn from fiber, the production of

2 FRAMEWORK

the relevant fabric from the yarn, pretreatment, dyeing and printing, as well as various auxiliary tasks. Some of these tasks are more cost-efficiently performed by machines, while others can only be performed by human labor. That is, labor and capital have comparative advantages in different tasks, meaning that their relative productivity varies across tasks. ‘Automation’ is then conceptualized as an expansion of the set of tasks which are cost-efficiently performed by machines. This process of substituting machines for labor in the production process has a negative effect on the demand for labor, as firms require humans to perform fewer tasks. This is what Acemoglu and Restrepo have labeled the ‘displacement effect’ of automation.

Notably, what we call the net displacement effect of AI on labor demand is also determined by the extent to which progress in AI creates new tasks in which human labor has a comparative advantage relative to machines. As documented by Autor et al. (2024), the last decades have seen a significant expansion of job titles, arguably allowed for by technological change. For example, the digital revolution has created a whole host of new occupations, e.g. those working in the information technology (IT) industry. This is what Acemoglu and Restrepo have labeled the ‘reinstatement effect’ of technical change, which increases the demand for labor as an input to production. Ultimately, when referring to the net task displacement effect of AI, we mean the adverse effect of displacement on labor demand less the positive reinstatement effect of AI on labor demand.

On the other hand, increases in productivity have a positive effect on the demand for labor in the model of Acemoglu and Restrepo. Their model contains multiple drivers of aggregate productivity growth, all of which we summarize as the ‘aggregate productivity effect of AI’. Firstly, due to advances in AI, machines might gain a comparative advantage vis-a-vis humans in certain tasks due to advances in AI, and firms will subsequently replace humans with machines within those tasks. As firms replace humans with more cost-efficient machines, production becomes more efficient. This is what Acemoglu and Restrepo call the ‘productivity effect’ of technological change on AI. As an additional boost to the productivity effect, we also include the ‘capital accumulation effect’ in this dimension: In the short-run, increased demand for capital due to automation meets a fixed supply of capital, driving up the rental rate for capital. In the medium-run, the supply of capital adjusts to the increased rental rate, which reduces the rental rate and increases the real wage, thereby giving the productivity effect an additional boost. Thirdly, technical change might also increase the productivity of tasks which are already being performed by machines. This is what the authors call ‘deepening of automation’. In either case, productivity increases allow firms to reduce the prices they charge for their products. Other things being equal, households become effectively richer, which leads to increased demand for goods and services, either within the sector experiencing

2 FRAMEWORK

the productivity increases or in other sectors of the economy. Faced with increased demand, firms expand production, thereby increasing demand for all of their factor inputs, including labor.

To build intuition, consider Figure 2 below which has the net displacement effect of AI on the x-axis and the aggregate productivity effect of AI on the y-axis. The 45 degree line starting at the origin represents constant labor demand. In all scenarios above this line, AI has an adverse effect on the demand on labor, as the net displacement effect outweighs the aggregate productivity effect of AI. In contrast, in scenarios below the line the aggregate productivity effect of AI is stronger than the net displacement effect, which means that aggregate labor demand is poised to decrease. Importantly, this does not mean that the scenarios above or below the line are equally likely to occur.

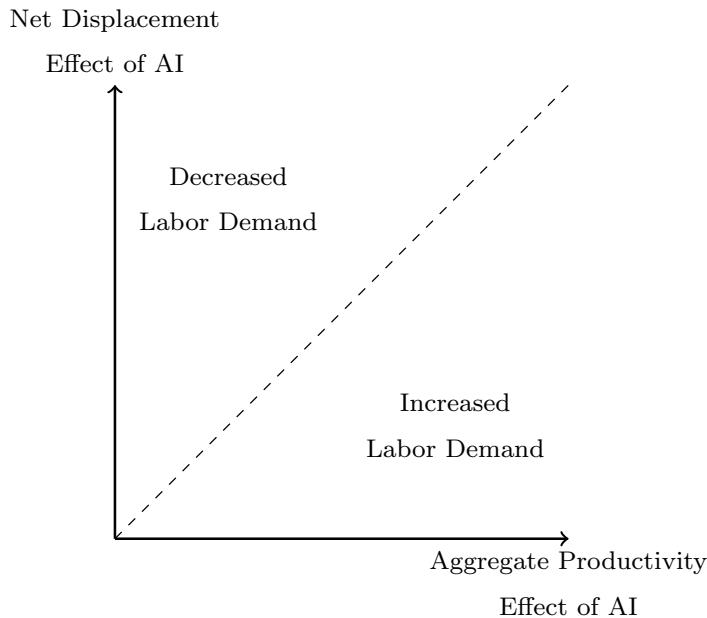


Figure 2: Labor demand outcomes based on the interaction of AI's net displacement and aggregate productivity effects. The dashed line divides scenarios of decreased and increased labor demand, highlighting the interplay between these two dimensions.

Furthermore, shifts in labor demand could manifest in changes in wages, changes in employment, or, what is most likely, a combination of both. How shifts in labor demand will ultimately affect our labor markets depends on a range of factors, such as the elasticities of labor demand and supply and the labor market institutions which are in place. One should also note that, even if AI has an adverse effect on labor demand as some economists believe, other effects which are independent of AI could outweigh these adverse effects. As an example, Manyika et al. (2017) argue that demographic changes and the resulting shortage in labor supply, and not emerging technologies, will dominate labor markets in the foreseeable future.

As visualized in Figure 1, we distinguish four types of scenarios for AI’s impact on labor markets, each corresponding to different economists’ predictions. In a low-effect AI scenario, AI displaces few workers and has minimal impact on productivity, thus having a negligible effect on labor demand. In a so-so automation scenario, AI substitutes machines for workers in many tasks, but these machines are only slightly more productive than humans, leading to decreased labor demand and potential reductions in employment or wages. A productive synergy scenario sees AI automating work in limited industries while significantly augmenting human labor, increasing productivity and creating new tasks, which boosts aggregate labor demand, wages, and employment. In a technological supremacy scenario, AI transforms economies by automating a wide range of tasks with highly capable and productive machines, making the effect on labor demand ambiguous and dependent on the balance between AI’s displacement and productivity effects. The ”Singularity” represents an extreme variant of this scenario.

In what follows, we discuss the characteristics of each of the four scenarios in greater detail. We also highlight the arguments brought forward in the literature on why each of the scenarios may be considered plausible.

2.1 Low-Effect AI

In a low-effect AI scenario both the aggregate productivity effect and the net displacement effect of AI remain minimal. Consequently, AI has little impact on labor demand. This scenario is most plausible if transformative AI turns out to be more hype than reality, with limited economic impact.

One argument brought forward in favor of this hypothesis is that aggregate productivity growth in recent decades has been relatively slow despite fast advancements in (AI-)technology e.g. measured by the number of patent issues and the spread of digital technologies (Bloom et al. 2020). This phenomenon has been labelled the productivity paradox (Andrews, Criscuolo, and Gal 2016). Such lagging productivity statistics throw doubt on the prediction that AI might transform our societies altogether. AI-skeptics have interpreted these findings as being the result of the inherent limited capabilities of AI technologies. Those who endorse such skeptical positions often endorse arguments such as Polanyi’s paradox. Polanyi had noted that computers face challenges in performing tasks relying on tacit knowledge (i.e. tradition, intuition, inherited practices, implied values, and pre-judgments), such as organizing a closet. Michael Polanyi’s observation is that “We can know more than we can tell” (Polanyi 2009). More precisely, there are many tasks which people understand intuitively how to perform, but cannot elicit the rules or procedures they follow. Notably, there is a debate among computer scientists about whether AI (and machine learning in particular) can in the future provide a solution to Polanyi’s paradox, by applying statistics and inductive reasoning

to supply best-guess answers in cases where formal procedural rules are unknown (Autor 2014).

As a result, Gordon (2018) argues that much of the impact of the digital revolution has been seen already and any further innovations are more likely to be marginal improvements of past technologies rather than technological breakthroughs. That is, while IT has truly revolutionized our economies and labor markets, this technology has come to maturity. AI, in turn, is not a new invention but simply a continuation of the digital revolution and will not lead to major distortions in the labor market. Gordon does expect some job displacement effects resulting from AI, but these will happen slowly and at a steady pace. Complementarities resulting from AI might materialize, but will produce only modest increases in productivity.

A different reason why AI might not turn out to be an impactful technology is that diffusion might remain limited. As an example, Calvino and Criscuolo (2022) have argued that productivity growth due to AI might be limited as most benefits from AI are reaped by a few “superstar” firms: While firms at the global frontier of productivity have continued to increase their productivity steadily, the rest of the business population has not kept pace. This is especially true in the most digital- and knowledge-intensive sectors, to whom AI-labs belong. Other institutional factors which might prevent the diffusion of AI technologies include labor market institutions, (e.g. unions protecting jobs) and regulation (e.g. prohibiting self-driving cars due to safety concerns), organizational frictions as well as constraints to capital accumulation (e.g. bottlenecks in the supply-chain for chips). However, if AI-capabilities turn out to be truly revolutionary, one can arguably question the extent to which these factors can prevent diffusion in the long run.

2.2 So-So Automation

In a scenario of so-so automation, AI advancements enable firms to replace humans in many tasks, leading to reduced labor demand. At the same time, the AI technologies which replace human workers are only marginally more productive, implying that the aggregate productivity effect remains weak.

Falling labor demand could result in rising unemployment, underemployment, lower labor force participation rates, but could also entail lower or stagnating wages. Autor (2022) and Johnson and Acemoglu (2023) have argued that, in the immediate future, AI is more likely to result in lower wages for workers. In their view, AI is not advancing fast enough to increase unemployment significantly. In either case, labor’s share of national income is determined to fall, which would entail an increasing concentration of power and wealth as well as a loss of economic and political bargaining power. With increased concentration comes the peril of being trapped in an equilibrium where workers cannot improve their outcomes on their own, an outcome Brynjolfsson (2023) called

the “Turing Trap”. Next to inequality, even outright poverty might become a concern. In addition, low productivity growth means that the amount of resources which can be freed up for redistribution are limited.

Johnson and Acemoglu (2023) name two reasons for the push towards automation: Firstly, the vision dominating the current developments in AI research is one biased towards automation. That is, knowingly or not, AI-researchers, engineers, and entrepreneurs are all biased and incentivized towards creating an intelligence which matches, and thereby replaces, human intelligence. Brynjolfsson (2023) identifies the erstwhile Turing Test as creating an automation mindset for AI research at the expense of potential augmentation paths. Alan Turing had famously proposed a test for whether a machine was intelligent: Can that machine imitate a human so well that its answers are indistinguishable from a human’s? Secondly, introducing machines which augment labor is not attractive to organizations and companies which are intent on cutting costs. With respect to low productivity effects, Johnson and Acemoglu (2023) argue that humans are already very good at what they do. AI-based automation is not going to have impressive results when simply replacing humans. As an example, humans possess a high degree of social and situational intelligence fully absent in modern AI technologies.

An additional argument speaking in favor of the hypothesis is that this is arguably the trend we have experienced throughout the digital age. Acemoglu and Restrepo (2019) document that, while automation and task reinstatement were roughly in balance between 1950 and 1987, automation subsequently outpaced task reinstatement during 1987-2017. I.e. the overall trend in the digital era seems to be one in which fewer tasks emerge than are being automated. In addition, economists have noted that task-automation is biased towards routine tasks, implying that computers substitute for low-skilled workers while augmenting highly skilled workers (Autor, Levy, and Murnane 2003). In general, economists who argue for the plausibility of the so-so automation scenario largely believe that AI will continue these trends.

2.3 Productive Synergy

‘Productive synergy’ describes a scenario in which the net displacement effects of AI remain limited, but where the productivity growth resulting from AI adoption is significant. In such a scenario, aggregate labor demand is being boosted, ensuring growth in wages and employment in most industries.

For an economy to experience a significant boost in productivity requires a fundamental re-thinking of the organization of production. Firms must fundamentally transform business processes (Brynjolfsson, Rock, and Syverson 2021). This includes the automation of a lot of tasks and occu-

2 FRAMEWORK

pations, but also the creation of many new and productive tasks in which humans continue to have an advantage. Yet, while such structural transformations have overall positive long-term effects on the economy as a whole, they also create laggards and losers. Thus, although this scenario has positive aspects, structural change remains an important challenge. We should also note here that such scenarios are still compatible with a falling labor share (i.e. share of labor in national income). That is, while complementarities ensure that wage growth is strong, capital income could grow even faster.

A strong argument speaking in favor of this scenario is that previous general-purpose technologies (GPTs) have arguably impacted labor markets in a similar way. The concept of a GPT has been popularized by Bresnahan and Trajtenberg (1995). It labels technologies which have the potential of application across a broad variety of sectors, the ability to improve over time and to generate complementary innovation. Such technological revolutions usually came along with boosted output and large welfare gains (Brynjolfsson, Rock, and Syverson 2021). While they also displaced a wide range of tasks previously performed by humans, they also created new labor-intensive tasks. The emergence of such new labor-intensive tasks has been decisive in ensuring long-term real wage growth.

A. Agrawal, Gans, and Goldfarb (2019) judge AI to qualify as a GPT due to its ability to produce predictions, which are an essential input into any decision-making task. In turn, decision-making is required across a diverse range of occupations, which means that AI transforms work processes across a wide range of sectors. Still, prediction is only one input required for effective decision-making. Decision-making also requires the collection and organization of data, the ability to take action based on a decision, and the judgment to evaluate the payoffs associated with different outcomes. As a result, AI is less of a substitute for, and more of a complement to human work. As a consequence, A. K. Agrawal, Gans, and Goldfarb (2023) reject the Brynjolfsson (2023) point that AI-development is biased towards automation.

Hopes that AI could augment workers have particularly focussed on low-skilled workers. Although his text is arguably “not a forecast but a claim about what is attainable”, Autor et al. (2024) has argued that AI will alter the nature of human expertise. Expertise or skill, understood as the knowledge or competency required to perform a particular task, is arguably the primary source of labor’s value in developed countries. While computerization has generally increased the premium paid to skilled-labor, AI might reverse that trend by ‘upskilling’ low-skilled workers. That is, as AI weaves together the information required for decision-making, it might complement a large set of workers in performing higher-stakes decisions, thereby boosting workers’ productivity. Autor admits that AI will certainly automate some existing work. But it will further instantiate new human

capabilities, new goods and services that create demand for expertise we have yet to foresee. An article in

The Economist (2024) applies this line of reasoning to the divide between the global South and North. The article argues that AI could raise productivity and shrink gaps in human capital in emerging markets. While currently, there are not nearly enough skilled teachers, doctors, engineers or managers (i.e. expertise) in many countries of the Global South, AI could ease this shortfall by augmenting the existing workforce. This echoes well the idea discussed above that AI might help to ‘upskill’ left-behind workers. Another article goes so far as to declare that a “Blue-Collar Bonanza” is underway (The Economist 2023).

2.4 Technological Supremacy

Under a technological supremacy scenario, AI takes over a large share of tasks currently performed by humans, which decreases aggregate demand for labor. At the same time, AI-technologies significantly outperform the human workers it replaced. As a result, productivity increases are significant, which makes economies expand and increases the overall demand for labor. A so-called “Singularity” would be an extreme form of such a scenario.

Notice that the resulting overall effect of AI on labor demand is somewhat ambiguous. As workers are being displaced with highly productive machines, capital-owners are likely to experience spikes in their income. The effect of AI on labor income is ambiguous. If the productivity shock resulting from our technological progress is comparably large it might undo the AI-driven labor displacement effect. However, Acemoglu and Restrepo (2018a) believe such scenarios to be unlikely. While productivity effects are considered to be important, such dynamics depend strongly on the elasticity of substitution and are in general not sufficient to retain a stable labor share in national income. In either case, as capital is distributed much more unequally within societies than labor, the labor share of national income drops significantly. Yet, the abundance created by productivity growth allows for more costly redistribution schemes, which we discuss in the next section.

What speaks in favor of such a scenario being plausible? Some economists have argued that the distinguishing feature of AI compared to other general-purpose technologies is its ability to self-improve and thereby expand the set of tasks which can be automated. This speaks in favor of large progress in productivity and automation. In its most extreme form, AI’s ability to self-improve could lead to a so-called Singularity, a point in time at which machine intelligence exceeds human intelligence (Bostrom 2009; Good 1966) and economic growth accelerates at ”an ever-accelerating pace of improvements cascades through the economy”. Labor is becoming redundant under such a scenario. As a result, Brynjolfsson (2023) argued that AI might be “the most general of all

3 POLICY IMPLICATIONS

general-purpose technologies: after all, if we can solve the puzzle of intelligence, it would help solve many of the other problems in the world.”

Nordhaus (2021) observes that proponents of Singularity theory are more frequently computer scientists than economists. Among the few economists who do align with Singularity theory, many advocate for a ”soft version” of the concept. One of the most prominent figures in this regard is Erik Brynjolfsson. Brynjolfsson and McAfee (2017) emphasize AI’s capability to self-improve and believe in the transformational power of the technology. Still, they point to the technology’s limitations too. In their view, the most important limitation of AI systems is that they are devices for answering questions, not for posing them. This means that entrepreneurs, innovators, scientists, creators and others who figure out what territory to explore next might remain essential. In the same manner, Baily, Brynjolfsson, and Korinek (2023) argue that AI will not only make workers more productive but also increase the rate of innovation, laying the foundation for a significant acceleration in economic growth. Others have pointed towards the interactions of task-automation and aggregate productivity. As Trammell and Korinek (2023) argue, large leaps in aggregate productivity become more likely when task-automation is high. Neoclassical economic models suggest that output growth accelerates significantly once AI has automatized all tasks performed in an economy, given that the scarcity of labor is no longer a constraint on output.

To explain the productivity paradox, Brynjolfsson, Rock, and Syverson (2021) as well as Brynjolfsson, Mitchell, and Rock (2018) attribute the productivity paradox primarily to lags in AI implementation and restructuring (i.e. the productivity J-curve). In their view, when a new general purpose technology is introduced, productivity growth is arguably underestimated in national accounts. That is because introducing these new technologies requires, next to measurable investments in physical equipment, “large intangible investments and a fundamental rethinking of the organization of production itself” (Brynjolfsson, Rock, and Syverson 2021). Later, when the benefits of intangible investments are harvested, productivity growth is overestimated. Notably, these findings are in line with the findings of economic historians who have studied earlier general purpose technologies (e.g. David (1990)).

3 Policy Implications

We argue that, to effectively address the broader implications of AI on the workforce, policymakers should consider two sets of policies. The first set focuses on ’steering’ technological progress towards ensuring strong labor demand. This includes promoting labor augmentation over automation and achieving high productivity gains. The second set of ’adaptation’-policies does not attempt to directly influence the trajectory of technological development but aims to cushion the poten-

3 POLICY IMPLICATIONS

tial negative impacts on the workforce. This comprehensive approach is broadly in line with the framework provided in Bernardi et al. (2024).

3.1 Steering Progress

It is a recurring theme in the literature that technological development is neither predetermined nor inevitable, and that policy-makers can steer technological change towards more favorable labor market outcomes (e.g. Korinek and Stiglitz (2020), Johnson and Acemoglu (2023), Autor et al. (2024)). Technological development depends on the actions of public and private actors, which, in turn, respond to incentives, follow standards and need to abide by regulations. This includes both policies for changing what AI systems are being developed, as well as policies for influencing how they diffuse within economies.

In suggesting policies for steering technological progress, we will assume that it is in policy-makers interest to strengthen the demand for labor. This is in line with current standard labor market policy views: Productive employment provides income for the employed, the employer and the state in the form of taxes. Employment also provides non-financial benefits such as meaning and purpose (Susskind 2023). Current institutions are not ready to provide permanent sources of income and meaning, that would fully compensate the losses, for redundant workers. Furthermore, fully automating labor disposes workers of political power, which in turn can result in deterioration of democracy and extreme economic inequality (Johnson and Acemoglu 2023).

One straightforward way governments can strengthen the demand for labor is by financially supporting the R&D and adoption of AI-technologies that lead to the augmentation of jobs, especially those associated with high job satisfaction and meaning; and discourage the automation of such jobs (Korinek and Stiglitz 2020). This also includes funding academics and auditors who aim to measure and predict AI capabilities, in order to better understand how AI technologies will ultimately affect workers. Yet, while this is an appealing approach in theory, the general purpose nature of many AI systems means that it may be impossible to determine in advance the extent to which they will automate jobs. Further research is needed on how to measure and predict AI capabilities and corresponding risks with respect to work displacement and augmentation of workers (Anwar et al. 2024).

Still, a significant portion of research on AI is performed by private companies. A subtle way of influencing private sector R&D is through tax incentives and subsidies that affect the deployment and use of technologies. Korinek and Stiglitz (2020) and Johnson and Acemoglu (2023) have argued that, by decreasing taxes on labor vis-a-vis capital, AI-automation, as a way of making production more capital-intensive, becomes less attractive to firms. This is arguably even more so since labor

3 POLICY IMPLICATIONS

is the most highly-taxed factor in our economic system. However, to avoid a race to the bottom, we note that international cooperation is paramount when it comes to taxation (Piketty, Saez, and Zucman 2018). Yet, such tax reforms would arguably not come without side-effects. Some have argued that increasing the tax rate on capital income would arguably discourage savings and investment relative to current law. With increased consumption and lower investments, we would end up on a lower growth trajectory.

Strengthening unions and work councils may also affect the trajectory of technological change (Korinek and Stiglitz 2020). Johnson and Acemoglu (2023) highlight that labor unions have at times been successful in countering excessive automation. Facilitating involvement of workers keeps important company choices, e.g. about the extent of automation, in their hands. This might be particularly important because workers are the best to judge the value of intangible benefits of work, such as purpose and meaning (Johnson and Acemoglu 2023; Autor, Dorn, and Hanson 2016; Belfield and Hua, n.d.; Korinek and Stiglitz 2020).

In addition, regulatory sandboxes may help to learn about the impacts of new technologies while limiting potential risks by initially deploying new technologies only in a subset of firms or industries (OECD 2023). Firms or industries where the social damages exceed a certain threshold could then be prevented from adopting the technology for some time (Acemoglu and Lensman 2024). This is similar to providing a gradual or staged access to AI systems, which allows balancing control of the risks with learning about the risks (Solaiman 2023).

Market concentration has an ambiguous effect on labor demand. The most capable AI foundation models have a natural tendency towards increased market concentration. This is due to their high initial investment costs coupled with relatively low operational expenses once those investments have been made. This concentration, in turn, can amplify incentives for collusion and misrepresentation, potentially harming workers. For example, mandatory reports on the labor market impacts of certain AI innovations might be manipulated to present a more favorable outlook (Sastry et al. 2024; Vipra and Korinek 2023). In addition, in monopolistic markets, tech giants may amass political and economic power and may ultimately disregard the interests of stakeholders (e.g. by pushing for excessive automation). This is especially so given that profits increase in market concentration, implying that more resources are disposable for lobbying purposes. On the other hand, controlling diffusion might be much more challenging when there are a multitude of AI labs. We also note that open source as a commonly suggested strategy towards tackling monopolistic competition (Brynjolfsson and Unger 2023) runs contrary to efforts of controlling diffusion. We note more generally that it is important to balance the benefits of open sourcing with negative externalities, like safety risks, that more competition in general and open source development in

3 POLICY IMPLICATIONS

particular can pose (Vipra and Korinek 2023; Seger et al. 2023).

However, these ‘steering’-policies have important limitations. One important issue is that many of the policies discussed above are fairly blunt instruments and will arguably have implications which go beyond the present context. For example, while equalizing the tax rate on capital and labor income may incentivise innovations that increase labor demand, it will likely also bring about a wide array of adverse effects on the saving and investment behavior of economic agents. Moreover, prohibiting open-sourcing AI technologies may also stifle innovation and lead to a concentration of market and decision-making power. We also note that there is an ongoing discussion on the extent to which technological progress can indeed be steered in the first place. Skeptics argue that it is hard to foresee the downstream use-case of specific innovations, making it hard for developers to guide their efforts into a specific direction. As a result, these policies can hardly be seen as failsafe.

3.2 Adaptation Policies

In light of these limitations, we argue that it is essential to complement ‘steering’ policies with a second set of ‘adaptation’ policies. These policies aim to mitigate the adverse effects of AI on employment and economic stability without actively altering the trajectory of technological progress. Yet, the high degree of uncertainty surrounding AI’s impact on labor markets, along with the diverse opinions among economists, necessitates a flexible and multifaceted approach. Following Korinek (2023), we employ a ‘portfolio-approach’ to structure our policy recommendations to better address the unique challenges presented by each potential scenario.

We begin by noting that, if we end up in a low-effect scenario, where the labor market impact of AI remains minimal, there should not be a major policy response.

In contrast, the so-so automation scenario necessitates a substantial policy response due to the potential for a falling labor share of national income, which, if left unchecked, could lead to increased inequality, poverty, and a concentration of political and economic power. The lack of productivity growth in such a scenario would also limit the resources available for remedies such as income and wealth redistribution. To counteract these imbalances, significant redistributive measures would be essential. For example, policymakers could increase taxes on higher income brackets and capital gains, using the additional revenue for redistribution and potentially reducing the tax rate on labor income, which also has a steering effect as discussed above (Ahn 2024). Given the resource constraints, public employment programs in infrastructure, green energy, education, and healthcare might be more cost-effective than universal basic income schemes. Moreover, job guarantee schemes, as highlighted by Susskind (2023), could be particularly beneficial for those who find work to be an important and non-substitutable source of meaning. Additionally, job-sharing initiatives could

3 POLICY IMPLICATIONS

help distribute available work more evenly across the workforce, maintaining employment levels, improving work-life balance, and enhancing productivity through better employee morale.

Although aggregate labor demand remains strong under the productive-synergy scenario brings about its own unique challenges. In particular, for an economy to experience a significant boost in productivity, a significant restructuring of production processes seems unavoidable. While such structural transformations have overall positive long-term effects on the economy as a whole, they will also create laggards and losers. Fortunately, high levels of productivity growth may allow for more expensive redistribution measures. One prominently discussed idea in response to structural change is to invest in the retraining or upskilling of workers, to help them adapt to the reorganization of the economy (Brynjolfsson 2023). This might involve acquiring AI-related skills and skills in areas AI cannot perform so well in. Effective retraining programs should be accessible, inclusive, customized, and aligned with the market demand. AI may also help with retraining, e.g. by making personalized learning much more accessible than it has been historically (Mindell, Reynolds, et al. 2023). Other policy proposals include regular national skills assessment, and robust support (income assistance and job placement services) for workers transitioning from automated roles (Chhabria and Marchant 2024). Note that reducing the diffusion speed of AI technologies, via regulatory sandboxes or gradual or staged deployment, may allow for more gradual adjustments of the labor force. Finally, note that if capital income grows faster than labor income, then labor income share falls, and inequality may rise. Consequently, we believe redistributive measures to still remain an important measure in the toolkit of adaptive policies.

Under a scenario of technological supremacy, the overall effect of AI on labor demand is ambiguous. As workers are being displaced with highly productive machines, capital-owners are likely to experience a rise in income. However, strong productivity effects could undo the adverse effects of AI at least partially. In either case, as labor's share in national income drops significantly, societies will experience large scale increases in inequality. Still, the abundance created by productivity growth allows for more costly redistribution schemes, including both income and wealth redistribution through policies such as universal basic income (Bell and Korinek 2023). Korinek (2023) recommends launching a “seed UBI” to prepare the infrastructure for such eventuality. Others view UBI as a defeatist and disempowering alternative to paid employment (Johnson and Acemoglu 2023). Wealth redistribution schemes could take the form of universal basic capital or the distribution of shares of large companies (Moser 2021). Arguably, antitrust policies may also become particularly important as economic power might be concentrated in the hands of tech giants. Finally, this is also the scenario where power imbalance and safety issues can become too pressing to leave the technologies in private hands, so that a nationalization of some parts of the supply

chain may be optimal (Naudé and Dimitri 2020). Clearly, the feasibility of nationalization depends on the context; it is more relevant in economies which are at or close to the technological frontier, and when international coordination helps to mitigate race dynamics from other sovereign actors.

It is easy to dismiss adaptation policies as futuristic and not immediately relevant. However, the rapid pace of technological advancement necessitates that policymakers begin preparing for these changes now. Implementing foundational measures today ensures that the necessary infrastructure and systems are in place when more significant impacts arise. For instance, a 'seed UBI' could be introduced to test and refine the mechanisms for a more comprehensive universal basic income program. By starting with a modest UBI, policymakers can evaluate its effects, identify potential issues, and fine-tune the delivery systems, making a future transition smoother and more efficient.

4 Discussion

This policy paper has introduced a framework for analyzing AI's impact on the labor market and suggested policies for mitigating potentially adverse effects. More specifically, we have argued that two factors will ultimately determine AI's effect on labor demand: (i) the net extent to which AI technologies will replace or automate human tasks, and (ii) the overall productivity effects of AI deployment. Within this framework, we then presented four distinct scenarios for how AI could affect the demand for labor. We call these scenarios low-effect AI, so-so automation, productive synergy and technological supremacy. To ensure AI's impact on labor demand is positive, we suggested a set of policies focussed on 'steering' technological progress towards ensuring strong labor demand. This included ideas for directing R&D towards labor-augmentation, for instance by reducing the tax on labor income, by strengthening unionization and by controlling diffusion of AI technologies. However, we argued that these policies should be complemented by a set of 'adaptation' policies, which aim to mitigate the adverse effects of AI on employment and economic stability without directly altering the trajectory of technological progress.

As an important limitation to our work, we stress that our analysis largely abstracts from important heterogeneities in how AI will affect workers: AI automation is not and will not be uniform but differ significantly across professions. Previous waves of automation and technological change too did not affect all workers alike but were biased, i.e. benefitting some workers and harming others (Autor, Levy, and Murnane 2003). An important strand of the literature concerns itself with attempting to predict the tasks which will be most affected by AI technologies. Such studies map the tasks required within a specific occupation (e.g. based on occupational dictionaries such as the O*Net) to current AI capabilities. A shared finding among these studies is that, while AI is likely to affect the workforce in its entirety, in the short-run AI appears to mainly impact

4 DISCUSSION

high-skilled and highly-paid workers (Eloundou et al. 2023; Felten, Raj, and Seamans 2021; Webb 2019)). Less clear is whether AI will substitute or complement these workers.

This underlines an important point: Ultimately, only time can answer the question of how AI will affect labor markets. Existing empirical evidence on the labor market effects of AI needs to be interpreted cautiously. AI-technologies are still nascent and evolving at a fast pace, and empirical studies can only give an account of the effects of current technologies. In addition, many of these studies are micro-econometric, which implies we should be careful in extrapolating onto the aggregate level. One important set of experimental studies has examined the effects of AI-assistance on the productivity of workers (Brynjolfsson and Unger 2023; Dell’Acqua et al. 2023; Hui, Reshef, and Zhou 2024; Noy and Zhang 2023; Otis et al. 2023; Peng et al. 2023). A key finding of this literature is that AI seems to significantly increase the productivity of workers. More worryingly, these studies also seem to suggest that AI tends to substitute human labor. For example, Noy and Zhang (2023) find that users seemed to take over the AI’s output relatively unaltered and Hui, Reshef, and Zhou (2024) find that earnings and income of freelancers in highly affected occupations tend to decrease as a result of introduction of generative AI. Taken together, these early results can be cautiously seen as being suggestive of a technological supremacy scenario.

We also note the predictions of economists and computer scientists tend to diverge fairly strongly (Nordhaus 2021). We believe that promoting interdisciplinary research is essential to better understand and anticipate AI’s impacts on labor markets. Collaboration between economists and technologists can yield valuable insights and guide effective policy decisions. These partnerships can help bridge the gap between technological advancements and economic policies, fostering a more inclusive and sustainable future. Furthermore, the successful implementation of the proposed policies will require careful planning and consideration of potential barriers. Policymakers must address logistical challenges, potential resistance from stakeholders, and the need for substantial funding. Continuous monitoring and evaluation of AI’s impact on labor markets will be critical to adapt policies as needed. Finally, given the global nature of AI development and deployment, international cooperation and policy harmonization will be essential to manage cross-border implications and share best practices.

References

Acemoglu, Daron, and Todd Lensman. 2024. “Regulating transformative technologies.” *American Economic Review: Insights* 6 (3): 359–376.

Acemoglu, Daron, and Pascual Restrepo. 2018a. “Artificial Intelligence, Automation, and Work.” In *The Economics of Artificial Intelligence: An Agenda*, 197–236. University of Chicago Press.

———. 2018b. “The Race Between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment.” *American Economic Review* 108 (6): 1488–1542.

———. 2019. “Automation and New Tasks: How Technology Displaces and Reinstates Labor.” *Journal of Economic Perspectives* 33 (2): 3–30.

Agrawal, Ajay, Joshua S Gans, and Avi Goldfarb. 2019. “Artificial Intelligence: the Ambiguous Labor Market Impact of Automating Prediction.” *Journal of Economic Perspectives* 33 (2): 31–50.

Agrawal, Ajay K, Joshua S Gans, and Avi Goldfarb. 2023. *The Turing Transformation: Artificial intelligence, intelligence augmentation, and skill premiums*. Technical report. National Bureau of Economic Research.

Ahn, Michael J. 2024. “Navigating the future of work: A case for a robot tax in the age of AI.”

Andrews, Dan, Chiara Criscuolo, and Peter N Gal. 2016. “The best versus the rest: the global productivity slowdown, divergence across firms and the role of public policy.”

Anwar, Usman, Abulhair Saparov, Javier Rando, Daniel Paleka, Miles Turpin, Peter Hase, Ekdeep Singh Lubana, Erik Jenner, Stephen Casper, Oliver Sourbut, et al. 2024. “Foundational challenges in assuring alignment and safety of large language models.” *arXiv preprint arXiv:2404.09932*.

Autor, David. 2014. *Polanyi’s paradox and the shape of employment growth*. Technical report. National Bureau of Economic Research.

———. 2022. “The Labor Market Impacts of Technological Change: From Unbridled Enthusiasm to Qualified Optimism to Vast Uncertainty.” *National Bureau of Economic Research Working Paper 30074*.

Autor, David, Caroline Chin, Anna Salomons, and Bryan Seegmiller. 2024. “New frontiers: The origins and content of new work, 1940–2018.” *The Quarterly Journal of Economics*, qjae008.

REFERENCES

Autor, David H, David Dorn, and Gordon H Hanson. 2016. “The China shock: Learning from labor-market adjustment to large changes in trade.” *Annual review of economics* 8 (1): 205–240.

Autor, David H, Frank Levy, and Richard J Murnane. 2003. “The Skill Content of Recent Technological Change: An Empirical Exploration.” *The Quarterly Journal of Economics* 118 (4): 1279–1333.

Baily, Martin, Erik Brynjolfsson, and Anton Korinek. 2023. “Machines of mind: The case for an AI-powered productivity boom.”

Belfield, Haydn, and Shin-Shin Hua. n.d. “Compute and Antitrust: Regulatory implications of the AI hardware supply chain, from chip design to foundation model APIs.”

Bell, Stephanie A, and Anton Korinek. 2023. “AI’s Economic Peril.” *Journal of Democracy* 34 (4): 151–161.

Bernardi, Jamie, Gabriel Mukobi, Hilary Greaves, Lennart Heim, and Markus Anderljung. 2024. “Societal Adaptation to Advanced AI.” *arXiv preprint arXiv:2405.10295*.

Bloom, Nicholas, Charles I Jones, John Van Reenen, and Michael Webb. 2020. “Are ideas getting harder to find?” *American Economic Review* 110 (4): 1104–1144.

Bostrom, Nick. 2009. “The future of humanity.” *Geopolitics, History, and international relations* 1 (2): 41–78.

Bresnahan, Timothy F, and Manuel Trajtenberg. 1995. “General purpose technologies ‘Engines of growth’?” *Journal of econometrics* 65 (1): 83–108.

Brynjolfsson, Erik. 2023. “The turing trap: The promise & peril of human-like artificial intelligence.” In *Augmented education in the global age*, 103–116. Routledge.

Brynjolfsson, Erik, and ANDREW McAfee. 2017. “Artificial intelligence, for real.” *Harvard business review* 1:1–31.

Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock. 2018. “What can machines learn and what does it mean for occupations and the economy?” In *AEA papers and proceedings*, 108:43–47. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2021. “The productivity J-curve: How intangibles complement general purpose technologies.” *American Economic Journal: Macroeconomics* 13 (1): 333–372.

REFERENCES

Brynjolfsson, Erik, and Gabriel Unger. 2023. “The Macroeconomics of Artificial Intelligence.”

Calvino, Flavio, and Chiara Criscuolo. 2022. “Gone digital: Technology diffusion in the digital era.”

Chhabria, Pooja, and Natalie Marchant. 2024. “Davos 2024: 6 innovative ideas on reskilling, upskilling and building a future-ready workforce.” *World Economic Forum*, <https://www.weforum.org/stories/2024/01/davos-2024-6-innovative-ideas-for-reskilling-upskilling-and-building-a-future-ready-workforce/>.

David, Paul A. 1990. “The dynamo and the computer: an historical perspective on the modern productivity paradox.” *The American economic review* 80 (2): 355–361.

Dell’Acqua, Fabrizio, Edward McFowland III, Ethan R Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Krayer, François Candelier, and Karim R Lakhani. 2023. “Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality.” *Harvard Business School Technology & Operations Mgt. Unit Working Paper 24-013*.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2023. “Gpts are gpts: An early look at the labor market impact potential of large language models.” *arXiv preprint arXiv:2303.10130*.

Felten, Edward, Manav Raj, and Robert Seamans. 2021. “Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses.” *Strategic Management Journal* 42 (12): 2195–2217.

Good, Irving John. 1966. “Speculations concerning the first ultraintelligent machine.” In *Advances in computers*, 6:31–88. Elsevier.

Gordon, Robert J. 2018. *Why has economic growth slowed when innovation appears to be accelerating?* Technical report. National Bureau of Economic Research.

Hui, Xiang, Oren Reshef, and Luofeng Zhou. 2024. “The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market.” *Organization Science*.

Johnson, Simon, and Daron Acemoglu. 2023. *Power and progress: Our thousand-year struggle over technology and prosperity*. Hachette UK.

Korinek, Anton. 2023. “Scenario planning for an A (G) I future.” *Finance and Development* 60 (4): 30–33.

REFERENCES

Korinek, Anton, and Joseph E Stiglitz. 2020. “Steering technological progress.” In *NBER Conference on the Economics of AI*.

Manyika, James, Michael Chui, Mehdi Miremadi, Jacques Bughin, Katy George, Paul Willmott, and Martin Dewhurst. 2017. “A future that works: AI, automation, employment, and productivity.” *McKinsey Global Institute Research, Tech. Rep* 60:1–135.

Mindell, David A, Elisabeth Reynolds, et al. 2023. *The work of the future: Building better jobs in an age of intelligent machines*. Mit Press.

Moser, Elias. 2021. “Machines and technological unemployment: Basic income vs. basic capital.” In *Machine law, ethics, and morality in the age of artificial intelligence*, 205–225. IGI Global.

Naudé, Wim, and Nicola Dimitri. 2020. “The race for an artificial general intelligence: implications for public policy.” *AI & society* 35:367–379.

Nordhaus, William D. 2021. “Are we approaching an economic singularity? information technology and the future of economic growth.” *American Economic Journal: Macroeconomics* 13 (1): 299–332.

Noy, Shakked, and Whitney Zhang. 2023. “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence.” *Science* 381 (6654): 187–192.

OECD. 2023. “Regulatory sandboxes in artificial intelligence,” no. 356, <https://doi.org/https://doi.org/https://doi.org/10.1787/8f80a0e6-en>. <https://www.oecd-ilibrary.org/content/paper/8f80a0e6-en>.

Otis, Nicholas, Rowan P Clarke, Solene Delecourt, David Holtz, and Rembrand Koning. 2023. “The uneven impact of generative AI on entrepreneurial performance.” *Available at SSRN 4671369*.

Peng, Baolin, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. 2023. “Instruction tuning with gpt-4.” *arXiv preprint arXiv:2304.03277*.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman. 2018. “Distributional national accounts: methods and estimates for the United States.” *The Quarterly Journal of Economics* 133 (2): 553–609.

Polanyi, Michael. 2009. “The tacit dimension.” In *Knowledge in organisations*, 135–146. Routledge.

Sastry, Girish, Lennart Heim, Haydn Belfield, Markus Anderljung, Miles Brundage, Julian Hazell, Cullen O’Keefe, Gillian K Hadfield, Richard Ngo, Konstantin Pilz, et al. 2024. “Computing Power and the Governance of Artificial Intelligence.” *arXiv preprint arXiv:2402.08797*.

REFERENCES

Seger, Elizabeth, Aviv Ovadya, Divya Siddarth, Ben Garfinkel, and Allan Dafoe. 2023. “Democratizing AI: Multiple meanings, goals, and methods.” In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, 715–722.

Solaiman, Irene. 2023. “The gradient of generative AI release: Methods and considerations.” In *Proceedings of the 2023 ACM conference on fairness, accountability, and transparency*, 111–122.

Susskind, Daniel. 2023. *Work and meaning in the age of AI*.

The Economist. 2023. “A new age of the worker will overturn conventional thinking,” <https://www.economist.com/leaders/2023/11/30/a-new-age-of-the-worker-will-overturn-conventional-thinking>.

———. 2024. “AI holds tantalising promise for the emerging world,” <https://www.economist.com/leaders/2024/01/25/the-tantalising-promise-of-ai-for-the-emerging-world>.

Trammell, Philip, and Anton Korinek. 2023. *Economic growth under transformative AI*. Technical report. National Bureau of Economic Research.

Vipra, Jai, and Anton Korinek. 2023. “Market concentration implications of foundation models.” *arXiv preprint arXiv:2311.01550*.

Webb, Michael. 2019. “The Impact of Artificial Intelligence on the Labor Market.” Available at *SSRN 3482150*.