Artificial Intelligence in the Knowledge Economy*

Enrique Ide & Eduard Talamàs

April 1, 2024

Abstract

How does Artificial Intelligence (AI) affect the organization of work? We incorporate AI into an economy where humans endogenously sort into hierarchical firms: Less knowledgeable agents become "workers" (execute routine tasks), while more knowledgeable agents become "managers" (specialize in problem-solving). We model AI as an algorithm that uses computing power to mimic the behavior of humans with a given knowledge. We show that AI not only leads to occupational displacement but also changes the endogenous matching between *all* workers and managers. This leads to new insights regarding AI's effects on productivity, firm size, and degree of decentralization.

1 Introduction

Artificial intelligence (AI) is a new, powerful form of automation based on machines that can perform cognitive tasks typically associated with the human mind. While the transformative impact of AI on the landscape of work is undeniable, its precise implications have become the center of a growing controversy. Indeed, predictions range from optimistic utopian visions to ominous apocalyptic scenarios (Meserole, 2018; Brynjolfsson, 2022; Johnson and Acemoglu, 2023; Autor, 2024).

Unlike previous waves of automation—that primarily led to the creation of tools proficient at handling repetitive tasks at scale—AI has demonstrated the capacity to address highly non-routine tasks previously reserved for highly skilled workers with sophisticated knowledge (Webb, 2020; Autor, 2024).[1] Consequently, it is not immediately apparent that experience with previous waves of au-

*IESE Business School, C/ de Arnús i de Garí 3-7, 08034 Barcelona, Spain (eide@iese.edu and etalamas@iese.edu). We have benefited from the comments of Nano Barahona, Andrea Galeotti, Benjamin Golub, Sebastián Otero, and seminar participants at the various venues where this work has been presented. All errors are our own. We declare we have no relevant or material financial interests that relate to the research described in this paper.

[1] According to Autor (2024) "AI's capacity to depart from script, to improvise based on training and experience, enables it to engage in expert judgment — a capability that, until now, has fallen within the province of elite experts. Though only in its infancy, this is a superpower."

tomation will seamlessly extrapolate to the case of AI.[2]

With this issue in mind, in this paper we propose a new framework to study the effects of AI on the labor market. Our approach is motivated by two key observations. First, AI holds the potential to automate knowledge work. Second, as underscored by Hayek (1945), a fundamental economic challenge for society is the efficient utilization of the available knowledge. Therefore, to uncover how firms use AI and the resulting effects on the future of work, it is important to consider how economic agents organize themselves in terms of knowledge and time to achieve common production goals.

The overarching theme of our findings is that AI not only creates occupational displacement but also changes the endogenous matching between all workers and managers in the economy. This leads to new insights. For example, while there is a growing interest in understanding the extent to which AI will substitute or augment humans (e.g., Brynjolfsson, 2022; Autor, 2024), we show how, once organizational effects are taken into account, human-like artificial intelligence necessarily does both: It complements some humans while substituting others.

The starting point of our analysis is the seminal papers by Antràs et al. (2006), and Fuchs et al. (2015), who consider an economy where labor is the sole input for production. Humans are endowed with one unit of time and are heterogeneous in terms of knowledge. Production occurs when a human dedicates her full unit of time to production, and her knowledge exceeds the difficulty of the problems she confronts.

The competitive equilibrium in this setting—which we take as our pre-AI outcome—involves humans either trying to solve problems on their own (becoming "independent producers") or sorting into hierarchical firms to make more efficient use of their time and knowledge. These firms have two key features. First, they exhibit *management by exception*: Less knowledgeable humans become "workers" who try to solve problems first, while more knowledgeable humans become "managers" specializing in the exceptional problems the first layer of workers was unable to solve. Second, there is *positive assortative matching*: More knowledgeable workers are always matched to more knowledgeable managers.[3]

Our innovation is to incorporate AI into this otherwise canonical setting. We model AI as an algorithm that uses computing power (or "compute") to mimic the behavior of a human with a given knowledge. Hence, AI is an automation technology.[4] As standard in the automation literature, we

[2]For instance, Muro et al. (2019) state: . . . " [When studying AI] most research has concentrated on an undifferentiated array of 'automation' technologies including robotics, software, and AI all at once . . . [As a result,] past 'automation' analyses—including our own—have likely obscured AI's distinctive impact."

[3]There is both anecdotal and systematic empirical evidence showing the emergence of such "knowledge hierarchies" (see, e.g., Garicano and Hubbard, 2012; Caliendo and Rossi-Hansberg, 2012; Caliendo et al., 2015, 2020). For instance, Alfred Sloan (1924), a former head of General Motors (GM), once wrote: "We do not do much routine work with details. They never get up to us. I work fairly hard, but on exceptions."

[4]We follow this approach—rather than modeling AI as a tool that directly complements humans—because it appears that there are currently stronger incentives for automation over pure augmentation technologies among technologists and

assume all firms have access to this technology. Firms that use AI are identical to those that do not, except that they rent compute instead of hiring labor to do either production or managerial work. Moreover, we assume that AI can be used *at scale* in the following two senses: (i) the same algorithm can be simultaneously used by all units of compute, and (ii) compute is large relative to human time.

We then characterize the competitive equilibrium when firms have access to AI. We show that if AI is used as a worker, it is necessarily the *most* knowledgeable worker in the economy, so it is supervised by the *most* knowledgeable humans. In addition, if AI is used as a manager, it is necessarily the *least* knowledgeable manager in the economy, so it assists the *least* knowledgeable humans.

A notable property of the equilibrium is that, as long as AI has unlimited potential applications, the rental rate of compute is equal to AI's knowledge, and, as a result, AI does not lead to the complete destruction of human routine jobs. This result arises despite the fact that AI is more knowledgeable than a subset of the human population and can be used at scale. Intuitively, even though compute is large relative to human time, it is still scarce relative to its potential applications. Hence, it continues to be worthwhile for every human to be employed in some capacity. The competitive equilibrium then allocates all those humans who are less knowledgeable than AI to be workers because that is their comparative advantage (since they are less likely than AI to succeed on their own).[5]

We then turn to our main endeavor: Analyzing the effects of AI on labor outcomes by comparing the pre- and post-AI equilibrium. In particular, we study the implications of AI for (i) occupational choice, (ii) the size and productivity of firms, (iii) the productivity of the non-displaced workers and the size of the firms supervised by the non-displaced managers, and (iv) labor income.

Regarding occupational choice, we show that when AI has the knowledge of a pre-AI worker, it displaces humans from routine to managerial work. In contrast, when AI has the knowledge of a pre-AI manager, the displacement goes in the opposite direction. Intuitively, when AI has the knowledge of a pre-AI worker, AI gives firms access to a relatively cheap technology to do routine work. This reduces workers' wages and increases the attractiveness of creating hierarchical firms. The result is a surge in the demand for managers, which induces the most knowledgeable routine workers of the pre-AI equilibrium to switch to managerial roles. A similar intuition explains why AI displaces the least knowledgeable managers to routine work when it has the knowledge of a pre-AI manager.

These displacement results have remarkable implications for the distribution of firm size (defined as firm output), span of control (defined as time/compute under its manager supervision), and productivity (defined as firm output divided by its span of control).[6] Indeed, when AI has the knowledge of a pre-AI worker, its introduction leads to smaller, less productive, and more centralized two-layer

executives (Acemoglu and Restrepo, 2019; Brynjolfsson, 2022; Johnson and Acemoglu, 2023).

[5] In Section 5, we consider the case where compute is so large that it is abundant not only relative to human time but also relative to its potential applications. In that case, the price of compute is zero and AI leads to the complete destruction of human routine jobs. Firms, however, still have a hierarchical structure: All problems are initially attempted by AI, and all those humans who are more knowledgeable than AI specialize in solving problems that AI cannot solve.

[6] Note that a firm's span of control is also equal to the time/compute it uses for production.

organizations. This is because, when the best pre-AI workers switch to managerial work, they create smaller and less productive firms while destroying the biggest and most decentralized firms. Since the opposite displacement occurs when AI has the knowledge of a pre-AI manager, AI leads to bigger, more productive, and more decentralized two-layer organizations in that case.

AI not only creates some firms and destroys others but affects all matches in the economy. This implies that AI also affects the productivity and span of control of the workers and managers *who have not been occupationally displaced*. We show that irrespective of AI's knowledge, AI decreases the productivity of non-displaced workers (except possibly the least knowledgeable ones) and increases the span of control of non-displaced managers (except possibly the most knowledgeable ones).

Intuitively, when AI has the knowledge of a pre-AI worker, the knowledge required to become a manager decreases, worsening the pool of managers available for non-displaced workers. At the same time, the match of the best managers worsens as they now supervise production by AI (while, pre-AI, they were managing humans more knowledgeable than AI), but the match of all the other managers improves because the newly appointed managers now supervise the worst workers.

Similar forces arise when AI has the knowledge of a pre-AI manager: The match of the worst workers improves because they are now assisted by AI (while, pre-AI, they were assisted by humans less knowledgeable than AI). However, the knowledge of the best workers increases, improving the pool of available workers for non-displaced managers and, therefore, leaving a worse pool of managers for the non-displaced workers that were originally managed by the humans who are more knowledgeable than AI.

Finally, we study the effects of AI on labor income. We show that AI increases total labor income but that it creates winners and losers in the labor market. Since the wage of every individual is equal to her marginal product, this implies that AI necessarily substitutes some humans (in the sense that it reduces their marginal product) and necessarily complements others (in the sense that it increases their marginal product).[7]

These distributional effects are shaped by two potentially countervailing forces: On the one hand, AI changes the composition of firms and, therefore, the quality of matches. On the other hand, AI changes the relative scarcities of different knowledge levels, affecting how each firm's output is divided between workers and managers. Which of these two forces dominates depends on the interaction between AI's knowledge and communication costs: Only the most knowledgeable humans benefit when AI's knowledge is low, while only the least knowledgeable humans benefit when AI's knowledge and communication costs are high. In any case, the winners from AI are always at the extremes of the knowledge distribution.

The rest of this paper is organized as follows. Following a brief discussion of the most closely re-

[7]Note that an agent's marginal product (defined as the output increase of introducing such an agent into the economy) is not necessarily equal to her productivity (defined as the agent's expected output). This is because introducing an agent into the economy might affect the output of *other agents in the economy* through changes in workers-manager matching.

lated literature, we present our model in Section 2. After describing the pre-AI equilibrium in Section 2.3, we characterize the equilibrium with AI in Section 3. In Section 4, we turn to our main endeavor: Comparing the pre- and post-AI equilibrium; this section contains our main results. Section 5 discusses several extensions, including the effects of AI on the labor share and labor inequality and the possibility of technological unemployment. Section 6 concludes.

Related literature

This paper contributes to two different literatures. First, it introduces automation and AI to the literature on knowledge hierarchies. Second, it incorporates organizations and a different way of thinking about AI to the automation literature. In this section, we discuss how these contributions relate to existing work.

The literature on knowledge hierarchies starts with Garicano (2000), which introduces the technology and describes the circumstances under which knowledge hierarchies are optimal when agents are homogenous. Garicano and Rossi-Hansberg (2004, 2006) embed this model in a setting with heterogenous agents to study how the endogenous organization of knowledge interacts with inequality. Fuchs et al. (2015), in turn, show that firm-like contractual arrangements uniquely deliver the first-best when there is double-sided asymmetric information about participants' ability to solve problems.[8] In contrast to our paper, none of the work in this literature considers automation or artificial intelligence.

In the context of knowledge hierarchies, the paper closest to us is Antràs et al. (2006), who study the effects of offshoring by comparing the equilibrium of a closed economy with one in which firms can form international teams. The main difference with our paper is that offshoring gives firms access to a population of humans with different knowledge, while artificial intelligence gives firms access to an algorithm that can solve problems *at scale*. As we discuss in detail in Section 5.6, this implies that the effects of AI are qualitatively different than those of offshoring.

Our paper also contributes to the literature on automation (e.g., Zeira, 1998; Acemoglu and Autor, 2011; Aghion et al., 2017; Acemoglu and Restrepo, 2018, 2022; Azar et al., 2023; Acemoglu and Loebbing, 2024; Korinek and Suh, 2024). This literature uses task-based models to study the effects of automation on labor outcomes, inequality, and economic growth.

The first important recent contribution to this literature is due to Zeira (1998), who shows how automation can lead to a decline in the labor share as the economy develops. Acemoglu and Restrepo (2018) counter that, by depressing the equilibrium wage, automation also encourages the creation of new tasks in which labor has a comparative advantage. Hence, even though the direct effect of

[8]Other important contributions to this literature include Garicano and Hubbard (2007), Caliendo and Rossi-Hansberg (2012), Garicano and Rossi-Hansberg (2012), Caicedo et al. (2019), and Carmona and Laohakunakorn (Forthcoming). See Garicano and Rossi-Hansberg (2015) for an excellent survey on knowledge hierarchies.

automation is to decrease the labor share, the economy might still exhibit a balanced growth path in which the labor share stays constant over time. Expanding on this, Acemoglu and Restrepo (2022) delve into the effects of automation on wage inequality and show that a substantial fraction of the changes in the US wage structure in recent decades can be attributed to relative wage declines among worker groups undergoing rapid automation. More recently, Acemoglu and Loebbing (2024) show how reductions in the cost of capital can cause employment and wage polarization.

We depart from this literature in two key respects. First, we explicitly integrate organizations into the analysis by employing knowledge hierarchies. As explained by Garicano and Rossi-Hansberg (2015), this is a particular specification of the task-based framework, where tasks are hierarchical and the relationship between them arises from an explicit organizational problem. By incorporating organizations into our analysis, we present novel insights about the impact of AI on the size and productivity distributions of firms. Moreover, our approach offers a new perspective on how AI influences labor outcomes through its effects on firm composition and organizational structure.

Second, we adopt a different approach to model AI. While existing literature often models automation as the capability of machines to replace workers in certain tasks, we assume that AI can perform the exact same tasks as humans but can only mimic the performance of a subset of the human population. This approach is motivated by the idea that—in contrast to previous waves of automation—artificial intelligence may be indistinguishable from human intelligence. Moreover, it offers a significant advantage: the tasks that are automated (i.e., AI's comparative advantage) are determined endogenously as a function of (i) the knowledge of the population, (ii) the level of information and communication technologies, and (iii) the advancement of AI.

2 The Model

We consider a perfectly competitive economy where producing output requires solving problems. The model builds on Garicano (2000), Antràs et al. (2006), and Fuchs et al. (2015). Our innovation is to introduce AI in this otherwise canonical setting. After describing the model in Section 2.1, we discuss its main assumptions in Section 2.2.

2.1 The Baseline Setting

The Pre-AI Economy.—There is a unit mass of humans, each endowed with one unit of time and exogenous knowledge $z \in [0,1]$. The distribution of knowledge in the population is given by a continuous probability distribution with full support on $[0,1]$, cumulative distribution function $G(z)$, and density $g(z)$. The knowledge of each individual is perfectly observable.

There is a large measure of potential and identical competitive firms that can enter the market. Entry is free, and firms are risk-neutral. Production occurs inside firms—who are the residual claimants of all output—and requires solving problems.

In particular, upon entering the market, firms hire humans to produce. Each of these "production workers" devotes her full unit of time and applies her knowledge to a single problem of difficulty x. The difficulty of each problem is ex-ante unknown and distributed uniformly on $[0, 1]$, independently across problems.[9] If the production worker's knowledge exceeds the problem's difficulty, she solves the problem and produces one unit of output. Otherwise, the production worker is unable to produce output by herself.

Additionally, firms can hire another human to act as their workers' "manager." In that case, production workers can ask their manager for help when they are unable to solve a problem on their own. This exchange consumes $h \in (0, 1)$ units of the manager's time. If the manager's knowledge exceeds the problem's difficulty, the manager communicates the solution to the problem to the corresponding worker, who then produces a unit of output. Otherwise, no production takes place. We normalize the value of each unit of output to one.

Artificial Intelligence.— We model AI as an algorithm that can use compute to mimic the behavior of a human with knowledge $z_{\text{AI}} \in [0, 1)$. Hence, AI is an automation technology. We refer to z_{AI} as "AI's knowledge" (in Section 5.1, we discuss the case of artificial "superintelligence," i.e., $z_{\text{AI}} = 1$).

All firms have access to AI. Thus, in contrast to the pre-AI economy, firms decide not only their organizational structure but also whether to use this technology. Firms that use AI are identical to those that do not, except that they use AI instead of humans in either production or managerial work. To do this, they must rent one unit of compute per production worker or manager they replace.[10] The amount of compute in the economy—which is exogenously given—is denoted by $\mu \geq 0$.

Wages, Prices, and Profits.—Let $w(z)$ be the wage of a human with knowledge z and denote by r the rental rate of one unit of compute. All agents in this economy are income maximizers.

The problem of an active firm is to decide (i) whether to use humans or AI in production (and the knowledge of the human workers, when appropriate), (ii) whether to operate as a one-layer or two-layer organization, and (iii) in the case of two-layer organizations, whether to use a human or AI as manager (and the knowledge of the human manager, when appropriate).

The profits of a single-layer organization are simply:

$$\Pi_1 = \begin{cases} z - w(z) & \text{if the firm hires a human with knowledge } z \\ z_{\text{AI}} - r & \text{if the firm uses AI} \end{cases}$$

[9]The assumption that x is uniformly distributed is without loss, given that the distribution of knowledge in the population is arbitrary (i.e., assuming a uniform distribution is simply a normalization).

[10]The assumption that the algorithm uses the same amount of compute irrespective of the difficulty of the problem it faces is in line with how current AI models operate. See Fridman, Lex. "Sam Altman: OpenAI, GPT-5, Sora, Board Saga, Elon Musk, Ilya, Power & AGI." *The Lex Fridman Podcast #419*, March 18, 2024. https://lexfridman.com/sam-altman-2-transcript (accessed March 23, 2024).

The profits of a two-layer organization, in turn, depend on whether it does not use AI (a "nA" firm), uses AI as a manager (i.e., automates the top layer, a "tA" firm), or uses AI as a worker (i.e., automates the bottom layer, a "bA" firm).[11] In either case, we restrict attention to matching arrangements in which all workers matched with a given manager have the same knowledge. This restriction is without loss because the equilibrium exhibits positive assortative matching.

Consequently, to exploit its manager's time fully, a two-layer organization that hires workers with knowledge z optimally hires exactly $n(z) = [h(1-z)]^{-1}$ workers. The profits of a two-layer organization as a function of its type (nA, tA, or bA) are thus:

$$\Pi_2^{n\text{A}}(z,m) = n(z)[m - w(z)] - w(m)$$
$$\Pi_2^{t\text{A}}(z) = n(z)[z_{\text{AI}} - w(z)] - r$$
$$\Pi_2^{b\text{A}}(m) = n(z_{\text{AI}})[m - r] - w(m)$$

where z and m (which must be greater than z) denote the knowledge of a human worker and manager, respectively.

For brevity, we reserve the term "worker" for the agents engaging in production in two-layer organizations, while we use the term "independent producer" to refer to the agents working in production in single-layer organizations. Also, note that as in Antràs et al. (2006), we identify a given firm either by its independent producer (in the case of single-layer organizations) or by the manager who runs it (in the case of two-layer organizations).

Competitive Equilibrium.— Denote by I the set of humans hired as independent producers, and let W_p and W_a be the human workers managed by other humans and managed by AI, respectively.[12] Similarly, denote by M_p the set of humans who manage other humans, and by M_a the set of humans managing AI. Finally, let μ_i, μ_w, and μ_m be the amount of compute rented for independent production, production in two-layer firms, and supervision of humans, respectively.

Definition (Competitive Equilibrium). An equilibrium consists of non-negative amounts (μ_i, μ_w, μ_m), sets (W_p, W_a, I, M_p, M_a), a feasible matching function $f : W_p \to M_p$, a wage schedule $w : [0,1] \to \mathbb{R}_+$ and a rental rate of compute $r \in \mathbb{R}_+$, such that:

1. Firms optimally choose their structure (while earning zero profits).
2. nA firms that hire workers with knowledge z hire a manager with knowledge $f(z)$.
3. tA firms hire workers with knowledge in W_a.
4. bA firms hire managers with knowledge in M_a.
5. Markets clear: (i) $\mu_i + \mu_w + \mu_m = \mu$, and (ii) the union of the sets (W_p, W_a, I, M_p, M_a) is $[0,1]$ and the intersection of any two of these sets has measure zero.

[11] Note that a firm will never use AI at both layers of the organization. This is because an AI manager solves the exact same problems as its AI worker.

[12] We use the subscript "p" (for people) instead of "h" (for human), to avoid any confusion with the helping cost h.

Compute is "Abundant" Relative to Time.— Our main goal is to analyze the effects of AI by comparing the pre-AI equilibrium with the post-AI equilibrium when compute is abundant relative to human time:

$$\int_0^{z_{\mathrm{AI}}} n(z)^{-1} dG(z) + n(z_{\mathrm{AI}})\big(1 - G(z_{\mathrm{AI}})\big) < \mu \tag{1}$$

This inequality guarantees that the binding constraint in human-AI interactions is human time, not compute.[13]

To understand this condition, note that a firm will never hire a manager who is less knowledgeable than its workers. This implies that the most intensive (though not necessarily optimal) way to use compute in organizations comprising humans and AI is to (i) use AI to manage everyone who is less knowledgeable than AI (which requires $\int_0^{z_{\mathrm{AI}}} n(z)^{-1} dG(z)$ units of compute) and (ii) make every human that is more knowledgeable than AI supervise $n(z_{\mathrm{AI}})$ units of compute (which requires $n(z_{\mathrm{AI}})(1 - G(z_{\mathrm{AI}}))$ units of compute). Hence, condition (1) implies that there are not enough humans to interact with all compute in the post-AI world. In Section 5.3, we explore the role of this assumption by comparing the pre- and post-AI equilibrium in the opposite extreme in which compute is arbitrarily small.

Some Notation.— For future reference, we define $W \equiv W_a \cup W_p$ and $M \equiv M_a \cup M_p$ as the overall set of human workers and managers of this economy, respectively. We also denote by $e(m)$ the inverse of the matching function $f(z)$. That is, $e(m)$ is the "employee matching function" denoting the knowledge of the human worker matched with a human manager with knowledge m. This function always exists given that, as shown below, the equilibrium matching function is strictly increasing.

Finally, we denote by $\mathrm{int}S$ the interior of the set S. We also use $S \preceq S'$ to symbolize the idea that the set $S \subseteq [0,1]$ "lies below" the set $S' \subseteq [0,1]$. Formally, $S \preceq S'$ if $\sup S \leq \inf S'$ (or either S or S' is empty). For example, $W_a \preceq W_p$ means that the best worker managed by AI is weakly less knowledgeable than the worst worker managed by humans.

2.2 Discussion of the Model

Before moving on to the analysis, we briefly comment on some assumptions underlying our model. First, we model AI as an automation technology, i.e., an algorithm that can mimic the behavior of a human with a given knowledge. An alternative would be to model AI as a tool that does not create value by itself but complements humans. We adopt the former stance because there are stronger incentives for automation over pure augmentation technologies (Acemoglu and Restrepo, 2019; Brynjolfsson, 2022; Johnson and Acemoglu, 2023).

[13] Note that for any distribution G and helping cost $h \in (0,1)$, there exist a finite μ that satisfies this condition for all $z_{\mathrm{AI}} \in (0,1)$. This follows from the fact that the left-hand side of (1) is continuous in $z_{\mathrm{AI}} \in (0,1)$ and is bounded as $z_{\mathrm{AI}} \to 0$ and $z_{\mathrm{AI}} \to 1$ (it converges to $1/h$ and $g(1)/h + \int_0^1 n(z)^{-1} dG(z)$, respectively).

Second, the distinguishing feature of AI relative to human intelligence is that it can be used *at scale*. This manifests in the model in two distinct ways: (i) compute is large relative to human time, and (ii) unlike human knowledge (whose application is constrained by the time of the individual who possesses it), AI can be leveraged across *all units* of compute (implying that all units of compute can solve problems up to the same difficulty). Our motivation for (i) is that, in contrast to time, compute has been growing exponentially over the past two centuries (Nordhaus, 2007).[14] Our motivation for (ii) is that digital information is non-rival and has nearly zero marginal cost of reproduction (Brynjolfsson and McAfee, 2016).

Third, our model assumes that AI has unlimited applications (i.e., compute is abundant relative to time but scarce relative to the number of society's problems) and that AI can perform the same roles in the knowledge economy as humans (i.e., AI can be a worker, a manager, or an independent producer). In Section 5.4, we discuss how our results change when compute is also abundant relative to its applications, while in Section 5.7, we provide current real-world examples of AI being used in the three different roles of the knowledge economy.

Fourth, our main goal with this paper is to analyze how AI affects human labor outcomes. For this reason, we take AI technology and the economy's compute as given. An important implication of this assumption is that compute is exclusively used for the deployment of AI systems rather than for their training.[15] Studying firms' incentives to develop AI or to increase the economy's compute are intriguing avenues for future research.

Fifth, we follow Antràs et al. (2006) and Fuchs et al. (2015) in assuming that the distribution of human knowledge is exogenous and that organizations have at most two layers. We opt for these assumptions primarily for the sake of simplicity, although we believe they offer a good first approximation to the problem at hand. Understanding the impact of AI when human retraining or more complex organizations are feasible is left for future research.

Sixth and finally, we have introduced a large measure of anonymous firms whose role in the economy is to organize production. The existence of such firms is not strictly needed: The same outcome in terms of allocations and income arises if workers or managers are the owners of firms. However, introducing this large measure of anonymous firms allows us to present the model in the most concise and cleanest way possible.

[14]For instance, Nordhaus (2007) documents that the compound logarithmic growth rate of compute from 1850 to 2006 was approximately 18.3 percent per year. Moreover, he notes a significant acceleration during the period from 1940 to 2006, with the growth rate nearly doubling to approximately 36 percent per year.

[15]The industry separates AI's use of compute between "training" (i.e., teaching AI systems how to respond) and "deployment" or "inferencing" (i.e., processing and reacting to new bits of information). As of 2023, more than 40% of Nvidia's data center business (the leading supplier of specialized microchips for AI) was for the deployment of AI systems, and that share is predicted to grow in the future. See Asa Fitch, "How a Shifting AI Chip Market Will Shape Nvidia's Future," *The Wall Street Journal*, February 25, 2024, https://www.wsj.com/tech/ai/how-a-shifting-ai-chip-market-will-shape-nvidias-future-f0c256b1 (accessed February 26, 2024).

2.3 Benchmark: The Pre-AI Equilibrium

We begin by presenting a partial characterization of the equilibrium without AI (for the full characterization, see Appendix A). This is also the equilibrium when compute is zero and was originally described by Fuchs et al. (2015).[16] Note that in this case $W_a = M_a = \emptyset$, so $W = W_p$ and $M = M_p$.

Proposition 1. *In the absence of AI, there is a unique equilibrium. The shape of this equilibrium depends on whether the "helping cost" h is above or below a cutoff $h_0 \in (0, 1)$:*

- *If $h \leq h_0$, there is $\hat{z} \in (0, 1)$ such that $W = [0, \hat{z}]$, $I = \emptyset$, and $M = [\hat{z}, 1]$.*
- *If $h > h_0$, there are cutoffs $0 < \underline{z} < \bar{z} < 1$ such that $W = [0, \underline{z}]$, $I = (\underline{z}, \bar{z})$ and $M = [\bar{z}, 1]$.*

In either case, there is strictly positive assortative matching (i.e., $f : W \to M$ is strictly increasing), and the wage function $w(z)$ is continuous, strictly increasing, and (weakly) convex. Moreover, it satisfies:

- *$w(z) > z$ for $z \in W$ (except possibly at $z = \sup W$), and $w'(z) \in (0, 1)$, and $w''(z) > 0$ for $z \in \mathrm{int}W$.*
- *$w(z) = z$ for all $z \in I$.*
- *$w(z) > z$ for $z \in M$ (except possibly at $z = \inf M$), and $w'(z) > 1$, and $w''(z) > 0$ for $z \in \mathrm{int}M$.*

Proof. See Appendix A. □

The equilibrium without AI—which we illustrate in Figure 1—has several salient features. First, it exhibits occupational stratification: Managers are more knowledgeable than independent producers (i.e., $I \preceq M$), who, in turn, are more knowledgeable than workers (i.e., $W \preceq I$). This is because the marginal value of knowledge is strictly higher than 1 for managers (as they can leverage their knowledge by applying it to more than one problem), exactly equal to 1 for independent producers (as expected output equals knowledge for them), and strictly lower than 1 for workers (as their knowledge is used to free up managerial time).

Second, there is positive assortative matching: More knowledgeable workers are matched with more knowledgeable managers. The reason is that worker and managerial knowledge are complements, as a more knowledgeable manager helps produce more output (so she should manage a large team), and team size is increasing in the knowledge of workers (as more knowledgeable workers ask fewer questions).

Third, workers and managers earn strictly more than their output as independent producers (except possibly in the case of the most knowledgeable worker and the least knowledgeable manager). Indeed, since the marginal value of manager knowledge is larger than 1, if managers earned their expected output as independent producers, all two-layer firms would only want to hire the most knowledgeable agents as managers. Similarly, since the marginal value of worker knowledge is less

[16] Note that an economy with $\mu = 0$ is different than an economy with $\mu > 0$ but $z_{\mathrm{AI}} = 0$. This is because, even if AI cannot solve any problems, it can still draw them, enlarging the production possibility frontier of the economy.

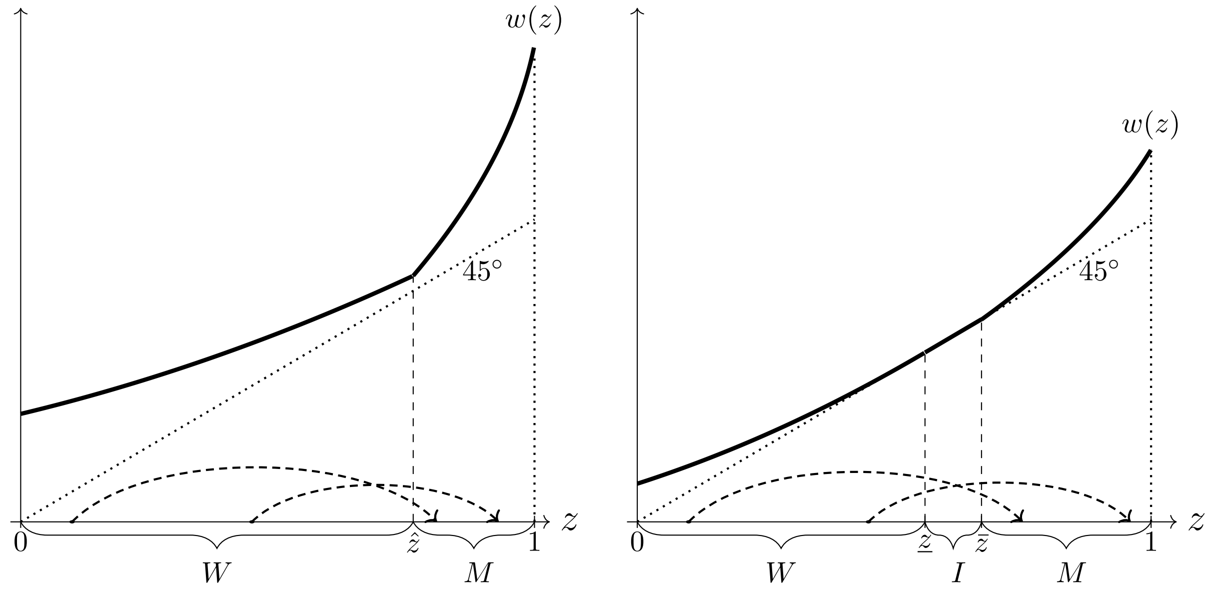

(a) Pre-AI Equilibrium when $h < h_0$ (b) Pre-AI Equilibrium when $h > h_0$

Figure 1: Illustration of the Pre-AI equilibrium.

Notes. Distribution of knowledge: $G(z) = z$. Parameter values: Panel (a) has $h = 1/2$ ($< h_0 = 3/4$), while panel (b) has $h = 0.8125$ ($> h_0 = 3/4$). The thick line depicts the wage function $w(z)$. The dashed arrows illustrate the matching function at two arbitrary points. See Section 3 of the Online Appendix for more details on the construction of this figure.

than 1, if workers earned their expected output as independent producers, all two-layer firms would only want to hire the agents with the least knowledge.

Fourth and finally, the income function $w(z)$ is continuous, strictly increasing, and weakly convex. Monotonicity arises because the marginal value of knowledge is strictly positive, and its convexity (which is strict when $z \in W \cup M$) is due to the existence of complementarities between managers' and workers' knowledge. Continuity, on the other hand, is a necessary condition for market clearing.

For simplicity, in what follows, we restrict attention to $h < h_0$. This implies that there are no independent producers in the pre-AI equilibrium. As discussed in Section 5.1, virtually all of our results extend to $h \geq h_0$.

3 The AI Equilibrium

In this section, we present a partial characterization of the AI equilibrium with the essential information needed to understand our main results. We relegate the complete characterization to Appendix B because many of its details can be safely skipped on a first reading (but are needed for the formal arguments). For future reference, we index this equilibrium using the superscript "$*$" (note that the pre-AI equilibrium has no superscript).

Proposition 2. *In the presence of AI, there is a unique equilibrium. The equilibrium has the following features:*

- *The price of compute is equal to AI's knowledge:* $r^* = z_{\text{AI}}$.
- *Occupational stratification:* $W^* \preceq I^* \preceq M^*$.
- *No worker better than AI; no manager worse than AI:* $W^* \preceq \{z_{\text{AI}}\} \preceq M^*$.
- *Positive assortative matching:* $f^* : W_p^* \to M_p^*$ *is strictly increasing and* $W_a^* \preceq W_p^*$ *and* $M_p^* \preceq M_a^*$.
- *If* $z_{\text{AI}} > 0$*, AI does not lead to the complete destruction of human routine jobs:* $W^* \neq \emptyset$.

Furthermore, AI is always used for independent production, and whether it is also used as a worker or as a manager depends on its knowledge ***relative to the pre-AI equilibrium****.*

- *If* $z_{\text{AI}} \in W$*, then AI is necessarily used as a worker (and possibly also as a manager).*
- *If* $z_{\text{AI}} \in M$*, then AI is necessarily used as a manager (and possibly also as a worker).*

Finally, the wage function $w^*(z)$ *is continuous, strictly increasing, and (weakly) convex, and satisfies:*

- $w^*(z) = z_{\text{AI}}(1 - 1/n(z)) > z$ *for all* $z \in W_a^*$.
- $w^*(z) = f^*(z) - w^*(f^*(z))/n(z) > z$ *for all* $z \in W_p^*$.
- $w^*(z) = z$ *for all* $z \in I^*$.
- $w^*(z) = m_p + \int_{m_p}^{z} n(e^*(u))du > z$*, for all* $z \in M_p^*$*, where* $m_p \equiv \inf M_p^*$.
- $w^*(z) = n(z_{\text{AI}})(z - z_{\text{AI}}) > z$*, for all* $z \in M_a^*$.

Proof. Note that the statement implies that in the knife-edge case where AI has the knowledge of both a pre-AI worker and a pre-AI manager, i.e., $z = \hat{z}$, AI is necessarily used in all three roles. For the proof, see Appendix B. □

Proposition 2 has three parts. The first one states fundamental properties of the equilibrium. The second one describes how firms use AI as a function of AI's knowledge. Finally, the third part characterizes the equilibrium wages. The remaining of this section provides intuition for each of these three parts.

3.1 Fundamental Properties of the AI Equilibrium

First, the equilibrium price of compute is equal to AI's knowledge because compute is abundant relative to human time. Indeed, by definition, this implies that there are not enough humans to interact with AI inside two-layer organizations. Hence, some compute must be used for independent production. The zero-profit condition of the single-layer firms using AI then requires $r^* = z_{\text{AI}}$..

Second, the equilibrium continues to exhibit occupational stratification (as the marginal value of knowledge continues to be the smallest for workers, the second smallest for independent producers, and the highest for managers) and positive assortative matching (as there are still complementarities between worker and managerial knowledge).

Third, occupational stratification plus the fact that some compute must be used for independent production implies that no worker can be better than AI and that no manager can be worse than AI, i.e., $W^* \preceq \{z_{\text{AI}}\} \preceq M^*$. Positive assortative matching then implies that if AI is used as a worker, then it is managed by the most knowledgeable humans, i.e., $M_p^* \preceq M_a^*$. Similarly, if AI is used as a manager, it manages the least knowledgeable humans, i.e., $W_a^* \preceq W_p^*$.

Finally, if $z_{\text{AI}} > 0$, AI does not completely destroy worker positions for humans.[17] This result arises even though compute is abundant and AI is more knowledgeable than a fraction of the human population. Intuitively, even though compute is large relative to human time, it is still scarce relative to its potential applications. Hence, it continues to be worthwhile for every human to be employed in some capacity. The competitive equilibrium then allocates all those humans who are less knowledgeable than AI to do routine work in two-layer organizations because that is their comparative advantage (since they are less likely than AI to succeed on their own).

3.2 How Firms Use AI

As mentioned above, some compute must necessarily be deployed in independent production. As stated in Proposition 2, whether AI is used in any other capacity depends on its knowledge relative to the pre-AI equilibrium. For brevity, here we only explain the case where $z_{\text{AI}} \in W$ since the case where $z_{\text{AI}} \in M$ is analogous.

To start, note that AI cannot be exclusively used as an independent producer. To see this, suppose otherwise for contradiction. Then, AI does not affect labor outcomes, as it does not interact with humans in the workplace. This implies that the human with AI's knowledge is still being employed as a worker and is, therefore, earning strictly more than z_{AI}. Hence, the firm hiring this worker has incentives to replace her with one unit of compute. A contradiction.

Next, we argue that AI must necessarily be used as a worker. To see this, suppose for contradiction that AI is exclusively used as a manager and as an independent producer. By occupational stratification, everyone with knowledge above z_{AI} is also a manager, so AI increases the number of managers and decreases the number of workers. This violates market clearing, a contradiction.

Let us now explain when AI is used (i) only as a worker and as an independent producer and (ii) in all three roles. To do this, we use Figure 2, which illustrates the effects of introducing AI starting from the pre-AI scenario depicted in Figure 1 panel (a). The introduction of AI can be thought of as unleashing the following chain of events.

First, immediately upon its introduction, AI forces those humans with knowledge $z \in [z_{\text{AI}}, \hat{z}]$ to switch from workers to independent producers (recall that AI is always the most knowledgeable worker in equilibrium) and reduces the wages of all those workers who are (weakly) less knowledgeable than AI, i.e., those with $z \leq z_{\text{AI}}$. In both cases, there is a decline in the earnings of humans

[17] If $z_{\text{AI}} = 0$, all humans become managers supervising AI and the resulting wage function is $w^*(z) = z/h$.

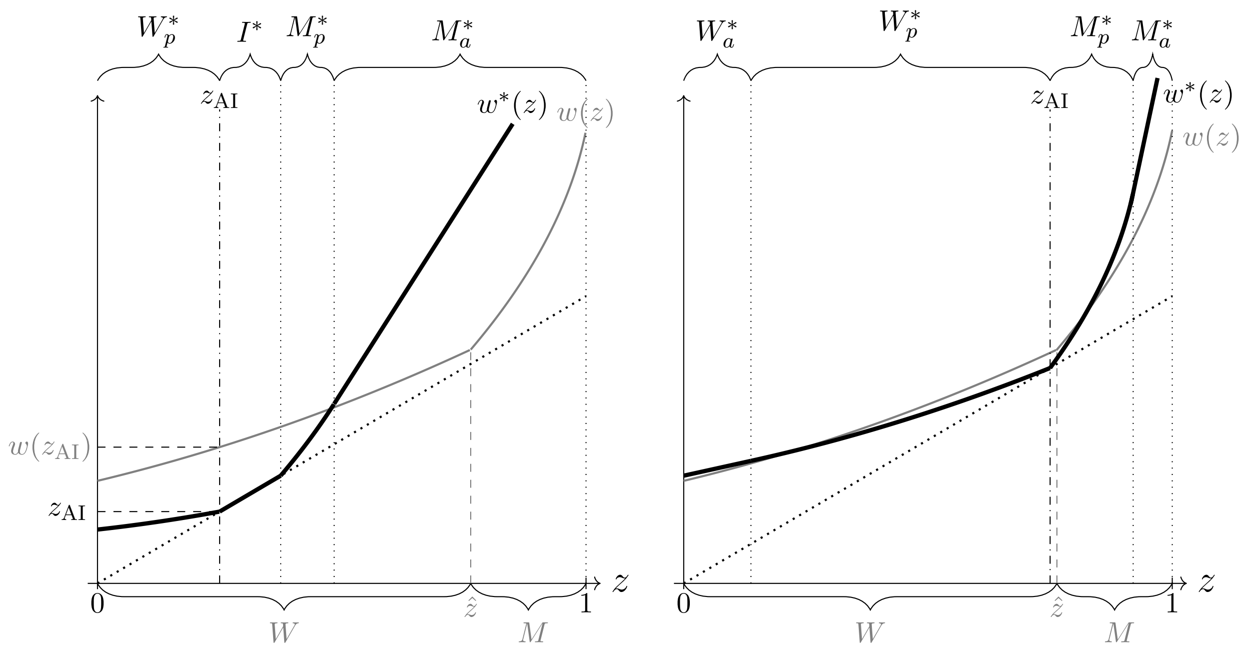

(a) AI is used as a Worker and as an Independent Producer

(b) AI is Used in All Three Roles

Figure 2: From the Pre- to the Post-AI Equilibrium

Notes. Distribution of knowledge: $G(z) = z$. Parameter values: Both panels have $h = 1/2$ ($< h_0 = 3/4$), so $\hat{z} = 3 - \sqrt{5}$. Panel (a) has $z_{\text{AI}} = 0.25$, while panel (b) has $z_{\text{AI}} = 0.75$. See Section 3 of the Online Appendix for more details on the construction of this figure.

both to the left and to the right of AI's knowledge level, as AI provides firms with a relatively cheap technology to do routine work.

The reduction in workers' wages, coupled with the diminished income of the newly created class of independent producers, then increases the attractiveness of creating two-layer organizations. Consequently, there is a surge in demand for managerial roles: Firms begin hiring managers who are less knowledgeable than the least knowledgeable managers before the advent of AI.

If the surge in managerial positions is relatively modest compared to the mass of humans in $[z_{\text{AI}}, \hat{z}]$, then the last newly hired manager is still more knowledgeable than AI, so AI is used as a worker and as an independent producer. This is the case depicted in Figure 2(a). Otherwise, the class of human independent producers initially brought into existence by AI is completely absorbed by the newly created managerial positions, and firms begin using AI to manage the least knowledgeable workers. Hence, in this case, AI plays all three roles in equilibrium. This situation is illustrated in Figure 2(b).

3.3 Equilibrium Wages

Finally, we turn to the third part of Proposition 2: Equilibrium wages. For insight, we informally construct them in the case in which AI is used in all three possible roles.

By the first part of Proposition 2, when AI is used in all three roles the human population must necessarily partition as follows:

$$W_a^* \preceq W_p^* \preceq \{z_{\mathrm{AI}}\} \preceq M_p^* \preceq M_a^* \text{ with } \sup W_p^* = z_{\mathrm{AI}} = \inf M_p^*$$

The wage of a worker with knowledge $z \in W_a^*$ is determined by the zero-profit condition of tA firms plus the fact that the equilibrium rental rate of compute is $r^* = z_{\mathrm{AI}}$. The wage of a manager with knowledge $z \in M_a^*$ is determined similarly, though using the zero profit condition of bA firms.

Constructing the wages of those hired by nA firms requires more work. First, recall that $f^* : W_p^* \to M_p^*$ is the function matching human workers with human managers and that $e^*(z) = (f^*)^{-1}(z)$. Now consider the problem of a nA firm that recruited $n(z)$ workers with knowledge $z \in W_p^*$ and is deciding which manager $m \in M_p^*$ to hire:

$$\max_{m \in M_p^*} n(z)[m - w(z)] - w(m)$$

The corresponding first-order condition evaluated at $m = f^*(z)$ implies that $w^{*\prime}(f^*(z)) = n(z)$, or, equivalently, $w^{*\prime}(z) = n(e^*(z))$ for any $z \in M_p^*$. Thus:

$$w^*(z) = C + \int_{z_{\mathrm{AI}}}^{z} n(e^*(u))du, \text{ for any } z \in M_p^*$$

The constant C is then determined by the fact that the least knowledgeable manager has the same knowledge as AI, $\inf M_p^* = z_{\mathrm{AI}}$, so her wage must equal the price of one unit of compute: $C^* = r^* = z_{\mathrm{AI}}$. Finally, the wages of the workers being managed by other humans come from the zero profit condition of nA firms:

$$w^*(z) = f^*(z) - \frac{w^*(f^*(z))}{n(z)}, \text{ for any } z \in W_p^*$$

4 How AI Reorganizes Work

In this section, we study how AI reorganizes work by comparing the pre-AI and post-AI equilibrium. For the analysis that follows, we use the following terminology:

- $z_{\mathrm{AI}} \in \mathrm{int}W$: "AI has the knowledge of a pre-AI worker."
- $z_{\mathrm{AI}} \in \mathrm{int}M$: "AI has the knowledge of a pre-AI manager."
- $z_{\mathrm{AI}} \in W \cap M = \{\hat{z}\}$: "AI has the knowledge of both a pre-AI worker and a pre-AI manager."

4.1 Occupational Displacement

We begin by analyzing the effects of AI on occupational choices.

Proposition 3.

- *If $z_{\mathrm{AI}} \in \mathrm{int}W$, then AI displaces humans from routine to managerial work, i.e., $W^* \subset W$ and $M^* \supset M$.*

- *If $z_{\mathrm{AI}} \in \mathrm{int}M$, then AI displaces humans from managerial to routine work, i.e., $W^* \supset W$ and $M^* \subset M$.*
- *If $z_{\mathrm{AI}} = \hat{z}$, then there is no human displacement between routine and managerial work, i.e., $W^* = W$ and $M^* = M$. However, AI leads to the creation of bA and tA firms, i.e., $W^*_a \neq \emptyset$ and $M^*_a \neq \emptyset$.*

Proof. Note that $M^* \supset M$ if and only if $(W^* \cup I^*) \subset (W \cup I)$, since $(W^* \cup I^*)$ is the complement of M^* and $(W \cup I)$ the complement of M. For the proof of the proposition, see Appendix C. □

According to the proposition, while it is true that AI might lead to the displacement of humans from routine to complex jobs—as popular wisdom usually states—it might also lead to the opposite outcome. Interestingly, the nature of the displacement is determined by the knowledge of AI *relative to the pre-AI equilibrium*: When AI has the knowledge of a pre-AI worker, it displaces humans from routine work to managerial positions. In contrast, when AI has the knowledge of a pre-AI manager, the displacement goes in the opposite direction.

To provide intuition, consider the case in which $z_{\mathrm{AI}} \in \mathrm{int}W$. The result that $W^* \subset W$ follows directly from the fact that in the AI equilibrium, no worker is better than AI, i.e., $W^* \preceq \{z_{\mathrm{AI}}\}$. The more interesting result is that $M^* \supset M$. The reason is that AI has two opposing effects on the demand for managers.

On the one hand, by destroying human worker positions, AI reduces manager demand. On the other hand, by incentivizing firms to use compute to do routine work, AI increases manager demand. However, the second effect more than compensates for the first effect. This is because the added "workers" (or, more precisely, the compute allocated to production) are less knowledgeable than the human workers displaced (i.e., those $z \in W$ and $z \geq z_{\mathrm{AI}}$), and therefore, it is more valuable to match them with managers to provide them help.

The intuition for when $z_{\mathrm{AI}} \in \mathrm{int}M$ is analogous, so let us discuss the knife-edge situation where $z_{\mathrm{AI}} = \hat{z}$. In this case, by lowering the wages of the best workers and the worst managers, AI increases the demand for managers (to match with the cheaper workers) and increases the demand for workers (to match with the cheaper managers). As a result, AI is used in both worker and managerial roles, so it is both the best worker and the worst manager. It follows that, in this case, $W^* = W$ and $M^* = M$.

The fact that the nature of the occupational displacement is determined by the knowledge of AI relative to the pre-AI equilibrium has an additional important implication. Indeed, as shown in Section 1 of the Online Appendix, in a pre-AI equilibrium without independent production, the knowledge cutoff to become a manager decreases in h. Moreover, as shown in the same appendix, this knowledge cutoff is pointwise higher for any admissible h if the human population is more knowledgeable (in a first-order stochastic dominance sense).

Hence, if more developed countries have better communication technologies and/or a more knowledgeable human population, then the same AI technology might displace humans from routine to managerial work in more developed countries but displace humans in the opposite direction in less developed ones. The reverse, however, cannot happen: If AI displaces humans into managerial roles in less-developed countries, then the same must happen in more developed ones.[18]

4.2 Distribution of Firm Size, Productivity, and Span of Control

We now describe the effects of AI on the distribution of firms' size, productivity, and span of control. For simplicity, we focus on two-layer organizations. We take *firm size* to be its output, and *firm productivity* to be its output divided by the units of time and/or compute it uses for production. We define a manager's *span of control* as the units of time or compute under her supervision. Note that firm size is equal to its productivity times its manager's span of control.

We begin with the following preliminary result:

Corollary 1. *AI necessarily increases the number of two-layer firms.*

This corollary is a direct implication of Proposition 3. Because its proof is illuminating and relatively straightforward, we provide it as part of the main text:

Proof. Given that any two-layer organization is identified by the manager who runs it, it suffices to show that AI increases the overall number of managerial positions in the economy. When $z_{\text{AI}} \in \text{int}W$, this follows because $M^* \supset M$. When $z_{\text{AI}} \in \text{int}M$, more humans become workers after AI's introduction (i.e., $W^* \supset W$). Hence, the overall number of managers—human plus AI—must increase as each worker requires the same amount of help post-AI as pre-AI. Finally, when $z_{\text{AI}} = \hat{z}$, the result follows because all the managers pre-AI continue to be managers post-AI, and some compute is allocated to managerial work. □

4.2.1 Productivity

Denote by $\mathcal{P}(x)$ and $\mathcal{P}^*(x)$ the measure of firms with productivity less than or equal to x pre- and post-AI, respectively. Since the productivity of each firm is equal to its manager's knowledge, the support of $\mathcal{P}$ is M and the support of $\mathcal{P}^*$ is M^*. Moreover:

$$\mathcal{P}(x) = \begin{cases} 0 & \text{if } x < \hat{z} \\ G(x) - G(\hat{z}) & \text{if } \hat{z} \leq x < 1 \\ 1 - G(\hat{z}) & \text{if } 1 \leq x \end{cases}$$

$$\mathcal{P}^*(z) = \begin{cases} 0 & \text{if } x < \inf M^* \\ G(x) - G(\inf M^*) + \mu_m^* & \text{if } \inf M^* \leq x < 1 \\ 1 - G(\inf M^*) + \mu_m^* & \text{if } 1 \leq x \end{cases} \tag{2}$$

[18] An implicit assumption underlying this argument is that there are no cross-national teams. Understanding the impact of AI on offshoring is an interesting and relevant area for future research.

The next corollary, which is illustrated in Figure 3, describes the effects of AI on the distribution of firms' productivity:

Corollary 2. *Let $\delta_{\mathcal{P}}(x,y)$ and $\delta^*_{\mathcal{P}}(x,y)$ denote the measure of firms with productivity between x and y pre-AI and post-AI, respectively.*[19] *We have that $\delta_{\mathcal{P}}(x,y) = \delta^*_{\mathcal{P}}(x,y)$ for all $x, y \in M \cap \text{int}M^*$. Moreover,*

- *If $z_{\text{AI}} \in \text{int}W$, then AI extends the support of the distribution to include $[\inf M^*, \hat{z})$ and may create a mass of firms with productivity z_{AI}.*
- *If $z_{\text{AI}} \in \text{int}M$, then AI eliminates all firms with productivity below z_{AI}, and creates a mass of firms with productivity z_{AI}.*
- *If $z_{\text{AI}} = \hat{z}$, then AI does not affect the support of the distribution but creates a mass of firms with productivity $\hat{z} = z_{\text{AI}}$.*

Proof. Immediate from (2) and Propositions 2 and 3. □

The intuition for this corollary is relatively straightforward. First, AI does not affect the measure of firms with productivity in $M \cap \text{int}M^*$ because such a measure is equal to the mass of humans with knowledge in that set. Second, whether AI creates a mass of firms with productivity z_{AI} depends on whether a positive mass of compute is allocated to managerial roles. This explains why such firms

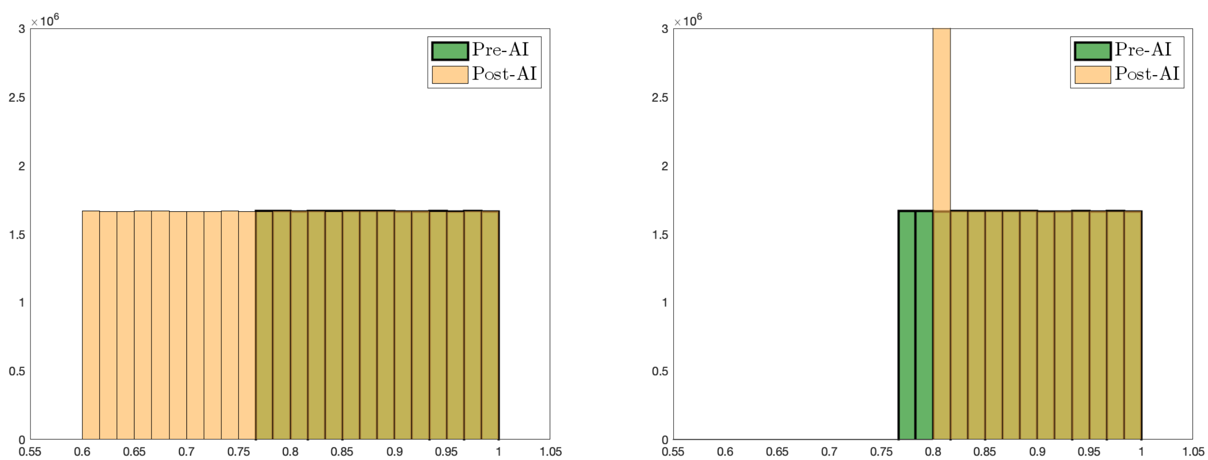

(a) $z_{\text{AI}} \in \text{int}W$ (b) $z_{\text{AI}} \in \text{int}M$

Figure 3: The Effects of AI on the Distribution of Firm's Productivity

Notes. Histogram based on a human population of $N = 100 \times 10^6$ individuals. Distribution of knowledge: $G(z) = z$. Both panels have $h = 1/2$ $(< h_0 = 3/4)$, so $\hat{z} = 3 - \sqrt{5}$. Moreover, panels (a) and (c) assume $z_{\text{AI}} = 0.5$, while panels (b) and (d) assume $z_{\text{AI}} = 0.8$. Pre-AI, there are 23.6×10^6 firms. Post-AI, the number of firms increases to 41.7×10^6 in panel (a) and to 30.3×10^6 in panel (b). See Section 3 of the Online Appendix for more details on the construction of this figure.

[19] That is, $\delta_{\mathcal{P}}(x,y) \equiv \mathcal{P}(x) - \mathcal{P}(y)$ and $\delta^*_{\mathcal{P}}(x,y) \equiv \mathcal{P}^*(x) - \mathcal{P}^*(y)$.

necessarily emerge when $z_{\rm AI} \in {\rm int}M$ or $z_{\rm AI} = \hat{z}$ (as AI is necessarily used as a manager in these cases), but not when $z_{\rm AI} \in {\rm int}W$ (as AI may only be a worker and an independent producer).

Finally, the effects of AI on the support of the productivity distribution are driven by occupational displacement: AI leads to the emergence of less productive firms when $z_{\rm AI} \in {\rm int}W$, as it induces the best pre-AI workers to become managers post-AI. Conversely, AI leads to the destruction of the economy's least productive firms when $z_{\rm AI} \in {\rm int}M$, as it induces the worst pre-AI managers to become workers post-AI.

4.2.2 Span of Control

Let $\mathcal{N}(x)$ and $\mathcal{N}^*(x)$ be the measure of firms with a span of control less than or equal to x pre- and post-AI, respectively. Note that a firm with span of control x has workers with knowledge $1-n(0)/x$. Hence, the mass of firms with span of control less than x is equal to the number of managers (humans or AI) required to supervise the workers with knowledge in $z \in [0, 1-n(0)/x]$. This implies that the supports of $\mathcal{N}$ of $\mathcal{N}^*$ are $N = [n(0), n(\sup W)]$ and $N^* = [n(0), n(\sup W^*)]$, respectively, and that:

$$\mathcal{N}(x) = \begin{cases} 0 & \text{if } x < n(0) \\ \int_0^{1-\frac{n(0)}{x}} h(1-z)dG(z) & \text{if } n(0) \le x < n(\hat{z}) \\ \int_0^{\hat{z}} h(1-z)dG(z) & \text{if } n(\hat{z}) \le x \end{cases}$$
$$\mathcal{N}^*(x) = \begin{cases} 0 & \text{if } x < n(0) \\ \int_0^{1-\frac{n(0)}{x}} h(1-z)dG(z) & \text{if } n(0) \le x < n(\sup W^*) \\ \int_0^{\sup W^*} h(1-z)dG(z) + h(1-z_{\rm AI})\mu_w^* & \text{if } n(\sup W^*) \le x \end{cases} \tag{3}$$

The following corollary describes the effects of AI on the distribution of span of control:

Corollary 3. *Let $\delta_{\mathcal{N}}(x,y)$ and $\delta^*_{\mathcal{N}}(x,y)$ denote the measure of firms with span of control between x and y pre-AI and post-AI, respectively. We have that $\delta_{\mathcal{N}}(x,y) = \delta^*_{\mathcal{N}}(x,y)$ for all $x, y \in N \cap {\rm int}N^*$. Moreover,*

- *If $z_{\rm AI} \in {\rm int}W$, then AI eliminates all firms with span of control above $n(\sup W^*)$, and creates a mass of firms with span of control $n(z_{\rm AI})$.*
- *If $z_{\rm AI} \in {\rm int}M$, then AI extends the support of the distribution to include $(n(\hat{z}), n(\sup W^*)]$ and may create a mass of firms with span of control $n(z_{\rm AI})$.*
- *If $z_{\rm AI} = \hat{z}$, AI does not affect the support of the distribution but creates a mass of firms with span of control $n(\hat{z}) = n(z_{\rm AI})$.*

Proof. Immediate from (3) and Propositions 2 and 3. □

Corollary 3 is illustrated in Figure 4. The intuition for this corollary is similar to that of Corollary 2 with some subtle twists. First, note that firms with span of control in $N \cap {\rm int}N^*$ hire humans that are

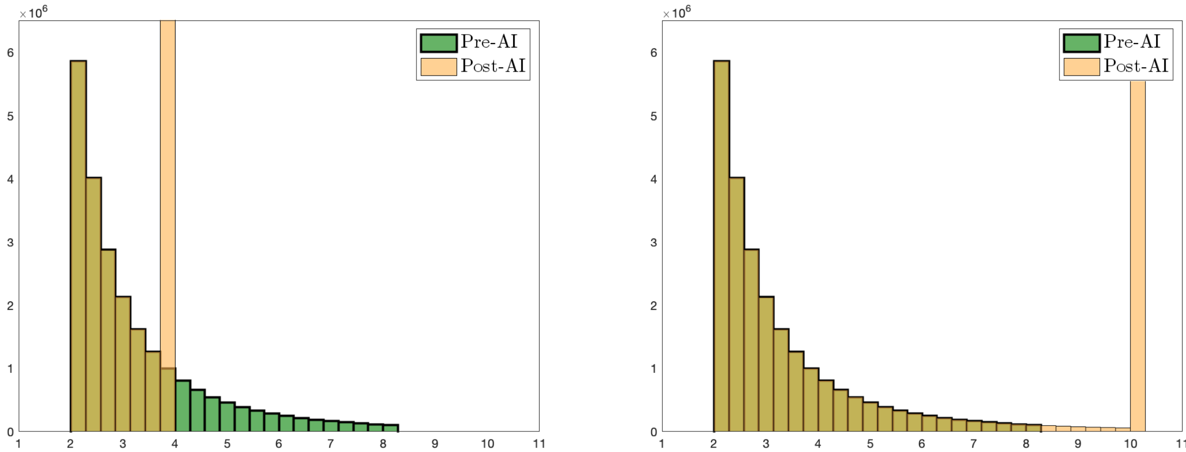

(a) $z_{\rm AI} \in {\rm int}W$

(b) $z_{\rm AI} \in {\rm int}M$

Figure 4: The Effects of AI on the Distribution of Firms' Span of Control

Notes. Histogram based on a human population of $N = 100 \times 10^6$ individuals. Distribution of knowledge: $G(z) = z$. Both panels have $h = 1/2$ ($< h_0 = 3/4$), so $\hat{z} = 3 - \sqrt{5}$. Moreover, panels (a) and (c) assume $z_{\rm AI} = 0.5$, while panels (b) and (d) assume $z_{\rm AI} = 0.8$. Pre-AI, there are 23.6×10^6 firms. Post-AI, the number of firms increases to 41.7×10^6 in panel (a) and to 30.3×10^6 in panel (b). See Section 3 of the Online Appendix for more details on the construction of this figure.

workers both pre- and post-AI. Since a given worker requires the same amount of managerial time irrespective of AI, there are exactly the same amount of firms hiring these workers pre- and post-AI. This explains why $\delta_{\mathcal{N}}(x, y) = \delta^*_{\mathcal{N}}(x, y)$ for all $x, y \in N \cap {\rm int}N^*$.

Second, whether AI creates a mass of firms with span of control $n(z_{\rm AI})$ depends on whether AI is used as a worker (which is only guaranteed when $z_{\rm AI} \in {\rm int}W$). Third, because a firm's span of control is determined by the knowledge of its workers, the equilibrium set of workers determines the support of the distribution of the span of control. The effects of AI on such support are then driven by occupational displacement: When $z_{\rm AI} \in {\rm int}W$, AI induces the best pre-AI workers to become managers post-AI, shrinking the support from above. Conversely, when $z_{\rm AI} \in {\rm int}M$, AI induces the worst pre-AI managers to become workers post-AI, expanding the support from above.

4.2.3 Firm Size

The effects of AI on the distribution of firm size are more involved. The reason is that size depends on both worker and managerial knowledge, and the reorganizations brought about by AI change all worker-manager matches in the economy. This implies that AI necessarily changes the mass of firms between any two points of the distribution. Nevertheless, we can still obtain sharp results regarding the effects of AI on its support:[20]

[20] See section 2 of the Online Appendix for the expression for the distribution of firm size pre- and post-AI.

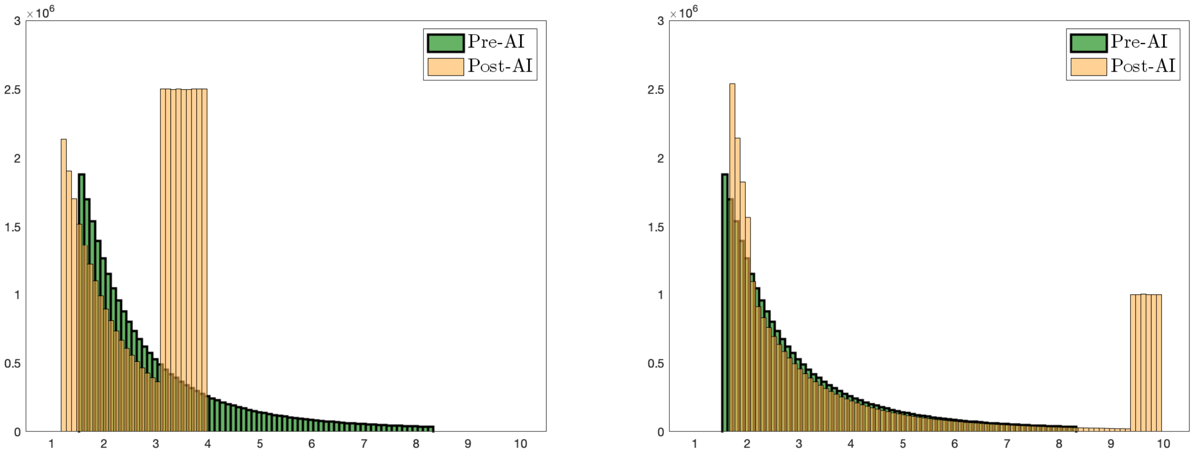

(a) $z_{\text{AI}} \in \text{int}W$ (b) $z_{\text{AI}} \in \text{int}M$

Figure 5: The Effects of AI on the Distribution of Firms' Size

Notes. Histogram based on a human population of $N = 100 \times 10^6$ individuals. Distribution of knowledge: $G(z) = z$. Both panels have $h = 1/2$ ($< h_0 = 3/4$), so $\hat{z} = 3 - \sqrt{5}$. Moreover, panels (a) and (c) assume $z_{\text{AI}} = 0.5$, while panels (b) and (d) assume $z_{\text{AI}} = 0.8$. Pre-AI, there are 23.6×10^6 firms. Post-AI, the number of firms increases to 41.7×10^6 in panel (a) and to 30.3×10^6 in panel (b). See Section 3 of the Online Appendix for more details on the construction of this figure.

Corollary 4. *Let S and S^* be the support of the size distribution of firms pre- and post-AI:*

- *If $z_{\text{AI}} \in \text{int}W$, then $\inf S^* < \inf S$ and $\sup S^* < \sup S$.*
- *If $z_{\text{AI}} \in \text{int}M$, then $\inf S^* > \inf S$ and $\sup S^* > \sup S$.*
- *If $z_{\text{AI}} = \hat{z}$, then $\inf S^* = \inf S$ and $\sup S^* = \sup S$.*

Proof. Follows from positive assortative matching, the fact that $z = 0$ is always a worker and $z = 1$ is always a manager, and Corollaries 2 and 3. □

Corollary 4 is illustrated in Figure 5. To understand it, notice that the least knowledgeable humans are always workers, while the most knowledgeable humans are always managers. Thus, positive assortative matching implies that irrespective of AI, (i) the size of the smallest firm is $n(0)$ times the knowledge of the worst manager, and (ii) the size of the biggest firm is equal to the largest span of control in the economy. Corollary 4 then follows directly from Corollaries 2 and 3.

4.3 Beyond Displacement: The Effects of AI on Non-Displaced Workers and Managers

We now analyze the effects of AI on the productivity and the span of control of non-displaced workers and managers (in the sense that they are workers and managers both pre- and post-AI, respectively). Although the results of the previous section might suggest that AI leaves some firms untouched (implying that the productivity and span of control of non-displaced workers and managers might stay the same), this is not the case. The reason is that AI's introduction affects all the matches in the economy.

Proposition 4. *AI has the following effects on the productivity of non-displaced workers and the span of control of non-displaced managers:*

- *If* $z_{\mathrm{AI}} \in \mathrm{int}W$*, then:*
  - *The productivity of* $z \in W^* \subset W$ *strictly decreases.*
  - *The span of control of* $z \in M \subset M^*$ *strictly increases if* $e(z) < z_{\mathrm{AI}}$*, and strictly decreases if* $e(z) > z_{\mathrm{AI}}$*.*
- *If* $z_{\mathrm{AI}} \in \mathrm{int}M$*, then:*
  - *The productivity of* $z \in W \subset W^*$ *strictly increases if* $z < e(z_{\mathrm{AI}})$*, and strictly decreases if* $z > e(z_{\mathrm{AI}})$*.*
  - *The span of control of* $z \in M^* \subset M$ *strictly increases.*
- *If* $z_{\mathrm{AI}} = \hat{z}$*, then:*
  - *The productivity of* $z \in W^* = W$ *decreases (strictly so for all* $z \neq 0$*).*
  - *The span of control of* $z \in M^* = M$ *increases (strictly so for all* $z \neq 1$*).*

Proof. See Appendix C. □

For intuition, suppose first that AI has the knowledge of a pre-AI worker. In this case, the knowledge required to become a manager decreases, worsening the pool of managers available for non-displaced workers. At the same time, the match of the best managers worsens as they are now supervising production by AI (while pre-AI, they were managing humans more knowledgeable than AI), but the match of all the other managers improves because the worst workers are now managed by the newly appointed managers.

Similarly, when AI has the knowledge of a pre-AI manager, the match of the worst workers improves because they are now assisted by AI (while pre-AI, they were assisted by humans less knowledgeable than AI). However, the knowledge of the best workers increases, improving the pool of available workers for non-displaced managers and, therefore, leaving a worse pool of managers for the non-displaced workers that were originally managed by humans more knowledgeable than AI.

Finally, when AI has the knowledge of both a pre-AI worker and a pre-AI manager, there is no human displacement across occupations. However, the match of the least knowledgeable workers worsens (as they are now assisted by AI), while the match of the most knowledgeable managers improves (as they are now supervising AI). The first result then implies that the pool of workers available for the remaining non-displaced managers improves with AI, while the second result implies that the pool of available managers for the remaining non-displaced workers worsens with AI.

4.4 Labor Income

In this section, we analyze the effects of AI on labor income. We begin with the following result:

Lemma 1. *Total output and total labor income increase with AI.*

The proof of this result is intuitive and relatively straightforward, so we provide it here as part of the main text.

Proof. The result that total output increases follows because (i) the First Welfare Theorem holds in this setting, and (ii) AI expands the production possibility frontier. The result that total labor income also increases with AI follows from two observations. First, if all compute is assigned to independent production, total labor income does not change with AI (as AI does not interact with humans in the workplace). Second, capital income is equal to μz_{AI} regardless of how compute is used, so:

$$\text{Total output post-AI} = \text{Total labor income post-AI} + \mu z_{\text{AI}}$$

Consequently, given that the AI equilibrium is efficient, unique, and does not allocate all compute to independent production, it must be that:

$$\text{Total output post-AI} > \text{Total labor income pre-AI} + \mu z_{\text{AI}}$$

Hence, total labor income must be strictly larger post-AI than pre-AI. □

We now turn to analyzing the distributional effects of AI. Given that each agent's wage is her marginal product, understanding whose wage increases or decreases with AI amounts to understanding which humans are complemented by the technology (in the sense that their marginal product increases with AI) and which humans are substituted by it (in the sense that their marginal product decreases with AI). Note that an agent's marginal product (defined as the output increase of introducing such an agent into the economy) is equal to her productivity (defined as the agent's expected output) only in the case of independent producers. This is because whenever an agent is introduced as either a worker or a manager, her introduction affects the output of other agents in the economy through changes in matching and firm composition.

Disentangling the distributional effects of AI is non-trivial due to the existence of two potentially countervailing forces: On the one hand, AI changes the composition of firms and, therefore, the quality of matches. On the other hand, by mimicking humans with knowledge z_{AI}, AI changes the relative scarcities of different knowledge levels, affecting how each firm's output is divided between workers and managers.

We first show that *if* a human with knowledge $z < z_{\text{AI}}$ wins from AI's introduction, then all those humans with knowledge $z' < z$ must also be winners. Similarly, *if* a human with knowledge $z > z_{\text{AI}}$ wins from AI's introduction, then all those humans with knowledge $z' > z$ are better off after the advent of AI. Given that a human with knowledge $z = z_{\text{AI}}$ is always worse-off, this implies that the winners from AI are necessarily located at the extremes of the knowledge distribution:

Lemma 2. *Define $\Delta(z) \equiv w^*(z) - w(z)$. Then $\Delta(z_{\text{AI}}) < 0$ and:*

- *If $\Delta(z) > 0$ for some $z \in [0, z_{\text{AI}}]$, then $\Delta(z') > 0$ for all $z' \in [0, z]$.*
- *If $\Delta(z) > 0$ for some $z \in [z_{\text{AI}}, 1)$, then $\Delta(z') \geq 0$ for all $z' \in [z, 1]$, with strict inequality for all $z' \in [z, 1)$.*

Proof. See Appendix C. □

Lemma 2 is subtle because, even though $\Delta(z_{\text{AI}}) < 0$, $\Delta(z)$ need not be v-shaped around z_{AI}, as illustrated in Figure 6. This is because AI can sometimes worsen the matches of the best managers and improve the matches of the worst workers. Nevertheless, Lemma 2 states that in the interval $[0, z_{\text{AI}}]$, the $\Delta(z)$ function can only cross zero from above, while in the interval $[z_{\text{AI}}, 1]$, the $\Delta(z)$ function can only cross zero from below. Intuitively, this is because wages decrease the most at z_{AI}, and AI is always used as the best worker and/or the worst manager. Hence, the least and most knowledgeable humans (matched with the worst managers and the best workers, respectively) are the ones whose share of output increases the most after the reduction of wages at z_{AI}.

The next step is characterizing whether (and when) winners exist below or above z_{AI}. For the following proposition, recall that $\hat{z}$ is the knowledge level of the worst pre-AI manager (which is also the best pre-AI worker):

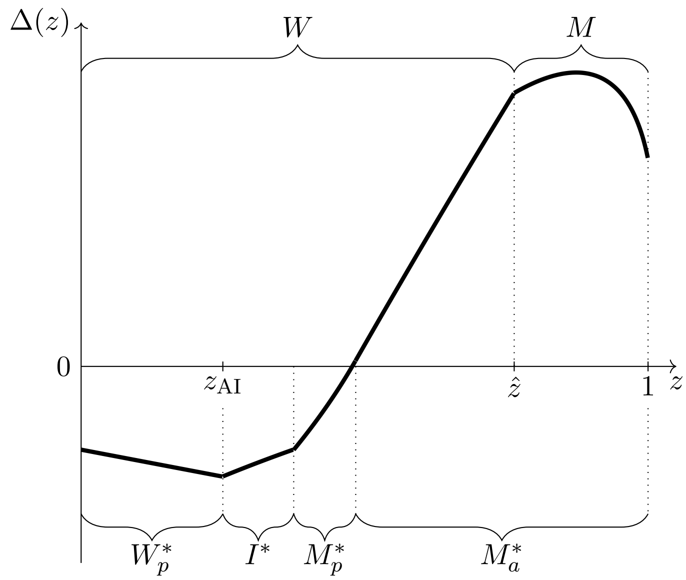

Figure 6: An Illustration of $\Delta(z)$ Function when $z_{\text{AI}} \in \text{int}W$

Notes. Distribution of knowledge: $G(z) = z$. Parameter values: $h = 1/2$ ($< h_0 = 3/4$), $z_{\text{AI}} = 1/4$. Thus, $\Delta(0) = -0.1695$, $\Delta(1) = 0.4222$, $\max_z \Delta(z) = 0.59823$ (with $\text{argmax}_z \Delta(z) = 0.87331$), and $\min_z \Delta(z) = -0.2243$ (with $\text{argmin}_z \Delta(z) = z_{\text{AI}} = 1/4$). See Section 3 of the Online Appendix for more details on the construction of this figure.

Proposition 5.

(i) There always exists z strictly greater than z_{AI} such that $\Delta(z) > 0$.

(ii) There exists z strictly smaller than z_{AI} such that $\Delta(z) > 0$ if and only if $z_{\text{AI}} > \bar{z}_{\text{AI}}$, where $\bar{z}_{\text{AI}} \in (0, \hat{z})$.

Proof. See Appendix C. □

According to Proposition 5, both extremes of the knowledge distribution can potentially benefit from AI's introduction. However, the proposition also reveals an interesting asymmetry between the effects of AI on the lowest and highest segments of the knowledge distribution. Indeed, while the most knowledgeable humans always benefit from the introduction of AI, the least knowledgeable humans benefit only when AI's knowledge is sufficiently high.

To understand this asymmetry, let us consider how changes in AI's knowledge affect the wages of the least and most knowledgeable humans in the post-AI world. When z_{AI} increases, (i) AI is a "closer substitute" to the most knowledgeable human, and (ii) the most knowledgeable human is better matched. While the first effect puts downward pressure on the wage of the most knowledgeable human, the second effect puts opposite pressure on her wage. Moreover, the latter effect is stronger the smaller is h, because a smaller helping cost allows each manager to supervise more workers. The assumption that $h < h_0$ then guarantees that the second effect more than compensates for the first.[21]

In contrast, when z_{AI} decreases, (i) AI is a closer substitute to the least knowledgeable human, and (ii) the least knowledgeable human is worse matched. In this case, both effects put downward pressure on her wage, explaining why the least knowledgeable human benefits from AI only when z_{AI} is sufficiently high.

5 Discussion and Extensions

5.1 Artificial Superintelligence and the Case $h \geq h_0$

For simplicity, we have assumed that $z_{\text{AI}} < 1$ because the equilibrium is discontinuous at $z_{\text{AI}} = 1$. Nevertheless, all of our results continue to hold when $z_{\text{AI}} = 1$ (noting that, in this case, AI has the knowledge of a pre-AI manager), except that most knowledge humans no longer benefit from AI's introduction.

The intuition for why this is so is closely connected to the equilibrium discontinuity at $z_{\text{AI}} = 1$. Indeed, when z_{AI} is arbitrarily close but strictly less than 1, the most knowledgeable humans benefit from AI as they leverage their knowledge by supervising AI (see Section 4.4). In contrast, when $z_{\text{AI}} = 1$, then AI supervises all humans. This implies that $w^*(z) = 1 - h(1 - z)$ (since $r^* = 1$), so the

[21] In section 5.1, we show that this is no longer the case when $h \geq h_0$, and hence, in this case, the most knowledgeable humans in the population can lose from AI.

most knowledgeable humans only earn $w^*(1) = 1 < w(1)$ (where the inequality follows because 1 was a manager pre-AI).

Also, for brevity, we have focused on the case $h < h_0$. In Section 4 of the Online Appendix, we show that all our results remain unchanged when $h \geq h_0$, with the following two exceptions.

First, AI can have the knowledge of a pre-AI independent producer, in which case it does not affect labor outcomes. Intuitively, when $z_{\text{AI}} \in I$, the first unit of compute is allocated in equilibrium to independent production (as prescribed by the pre-AI equilibrium), so it does not affect two-layer firms in any way. Hence, the same argument applies to the second unit, the third unit, and so on, so all units of compute end up being allocated to independent production. Consequently, AI does not affect wages (as the $w(z_{\text{AI}}) = z_{\text{AI}}$ already in the pre-AI equilibrium) nor other labor outcomes because it does not interact with humans in the workplace.

Second, the wage of all humans who are more knowledgeable than AI can decrease with AI when $h \geq h_0$. In particular, this occurs when $z_{\text{AI}} \in \text{int}M$. Intuitively, recall that a more knowledgeable AI has two opposing effects on the wages of the most knowledgeable humans: On the one hand, it reduces their wages, as AI is more similar to them; on the other hand, it improves their matches. While the second effect always dominates when $h < h_0$ (as mentioned in Section 4.4), this is not necessarily the case when $h \geq h_0$, since each manager now supervises a small number of workers.

5.2 Labor Income Inequality and Polarization

In section 4.4, we described the winners and losers resulting from AI's introduction but refrained from discussing the effects of AI on labor income inequality. This omission stems from the complexity surrounding an accurate measurement of wage inequality in the context of our model.

In particular, two issues arise when evaluating AI's effects on wage inequality. First, there are different valid metrics, and the effects of AI on inequality often depend on the particular metric chosen. For instance, while empirical studies often focus on the ratio of the wages of different income brackets (Acemoglu and Autor, 2011), the theoretical literature has also suggested considering absolute wage difference (Antràs et al., 2006). Second, measuring inequality requires specifying the scope of interest (e.g., within-worker, within-manager, or overall labor income inequality). Although all three are relevant, within-worker and within-manager inequality pose challenges due to AI-induced occupational displacement. Indeed, this displacement can significantly alter the identity of the best workers and the worst managers, resulting in mechanical changes in inequality.

To circumvent these issues, here we take a conservative approach: (i) we state the results that do not depend on the metric used (i.e., ratios or absolute differences), and (ii) we focus on the effects of AI on overall wage inequality (measured as either the difference or the ratio of the wages of the most and least knowledgeable humans in the population).

When $h < h_0$, the wages at the top of the knowledge distribution always increase, and the wages at the bottom of the knowledge distribution increase if and only if $z_{\mathrm{AI}} > \bar{z}_{\mathrm{AI}} \in (0, \hat{z})$. Hence, AI unambiguously increases human labor inequality when $z_{\mathrm{AI}} \leq \bar{z}_{\mathrm{AI}}$ (while when $z_{\mathrm{AI}} > \bar{z}_{\mathrm{AI}}$, the effect depends on the parameters and the metric used). When $h \geq h_0$, AI increases inequality if $z_{\mathrm{AI}} \in \mathrm{int}W$ but decreases inequality if $z_{\mathrm{AI}} \in \mathrm{int}M$. This follows from the fact that when $z_{\mathrm{AI}} \in \mathrm{int}W$, the wages at the top and bottom of the knowledge distribution increase and decrease, respectively, while the opposite occurs when $z_{\mathrm{AI}} \in \mathrm{int}M$.

Finally, one recurring theme in the literature is the effect of automation on wage polarization (meaning that the negative effects of AI on wages are concentrated in the middle of the income distribution). Our results show that wage polarization may or may not arise depending on z_{AI} and h. This can be seen, for instance, in Figure 2, which depicts the pre- and post-AI equilibrium for two different values of z_{AI}, $z_{\mathrm{AI}} = 1/4$ in panel (a) and $z_{\mathrm{AI}} = 3/4$ in panel (b), when human knowledge is uniformly distributed and $h = 1/2$.

In particular, panel (a) shows that humans around the median knowledge benefit from AI's introduction as they become managers of relatively cheap workers (instead of working for average managers pre-AI). Hence, in this case, there is no wage polarization. In contrast, wage polarization does arise in panel (b) as both extremes of the knowledge distribution benefit from AI, while those in the middle of the distribution are worse off.

5.3 The Role of Abundance of Compute Relative to Human Time

Until now, we have considered the case in which compute is large relative to human time. To understand the role of this assumption, in Section 5 of the Online Appendix we characterize the equilibrium in the opposite extreme in which compute μ is arbitrarily small.

In this case, the nature of the displacement still depends on AI's knowledge relative to the pre-AI equilibrium: When AI has the knowledge of a pre-AI worker, AI displaces humans from routine to managerial work. In contrast, when AI has the knowledge of a pre-AI manager, humans are displaced in the opposite direction. This also implies that our results regarding the distributions of firm size, productivity and span of control continue to hold.

However, when μ is small, it is no longer the case that AI is the best worker and/or the worst manager. This difference has two important implications. The first one relates to the quality of matches of those who continue to be workers and/or managers. The second relates to who wins and who loses with the introduction of AI.

Regarding the quality of matches, it is no longer the case that the productivity of all workers decreases when $z_{\mathrm{AI}} \in \mathrm{int}W$, or that all managers supervise larger firms when $z_{\mathrm{AI}} \in \mathrm{int}M$. In particular, there is now a set of workers—those with knowledge above z_{AI}—who have better managers post-AI than pre-AI when $z_{\mathrm{AI}} \in \mathrm{int}W$. Similarly, there is a set of managers—those with knowledge below z_{AI}—who manage smaller firms when $z_{\mathrm{AI}} \in \mathrm{int}M$. Intuitively, these differences arise because, when compute is abundant, there are no workers with knowledge above z_{AI} nor managers with knowledge below z_{AI}.

With regards to labor income, when compute is small, the winners from AI are not necessarily at the extremes of the knowledge distribution. The reason is that AI can now be among the least knowledgeable workers, so its introduction might increase the wages of those in the middle of the knowledge distribution, thus harming the most knowledgeable humans (who manage them).[22] This intuition highlights the key role that abundant compute plays in our analysis of the effects of AI on labor income.

5.4 What if Compute Exceeds its Potential Applications?

In the baseline setting, we assumed that compute is abundant relative to human time but scarce relative to its potential applications. This has two noteworthy implications. First, the equilibrium rental rate of compute is equal to AI's knowledge. Second, in the AI equilibrium, some humans are still doing routine work.

In Section 6 of the Online Appendix, we relax the assumption that compute is scarce relative to its applications. In particular, we consider the case in which (i) there is a large but finite amount Q of potential problems to be solved, and (ii) the compute is abundant not only relative to human time but also relative to Q.[23]

In this case, and in contrast to our baseline setting, the equilibrium price of compute is zero, and all agents that are less knowledgeable than AI are unemployed. Organizations, however, still have a hierarchical structure: All problems are initially attempted by AI, and all humans who are more knowledgeable than AI specialize in solving problems that AI cannot solve. Intuitively, only the time of the humans that are more knowledgeable than AI is scarce, so only they get rewarded for their work.

Nevertheless, our results concerning occupational displacement still hold: If AI has the knowledge of a pre-AI worker, it still shifts humans from routine to managerial work, while if AI has the knowledge of a pre-AI manager, it still reduces the number of humans doing managerial work. In this extension, however, AI displaces everyone who is less knowledgeable than AI to unemployment.

Moreover, our results concerning the effects of AI on the distributions of firms' size, productivity, and span of control continue to hold with the following minor change: The size of the smallest two-layer firm can increase or decrease with AI when $z_{\mathrm{AI}} \in \mathrm{int} W$ (while in the baseline, it always decreases), and necessarily increases when $z_{\mathrm{AI}} = \hat{z}$ (in the baseline, its size remains unchanged). In-

[22]Similarly, AI can now be among the most knowledgeable managers, so their introduction can increase the wages of those in the middle of the knowledge distribution, harming the least knowledgeable humans (who are matched with them).

[23]Note that the equilibrium of the baseline setting is still the equilibrium of this alternative model in the region of the parameter space where μ is abundant relative to time but not to Q.

tuitively, the difference arises because, in the baseline setting, the least knowledgeable worker has knowledge $z = 0$ both pre- and post-AI. In contrast, in this extension, AI is both the least and most knowledgeable worker, so AI increases the knowledge of the workers of the smallest/least productive firm.[24]

Regarding the quality of matches, the results about non-displaced workers are vacuous because the set of non-displaced workers is empty in this case. In turn, the results concerning the size of the firms being supervised by the managers who are not displaced remain unchanged. Finally, regarding labor income, the wages of everyone who is less knowledgeable than AI drop to zero. Hence, the winners, if there are any, are always at the high end of the knowledge distribution.

5.5 The Effects of AI on the Labor Share

Pre-AI, the labor share is equal to one because labor is the sole input in production. Interestingly, whether AI lowers the labor share or not depends on whether the set of potential applications for AI is larger or smaller than the amount of compute in the economy.

In the baseline setting—where the unlimited applications of compute imply that $r^* = z_{\text{AI}}$—the labor share decreases when moving from the pre-AI to the post-AI equilibrium. Moreover, it converges to zero as $\mu \to \infty$ since, when compute is abundant relative to human time, any additional unit of compute is exclusively allocated to independent production. As a result, further increases in μ raise capital income but not labor income.

This result dramatically changes when there is a large but finite amount Q of potential problems. In this case, the price of compute is 0 when μ is large enough (as discussed in the previous subsection). This has two implications. First, the post-AI labor share can be equal to one, just like pre-AI. Second, in a world with finitely many problems, the labor share is not monotone in μ: it decreases with μ while compute is abundant relative to time but not to problems and then increases with μ as compute catches up with the number of problems. We leave a more detailed study of this race between compute and its potential applications for future research.

5.6 AI vs. Improvements in Communication Technologies vs. Globalization

We now discuss how the effects of AI compare with the effects of two other important "shocks" previously studied in the literature: Improvements in communication technologies and the possibility of forming international teams (i.e., globalization).

Consider first a reduction in communication costs h, starting from $h < h_0$. As explained by Gari-

[24] As a result, in this extension, the smallest firm has a worse manager but better workers post-AI than pre-AI when $z_{\text{AI}} \in \text{int}W$ (while in the baseline, it only has a worse manager). Similarly, in this extension, the smallest firm has the same manager but better workers post-AI than pre-AI when $z_{\text{AI}} = \hat{z}$ (while in the baseline it has the same workers and manager).

cano and Rossi-Hansberg (2004), this decrease can be attributed to the widespread adoption of e-mail, cellular phones, and wireless technology during the late 1990s. The impacts of this change differ from those caused by AI.

Indeed, a reduction in h naturally allows managers to supervise larger teams, thus reducing the demand for managers. The latter has two immediate implications: There is occupational displacement from managerial to routine work (as the worst managers switch to worker roles), and all workers become better matched (and hence more productive). In contrast, the displacement generated by AI goes in the opposite direction when $z_{\mathrm{AI}} \in \mathrm{int}W$. Moreover, the introduction of AI (i) reduces the productivity of the best non-displaced workers irrespective of AI's knowledge (and the productivity of all workers when $z_{\mathrm{AI}} \in \mathrm{int}W$), and (ii) decreases the size of the firms supervised by the best managers when $z_{\mathrm{AI}} \in \mathrm{int}W$.

Second, consider the effects of globalization. Antràs et al. (2006) consider a two-country model in which countries only differ in their knowledge distributions. In particular, one country, the North, has a distribution of knowledge with a relatively high mean, while the other country, the South, has a distribution of knowledge with a relatively low mean. They show that allowing the formation of international teams (or "offshoring") shifts humans from routine to managerial work in the North, while it shifts humans from managerial to routine work in the South.

Given the similarity between the occupational displacement effects of AI and offshoring, one might conjecture that the effects of AI when $z_{\mathrm{AI}} \in \mathrm{int}W$ are qualitatively similar to the effects of offshoring from the North's perspective and that the effects of AI when $z_{\mathrm{AI}} \in \mathrm{int}M$ are qualitatively similar to the effects of offshoring from the South's perspective. This, however, is not the case. For instance, while offshoring increases the productivity of the best workers in the North (Antràs et al., 2006, Proposition 1), AI reduces the productivity of all workers when $z_{\mathrm{AI}} \in \mathrm{int}W$. Similarly, while offshoring decreases the span of control of all southern managers (Antràs et al., 2006, Proposition 1), AI increases the span of control of all managers when $z_{\mathrm{AI}} \in \mathrm{int}M$.

Intuitively, the key difference between AI and offshoring is AI's capacity to operate at scale (in the two senses discussed in Section 2.2). This implies that, although both globalization and AI induce the best northern workers pre-globalization/pre-AI to switch to managerial roles, the best northern managers switch to supervising the best non-displaced northern workers in the case of globalization, while they switch to supervising AI in the case of artificial intelligence.

Similarly, when $z_{\mathrm{AI}} \in \mathrm{int}M$, both globalization and AI induce the worst southern managers to switch to routine work, improving the overall pool of southern workers. However, in the case of globalization, the best southern workers are matched with the best northern managers, leaving a worse pool of workers for the non-displaced southern managers. In contrast, in the case of AI, it is the worst southern workers who end up being supervised by AI, leaving the best southern workers for the non-displaced southern managers.

5.7 Examples of AI Worker, Independent Producer, and Manager

We finish this section by providing real-world examples in which AI acts as a "worker," a "manager," or an "independent producer" in the sense of our framework.

To understand how AI can play the role of a "worker," consider legal services. Law firms are often hierarchical, partly reflecting the division of labor within the organization: Individuals with less knowledge and experience, such as associates and paralegals, execute tasks that are less knowledge-intensive, like reviewing old cases and searching records, among others, allowing senior lawyers and partners to use their expertise to solve the most important exceptions (Garicano and Hubbard, 2007).

These less knowledge-intensive tasks are the ones being increasingly automated by AI. For instance, according to a recent survey by Thomson Reuters (2021), 64% of law firms are deploying AI for legal research, 47% for document review, and 38% for document automation. In fact, as evidenced by the following quote from the American Bar Association, such automation has been occurring for some time:[25]

> AI is already having an impact on firms in the U.S. and around the world. Robots or machines are being utilized to do tedious, time-consuming tasks like collecting data, searching records, going through old cases, verifying facts, etc.— work currently done by junior lawyers and paralegals.

The legal services industry also provides vivid examples of how AI can be used as an "independent producer." Indeed, tax software like TurboTax, H&R Block, and SprintTax provide legal advice to millions of Americans every year—in most circumstances without any human assistance. More sophisticated AI independent producers can be found in U.S. cities like San Francisco, Phoenix, and Austin, where companies like Waymo and Cruise offer driverless taxi services ("robotaxis").[26]

Finally, to understand how AI can play the role of a "manager," consider its deployment in the customer service industry. Brynjolfsson et al. (2023) empirically study the effects of introducing a generative AI-based "conversational assistant" designed to help technical support agents of a Fortune 500 software company. As the authors emphasize, this AI system was designed to help these agents—for instance, by providing suggested answers to problems in real-time—rather than replace them. They show that AI increases the productivity of the least experienced and less knowledgeable

[25]"Legal Tech Gurus Forecast how AI will Impact your Practice," American Bar Association, July 2017. https://www.americanbar.org/news/abanews/publications/youraba/2017/july-2017/artificial-intelligence-and-the-future-of-law-practice/ (accessed February 5, 2024)

[26]On October 26, 2023, Cruise halted its US robotaxi services nationwide (grounding its fleet of 400 robotaxis) after the California Public Utility Commission suspended its permit for service after an incident in San Francisco. Meghan Bobrowsky, "GM's Cruise Says U.S. Is Investigating Driverless Car's Collision With Pedestrian." *The Wall Street Journal*, January 25, 2024. https://www.wsj.com/business/autos/gms-cruise-confirms-doj-investigation-of-driverless-car-incident-b249c13b?st=v4m4h0mu9udh7ct (accessed February 12, 2024).

agents by 35%, but that the effects on more experienced and knowledgeable workers were minimal.[27]

In the context of our model, we can think of this experiment as providing an AI "manager" who can help, in real-time, the technical support agents.[28] In particular, in contrast to the legal application described above, AI's role, in this case, is not to free up time by handling the easiest problems—this should create productivity gains across the board—but rather to help agents resolve the most difficult problems they encounter in their work. Brynjolfsson et al. (2023) findings, moreover, are consistent with the results of Section 4.4: The least knowledgeable technical support agents are better off from the introduction of AI, while the most knowledgeable ones—whom the AI technology closely imitates—likely suffered a reduction in their current pay, as the firm's practice is to calculate their bonuses relative to other agents' performance.

6 Final Remarks

The transformative impact of Artificial Intelligence (AI) is undeniable, yet its precise implications for the nature of work are controversial. In this paper, we introduce a novel framework to examine the impact of AI on labor outcomes. Our approach is motivated by the observation that AI holds the potential to automate knowledge work and that a fundamental problem of society is the efficient utilization of knowledge.

Our framework provides novel insights about how AI influences labor outcomes through its effects on firm composition and organizational structure. Moreover, it has the attractive feature that the role that AI plays in the economy—and hence its effects on labor outcomes—is determined endogenously as a function of the knowledge of the population, the level of information and communication technologies, and the degree of advancement of AI.

This paper opens up several avenues for future research. For example, we believe that our approach may be fruitfully applied to investigate (i) the effects of AI on capital markets, (ii) firms' incentives to develop AI, (iii) the implication and effectiveness of different retraining schemes, and (iv) the organization of international trade, offshoring and, more generally, economic development.

[27] Noy and Zhang (2023) finds similar results in a different experimental setting involving midlevel professional writing tasks.

[28] The technical support agents did have a human manager before the introduction of AI. However, they could not consult with her in real-time; the manager's role was to provide training and weekly feedback.

Appendix

A The Pre-AI Equilibrium: Complete Characterization

In this Appendix, we provide the full characterization of the equilibrium without AI, including the proof of Proposition 1. As noted in the main text, the equilibrium of an economy without AI was first described in general by Fuchs et al. (2015).

A.1 Proof of Proposition 1

Proposition 1 follows from Lemmas A.1, A.2, and Corollary A.1 (plus the fact that the First Welfare Theorem implies that the competitive equilibrium is efficient in our setting).

Lemma A.1. *In the absence of AI, there is a unique surplus maximizing allocation:*

- *When $h > h_0 \in (0,1)$, then $W = [0, \underline{z}]$, $I = (\underline{z}, \bar{z})$, and $M = [\bar{z}, 1]$. Moreover, $f(z;\bar{z})$ is strictly increasing and given by $\int_{\bar{z}}^{f(z;\bar{z})} dG(u) = \int_0^z h(1-u)dG(u)$, and the cutoffs $0 < \underline{z} < \bar{z} < 1$ satisfy:*

$$\frac{1}{h} - \bar{z} = \int_{\bar{z}}^{1} n\big(e(u;\bar{z})\big)du \quad \text{and} \quad f(\underline{z};\bar{z}) = 1 \quad \text{where } e(z;\bar{z}) = f^{-1}(z;\bar{z}) \tag{4}$$

- *When $h \leq h_0$, then $W = [0, \hat{z}]$, $I = \emptyset$, and $M = [\hat{z}, 1]$. Moreover, $f(z;\hat{z})$ is strictly increasing and given by $\int_{\hat{z}}^{f(z;\hat{z})} dG(u) = \int_0^z h(1-u)dG(u)$ and the cutoff $\hat{z} \in (0,1)$ satisfies $f(0;\hat{z}) = \hat{z}$, $f(\hat{z};\hat{z}) = 1$, and:*

$$\begin{aligned} \tfrac{1}{h} - \hat{z} &> \textstyle\int_{\hat{z}}^{1} n(e(z;\hat{z}))dz \text{ for all } h < h_0 \\ \tfrac{1}{h} - \hat{z} &= \textstyle\int_{\hat{z}}^{1} n(e(z;\hat{z}))dz \text{ when } h = h_0 \end{aligned} \tag{5}$$

Proof. For the proof see Fuchs et al. (2015, Lemma 2). □

Lemma A.2. *In the absence of AI, the equilibrium wage function is given by:*

- *When $h \geq h_0$:*

$$w(z) = \begin{cases} f(z;\bar{z}) - \dfrac{1}{n(z)} w(f(z;\bar{z})) & \text{if } z \in W \\ z & \text{if } z \in I \\ \bar{z} + \int_{\bar{z}}^{z} n\big(e(u;\bar{z})\big)du & \text{if } z \in M \end{cases}$$

- *When $h < h_0$:*

$$w(z) = \begin{cases} f(z;\bar{z}) - \dfrac{1}{n(z)} w(f(z;\bar{z})) & \text{if } z \in W \\ \frac{1}{1+n(\hat{z})}\left\{n(\hat{z}) - \int_{\hat{z}}^{1} n\big(e(u;\hat{z})\big)du\right\} + \int_{\hat{z}}^{z} n\big(e(u;\hat{z})\big)du & \text{if } z \in M \end{cases}$$

Proof. See the Online Supplement of Fuchs et al. (2015, specifically pp. 1–4). □

Corollary A.1. *Irrespective of whether $h \gtreqless h_0$, the equilibrium wage function $w(z)$ is continuous, strictly increasing, and (weakly) convex. Moreover, it satisfies:*

- $w(z) > z$ *for* $z \in W$ *(except possibly at* $z = \sup W$*), and* $w'(z) \in (0, 1)$*, and* $w''(z) > 0$ *for* $z \in \text{int}W$*.*
- $w(z) = z$ *for all* $z \in I$*.*
- $w(z) > z$ *for* $z \in M$ *(except possibly at* $z = \inf M$*), and* $w'(z) > 1$*, and* $w''(z) > 0$ *for* $z \in \text{int}M$*.*

Proof. Consider first $h \geq h_0$. From Lemma A.2, it is immediate that $\lim_{z \downarrow \bar{z}} w(z) = w(\bar{z}) = \bar{z}$. Note than that for $z \in \text{int}M$, then $w'(z) = n(e(z; \bar{z})) > 1$, which also implies that $w''(z) > 0$ as both $n(z)$ and $e(z; \bar{z})$ are strictly increasing in their arguments. Given that $\lim_{z \downarrow \bar{z}} w(z) = \bar{z}$, the previous results then imply that $w(z) > z$ for all $z \in (\bar{z}, 1]$.

On the other hand, $\lim_{z \uparrow \underline{z}} w(z) = 1 - w(1)h(1 - \underline{z}) = \underline{z} = \lim_{z \downarrow \underline{z}} w(z)$, where we are using the fact that $f(\underline{z}; \bar{z}) = 1$ and that $w(1) = \int_{\bar{z}}^{1} n\big(e(u; \bar{z})\big)du + \bar{z} = \frac{1}{h}$ (due to condition (4)). Note, moreover, that for $z \in \text{int}W$, then $w'(z) = hw(f(z; \bar{z})) > 0$, which immediately implies that $w''(z) > 0$, since both $w(z)$ and $f(z; \bar{z})$ are strictly increasing in their arguments. That $w(z) > z$ for all $z \in [0, \underline{z})$ then follows because $w(z)$ is strictly increasing and convex and $\lim_{z \uparrow \underline{z}} w(z) = \underline{z}$. Finally, that $w'(z) \in (0, 1)$ comes from the fact that:

$$w'(z) = hw(f(z; \bar{z})) = h \left[\frac{f(z; \bar{z}) - w(z)}{h(1 - z)} \right] < 1$$

where the second-to-last equality comes from the firms' zero-profit condition, and the last inequality because $f(z; \bar{z}) \leq 1$ and $w(z) > z$ when $z \in \text{int}W$.

Now consider $h \leq h_0$. Given that $f(\hat{z}; \hat{z}) = 1$, from Lemma A.2, we have that:

$$\lim_{z \uparrow \hat{z}} w(z) = 1 - \tfrac{1}{n(\hat{z})} w(1) \quad \text{and} \quad \lim_{z \downarrow \hat{z}} w(z) = \tfrac{1}{1 + n(\hat{z})} \Big\{ n(\hat{z}) - \textstyle\int_{\hat{z}}^{1} n\big(e(u; \hat{z})\big)du \Big\}$$

Note then that $w(1) = \lim_{z \downarrow \hat{z}} w(z) + \int_{\hat{z}}^{1} n\big(e(u; \hat{z})\big)du$. Combining the latter with the expression for $\lim_{z \uparrow \hat{z}} w(z)$ above and rearranging terms yields $\lim_{z \uparrow \hat{z}} w(z) = \lim_{z \downarrow \hat{z}} w(z)$. The proof that $\lim_{z \downarrow \hat{z}} w(z) > \hat{z}$, in turn, can be found in Fuchs et al. (2015, Online Supplement, p. 4). Finally, the proofs for the remaining properties of $w(z)$ (i.e., their monotonicity and convexity, among others) follow the exact same logic as in the case when $h > h_0$. □

B The AI Equilibrium: Complete Characterization

In this Appendix, we provide a complete characterization of the AI equilibrium. As noted in the main text, we focus on $h < h_0$ (see Section 4 of the Online Appendix for $h \geq h_0$). The proof of Proposition 2 is a direct implication of the results that follow.

B.1 Equilibrium Characterization

We begin the characterization with the following set of results:

Lemma B.1. *Any equilibrium with AI has the following features:*

- *Some compute must be allocated to independent production:* $\mu_i^* > 0$.
- *The price of compute is equal to AI's knowledge:* $r^* = z_{\mathrm{AI}}$.
- *Occupational stratification:* $W^* \preceq I^* \preceq M^*$.
- *No worker better than AI; no manager worse than AI:* $W^* \preceq \{z_{\mathrm{AI}}\} \preceq M^*$.
- *Positive assortative matching:* $f^* : W_p^* \to M_p^*$ *is strictly increasing and* $W_a^* \preceq W_p^*$ *and* $M_p^* \preceq M_a^*$.

Proof. • *Some compute must be allocated to independent production.*— This result follows because compute is abundant relative to human time. Hence, there are not enough humans to interact with AI inside two-layer organizations.

• *The price of compute is equal to AI's knowledge.*— This follows because the single-layer firms using AI must obtain zero profits.

• *Occupational stratification.*— Notice that the First Welfare Theorem holds in our setting. Hence, a competitive equilibrium must be efficient. Occupational stratification then follows because any surplus maximizing allocation must satisfy it. The proof of this last result is analogous to the proof of Lemma 1 in Fuchs et al. (2015).

• *No worker better than AI; no manager worse than AI.*— This result follows from occupational stratification and the fact that some compute must necessarily be used for independent production.

• *Positive assortative matching.*— The emergence of positive assortative matching—which follows from the supermodularity of the profits of two-layer organizations—is proven in Eeckhout and Kircher (2018, Proposition 1, p. 94) in a more general setting that encompasses ours. Positive assortative matching then implies that the matching function is strictly increasing and that $W_a^* \preceq W_p^*$ and $M_p^* \preceq M_a^*$ (since no worker is better than AI and no manager is worse than AI). □

The next corollary is a direct implication of Lemma B.1:

Corollary B.1. *An equilibrium allocation must take one of the following four potential configurations:*

- ***Type 1 configuration:***

$$W_a^* = \emptyset,\ W_p^* = [0, z_{\mathrm{AI}}],\ I^* = (z_{\mathrm{AI}}, \underline{z}_1^*),\ M_p^* = [\underline{z}_1^*, \bar{z}_1^*],\ M_a^* = [\bar{z}_1^*, 1],\ \textit{where } z_{\mathrm{AI}} < \underline{z}_1^* \le \bar{z}_1^* \le 1$$

$$\textit{So } \mu_w^* = \int_{\bar{z}_1^*}^1 n(z_{\mathrm{AI}}) dG(z),\ \mu_m^* = 0,\ \mu_i^* = \mu - \mu_w^*$$

- ***Type 2 configuration:***

$$W_a^* = [0, \underline{z}_2^*],\ W_p^* = [\underline{z}_2^*, z_{\mathrm{AI}}],\ I^* \subseteq \{z_{\mathrm{AI}}\},\ M_p^* = [z_{\mathrm{AI}}, \bar{z}_2^*],\ M_a^* = [\bar{z}_2^*, 1],\ \textit{where } 0 \le \underline{z}_2^* \le z_{\mathrm{AI}} \le \bar{z}_2^* \le 1$$

$$\textit{So } \mu_w^* = \int_{\bar{z}_2^*}^1 n(z_{\mathrm{AI}}) dG(z),\ \mu_m^* = \int_0^{\underline{z}_2^*} n(z)^{-1} dG(z),\ \mu_i^* = \mu - \mu_w^* - \mu_m^*$$

- ***Type 3 configuration:***

$$W_a^* = [0, \underline{z}_3^*],\ W_p^* = [\underline{z}_3^*, \bar{z}_3^*],\ I^* = (\bar{z}_3^*, z_{\mathrm{AI}}),\ M_p^* = [z_{\mathrm{AI}}, 1],\ M_a^* = \emptyset,\ \textit{where } 0 < \underline{z}_3^* \le \bar{z}_3^* < z_{\mathrm{AI}}$$

$$\textit{So } \mu_w^* = 0,\ \mu_m^* = \int_0^{\underline{z}_3^*} n(z)^{-1} dG(z),\ \mu_i^* = \mu - \mu_m^*$$

- ***Type 4 configuration:***

$$W_a^* = \emptyset,\ W_p^* = [0, \underline{z}_4^*],\ I^* = (\underline{z}_4^*, \bar{z}_4^*) \ni z_{\text{AI}},\ M_p^* = [\bar{z}_4^*, 1],\ M_a^* = \emptyset,\ \textit{where}\ 0 \leq \underline{z}_4^* < \bar{z}_4^* \leq 1$$

$$\textit{So}\ \mu_w^* = 0,\ \mu_m^* = 0,\ \mu_i^* = \mu$$

Proof. As mentioned above, the proof of this corollary is a direct implication of Lemma B.1. Note that in a Type 2 configuration, I^* can either be $\{z_{\text{AI}}\}$ or $\emptyset$ because the human with knowledge z_{AI} is indifferent between any of the three roles. However, this is irrelevant for all practical purposes because I^* has measure zero. □

Intuitively, in a Type 1 configuration, AI is used as a worker and as an independent producer. In a Type 2 configuration, AI is used in all three possible roles (i.e., as a worker, an independent producer, and a manager). In a Type 3 configuration, AI is used as a manager and as an independent producer, while in a Type 4 configuration, AI is used exclusively as an independent producer.

Now, recall that W and M are the sets of human workers and managers in the pre-AI equilibrium. For $z_{\text{AI}} \in W$, define the function $f_w : [0, z_{\text{AI}}] \to [z_{\text{AI}}, 1]$ by $\int_{z_{\text{AI}}}^{f_w(z;z_{\text{AI}})} dG(u) = \int_0^z h(1-u)dG(u)$ and note that $z_{\text{AI}} \in W$ implies that $f_w(z_{\text{AI}}; z_{\text{AI}}) \leq 1$. Let then $e_w(z; z_{\text{AI}}) \equiv f_w^{-1}(z; z_{\text{AI}})$ and define:

$$\Gamma_w(x) \equiv n(x)(f_w(x;x) - x) - x - \int_x^{f_w(x;x)} n(e_w(u;x))du$$

Similarly, for $z_{\text{AI}} \in M$, define the function $e_m : [z_{\text{AI}}, 1] \to [0, z_{\text{AI}}]$ by $\int_z^1 dG(u) = \int_{e_m(z;z_{\text{AI}})}^{z_{\text{AI}}} h(1-u)dG(u)$, note that $z_{\text{AI}} \in M$ implies that $e_m(z_{\text{AI}}; z_{\text{AI}}) \geq 0$, and define:

$$\Gamma_m(x) \equiv \tfrac{1}{h} - x - \int_x^1 n(e_m(u;x))du$$

Consider the following partition of the knowledge space (note that $W \cup M = [0,1]$ when $h < h_0$): $\mathcal{R}_1 \equiv W \cap \{z \in W : \Gamma_w(z) \leq 0\}$, $\mathcal{R}_2 \equiv \{z \in W : \Gamma_w(z) > 0\} \cup \{z \in M : \Gamma_m(z) > 0\}$, and $\mathcal{R}_3 \equiv M \cap \{z \in M : \Gamma_m(z) \leq 0\}$. The next lemma provides some important properties of this partition that will be useful later on:

Lemma B.2. *$z \in \mathcal{R}_1$ if $z \in [0, \epsilon)$ with $\epsilon \downarrow 0$. Moreover, $\hat{z} \in \mathcal{R}_2$, where $\hat{z}$ is the knowledge of the best pre-AI worker/worst pre-AI manager.*

Proof. Note that $\Gamma_w(0) = 0$ (as $f_w(z;0) = 0$), and that $\Gamma'_w(0) = -1$, as:

$$\Gamma'_w(x) = \tfrac{1}{h} - 1 - \tfrac{1}{h(1-x)} + \tfrac{f_w(x;x)-x}{h(1-x)^2} + h\int_x^{f_w(x;x)} \tfrac{n(e_w(z;x))^3 g(x)}{g(e_w(z;x))} dz$$

Hence, if $z \in [0, \epsilon)$ with $\epsilon \downarrow 0$, then $z \in W$ and $\Gamma_w(z) \leq 0$, so $z \in \mathcal{R}_1$.

We now prove that $\hat{z} \in \mathcal{R}_2$ by showing that $\Gamma_m(\hat{z}) = \Gamma_w(\hat{z}) > 0$. Note that $f_w(z;\hat{z}) = f(z;\hat{z})$, where $f(z;\hat{z})$ is the matching function of the pre-AI equilibrium. This implies that $f_w(\hat{z};\hat{z}) = 1$, so $\Gamma_w(\hat{z}) = (1/h) - \hat{z} - \int_{\hat{z}}^1 n(e(z;\hat{z}))dz > 0$, where the last inequality follows from Lemma A.1. Similarly, note that $e_m(z;\hat{z}) = e(z;\hat{z})$, where $e(z;\hat{z}) = f^{-1}(z;\hat{z})$. Thus, $\Gamma_m(\hat{z}) = (1/h) - \hat{z} - \int_{\hat{z}}^1 n(e(z;\hat{z}))dz$, so $\Gamma_m(\hat{z}) = \Gamma_w(\hat{z}) > 0$. □

We then have the following result:

Lemma B.3. *In the presence of AI, there is a unique competitive equilibrium. It is given as follows:*

- *If $z_{\rm AI} \in \mathcal{R}_1$, then the equilibrium allocation is Type 1. The equilibrium cutoffs $\underline{z}_1^*$ and $\bar{z}_1^*$ satisfy:*

$$\bar{z}_1^* = f_1^*(z_{\rm AI}; \underline{z}_1^*) \quad \textit{and} \quad n(z_{\rm AI})(f_1^*(z_{\rm AI}; \underline{z}_1^*) - z_{\rm AI}) = \underline{z}_1^* + \int_{\underline{z}_1^*}^{f_1^*(z_{\rm AI};\underline{z}_1^*)} n\big(e_1^*(z; \underline{z}_1^*)\big) dz$$

where $f_1^* : [0, z_{\rm AI}] \to [\underline{z}_1^*, \bar{z}_1^*]$ is given by $\int_{\underline{z}_1^*}^{f_1^*(z;\underline{z}_1^*)} dG(u) = \int_0^z h(1-u)dG(u)$ and $e_1^*(z;\cdot) = (f_1^*)^{-1}(z;\cdot)$

- *If $z_{\rm AI} \in \mathcal{R}_2$, then the equilibrium allocation is Type 2. The equilibrium cutoffs $\underline{z}_2^*$ and $\bar{z}_2^*$ satisfy:*

$$\bar{z}_2^* = f_2^*(z_{\rm AI}; \underline{z}_2^*) \quad \textit{and} \quad n(z_{\rm AI})(f_2^*(z_{\rm AI}; \underline{z}_2^*) - z_{\rm AI}) = z_{\rm AI} + \int_{z_{\rm AI}}^{f_2^*(z_{\rm AI};\underline{z}_2^*)} n\big(e_2^*(z; \underline{z}_2^*)\big) dz$$

where $f_2^* : [\underline{z}_2^*, z_{\rm AI}] \to [z_{\rm AI}, \bar{z}_2^*]$ given by $\int_{z_{\rm AI}}^{f_2^*(z;\underline{z}_2^*)} dG(u) = \int_{\underline{z}_2^*}^z h(1-u)dG(u)$ and $e_2^*(z;\cdot) = (f_2^*)^{-1}(z;\cdot)$

- *If $z_{\rm AI} \in \mathcal{R}_3$, then the equilibrium allocation is Type 1. The equilibrium cutoffs $\underline{z}_3^*$ and $\bar{z}_3^*$ satisfy:*

$$\bar{z}_3^* = e_3^*(1; \underline{z}_3^*) \quad \textit{and} \quad \tfrac{1}{h} = z_{\rm AI} + \int_{z_{\rm AI}}^1 n\big(e_3^*(z; \underline{z}_3^*)\big) dz$$

where $f_3^* : [\underline{z}_3^*, \bar{z}_3^*] \to [z_{\rm AI}, 1]$ given by $\int_{z_{\rm AI}}^{f_3^*(z;\underline{z}_3^*)} dG(u) = \int_{\underline{z}_3^*}^z h(1-u)dG(u)$ and $e_3^*(z;\cdot) = (f_3^*)^{-1}(z;\cdot)$

The equilibrium matching function is given by $f^*(z) = f_j^*(z; \underline{z}_j^*)$ if $z_{\rm AI} \in \mathcal{R}_j$, while the equilibrium wage $w^*(z)$ is continuous, strictly increasing, and (weakly) convex, and satisfies:

(i) $w^*(z) = z_{\rm AI}(1 - 1/n(z)) > z$ *for all* $z \in W_a^*$.
(ii) $w^*(z) = f^*(z) - w^*(f^*(z))/n(z) > z$ *for all* $z \in W_p^*$.
(iii) $w^*(z) = z$ *for all* $z \in I^*$.
(iv) $w^*(z) = m_p + \int_{m_p}^z n(e^*(u))du > z$, *for all* $z \in M_p^*$, *where* $m_p \equiv \inf M_p^*$.
(v) $w^*(z) = n(z_{\rm AI})(z - z_{\rm AI}) > z$, *for all* $z \in M_a^*$.

Before formally proving this lemma, we informally derive the equilibrium in one of the regions to provide insight into its construction:

Informal Construction of the Equilibrium.—- Suppose that $z_{\rm AI} \in \mathcal{R}_2$. By Corollary B.1, we know that such equilibrium must lead to the following partition of the human population:

$$W_a^* = [0, \underline{z}_2^*],\ W_p^* = [\underline{z}_2^*, z_{\rm AI}],\ I^* = \emptyset,\ M_p^* = [z_{\rm AI}, \bar{z}_2^*],\ M_a^* = [\bar{z}_2^*, 1],\ \text{where } 0 \le \underline{z}_2^* \le z_{\rm AI} \le \bar{z}_2^* \le 1$$

As mentioned in the main text, given that the equilibrium price of compute is $r^* = z_{\rm AI}$, the zero-profit condition of a tA firm pins down the wage $w^*(z) = z_{\rm AI}(1 - 1/n(z))$ of a human worker with knowledge $z \in W_a^*$. Similarly, the zero profit condition of a bA firm determines the wage $w^*(z) = n(z_{\rm AI})(z - z_{\rm AI})$ of a human manager with knowledge $z \in M_a^*$.

Now consider the firms that do not use AI, i.e., the nA firms. First, let $f_2^*(z)$ be the equilibrium matching function in this case. This function must be strictly increasing and satisfy the following resource constraint $\int_{z_{\rm AI}}^{f_2^*(z)} dG(u) = \int_{\underline{z}_2^*}^z h(1-u)dG(u)$ for all $z \in [\underline{z}_2^*, z_{\rm AI}]$. This constraint states that the total time required to consult on the problems left unsolved by workers in the interval $[\underline{z}_2^*, z]$ must equal the total time available of managers in the interval $[z_{\text{AI}}, f_2^*(z)]$. Moreover, given that $\sup W_p^* = z_{\text{AI}}$ and $\sup M_p^* = \bar{z}_2^*$, it must also be that $f_2^*(z_{\text{AI}}) = \bar{z}_2^*$.

Note then that for any given $z \in W_p$, there exists a unique increasing function $f_2^*(z)$ that satisfies both constraints. It is given by the solution to the differential equation $f_2^{*\prime}(z) = h(1-z)g(z)/g(f_2^*(z))$ with border condition $f_2^*(z_{\text{AI}}) = \bar{z}_2^*$ (which comes from differentiating both sides of the resource constraint). We denote such a unique function by $f_2^*(z; \underline{z}_2^*)$ (as it depends on $\underline{z}_2^*$ through its domain).

With this in mind, consider the problem of a nA firm that recruited $n(z)$ workers of type $z \in W_p^*$ and is deciding which manager $z \in M_p^*$ to hire: $\max_{m \in M_p^*} \Pi_2^{n\text{A}}(z, m) = n(z)[m - w(z)] - w(m)$. As mentioned in the main text, the corresponding first-order condition evaluated at $m = f_2^*(z; \underline{z}_2^*)$ implies that $w^{*\prime}(z) = n(e_2^*(z; \underline{z}_2^*))$ for any $z \in M_p^*$. Thus, $w^*(z) = C + \int_{z_{\text{AI}}}^{z} n(e_2^*(u; \underline{z}_2^*))du$ for any $z \in M_p^*$. The wages of the workers of such firms then come from the zero profit condition of nA firms: $w^*(z) = f_2^*(z; \underline{z}_2^*) - w^*(f_2^*(z; \underline{z}_2^*))/n(z)$ for any $z \in W_p^*$.

The final step is determining the constant C, the cutoff $\underline{z}_2^*$, and arguing that no firms have incentives to deviate. To do this, note that the least knowledgeable human manager has the same knowledge as AI, i.e., $\inf M_p^* = z_{\text{AI}}$. Hence, her wage must be equal to the price of one unit of compute, so $C^* = z_{\text{AI}}$. Moreover, the most knowledgeable manager hired by a nA firm has the same knowledge as the least knowledgeable manager of a bA firm. As a result, they must also receive the same wage, i.e., $\lim_{z \uparrow \bar{z}_2^*} w^*(z) = \lim_{z \downarrow \bar{z}_2^*} w^*(z)$. Since $\bar{z}_2^* = f_2^*(z_{\text{AI}}; \underline{z}_2^*)$, we obtain:

$$n(z_{\text{AI}})(f_2^*(z_{\text{AI}}; \underline{z}_2^*) - z_{\text{AI}}) = z_{\text{AI}} + \int_{z_{\text{AI}}}^{f_2^*(z_{\text{AI}}; \underline{z}_2^*)} n(e_2^*(z; \underline{z}_2^*))dz \tag{6}$$

which is the equilibrium condition in the statement of Lemma B.3. It is then possible to prove that there is a unique cutoff $\underline{z}_2^*$ that satisfies this condition and that such a cutoff is contained in $(0, z_{\text{AI}}]$ if and only if $z_{\text{AI}} \in \mathcal{R}_2$. This explains why this equilibrium can only arise in such a region of the parameter space. The fact that $C^* = z_{\text{AI}}$ and that $\underline{z}_2^*$ satisfies (6) then implies that $w^*(z)$ is also continuous at the juncture between W_p^* and M_p^*:

$$\lim_{z \downarrow \underline{z}_2^*} w^*(z) = f_2^*(\underline{z}_2^*; \underline{z}_2^*) - \frac{w^*(f_2^*(\underline{z}_2^*; \underline{z}_2^*))}{n(\underline{z}_2^*)} = z_{\text{AI}} \left(1 - \frac{1}{n(\underline{z}_2^*)}\right) = \lim_{z \uparrow \underline{z}_2^*} w^*(z)$$
$$\lim_{z \uparrow z_{\text{AI}}} w^*(z) = f_2^*(z_{\text{AI}}; \underline{z}_2^*) - \frac{w^*(f_2^*(z_{\text{AI}}; \underline{z}_2^*))}{n(z_{\text{AI}})} = z_{\text{AI}} = \lim_{z \downarrow z_{\text{AI}}} w^*(z)$$

This is sufficient sufficient to guarantee that $w^*(z)$ is continuous in all its domain, i.e., for all $z \in [0, 1]$.

From here, arguing that no firm has incentives to deviate is straightforward. Indeed, note that the wage function is continuous, strictly increasing, and weakly convex. This implies that if a firm does not have incentives to deviate "locally," then it does not have incentives to deviate globally either. That bA firms do not want to deviate locally, i.e., hire a different human manager in M_p^*, follows because such deviation also leads to no profits. Similar reasoning also explains why tA and nA firms do not have incentives to deviate locally either. □

We now formally prove the lemma. This is done in two steps. First, we show that the outcomes described in the lemma are indeed an equilibrium by verifying that (i) markets clear, and (ii) firms are maximizing their profits while obtaining zero profits. Then, we prove that the equilibrium is unique.

Proof of Step 1. We will only show this part for $z_{\rm AI} \in \mathcal{R}_1$, as the other two cases are analogous. We begin by verifying market clearing in the market for compute. By Corollary B.1, it is immediate that $\mu_i^* + \mu_w^* + \mu_m^* = \mu$. Moreover, the total time required to consult on the problems left unsolved by AI is equal to the total time available of managers in M_a^*, i.e., $h(1-z_{\rm AI})\mu_w^* = \int_{\bar{z}_1^*}^1 dG(z)$.

We now move to market clearing of the labor market. First, it must be that the total time required to consult on the problems left unsolved by the human workers in the interval $[0, z] \subseteq W_p^*$ is equal to the total time available of human managers in the interval $[\underline{z}_1^*, f_1^*(z; \underline{z}_1^*)] \subseteq M_p^*$. This resource constraint is satisfied as $f_1^*(z; \underline{z}_1^*)$ is given by $\int_{\underline{z}_1^*}^{f_1^*(z;\underline{z}_1^*)} dG(u) = \int_0^z h(1-u)dG(u)$ and $\bar{z}_1^* = f_1^*(z_{\rm AI}; \underline{z}_1^*)$.

Second, it must be that the union of the sets $(W_p^*, W_a^*, I^*, M_p^*, M_a^*)$ is $[0, 1]$, and the intersection of any two of these sets has measure zero. By Corollary B.1, this occurs if and only if $z_{\rm AI} < \underline{z}_1^* \le \bar{z}_1^* \le 1$. Verifying that that $\underline{z}_1^* \le \bar{z}_1^*$ is straightforward: It follows because $f_1^*(z; \underline{z}_1^*)$ is strictly increasing in z plus the fact that $\underline{z}_1^* = f_1^*(0; \underline{z}_1^*)$ and $\bar{z}_1^* = f_1^*(z_{\rm AI}; \underline{z}_1^*)$.

Showing that $z_{\rm AI} < \underline{z}_1^*$ requires more work. Note that $\underline{z}_1^*$ is given by the solution $\Gamma_1(\underline{z}_1^*; z_{\rm AI}) = 0$, where $\Gamma_1(x; z_{\rm AI}) \equiv n(z_{\rm AI})(f_1^*(z_{\rm AI}; x, z_{\rm AI}) - z_{\rm AI}) - x - \int_x^{f_1^*(z_{\rm AI};x,z_{\rm AI})} n\big(e_1^*(z; x, z_{\rm AI})\big)dz$ (here we are making explicit that $f_1^*(\cdot)$ and $e_1^*(\cdot)$ also depend indirectly on $z_{\rm AI}$ through the boundary of the set $W_p^* = [0, z_{\rm AI}]$ to avoid any type of confusion[29]). It is then not difficult to prove that $\Gamma_1(x; z_{\rm AI})$ is strictly increasing in x, so $z_{\rm AI} < \underline{z}_1^*$ if and only if $\Gamma_1(0; z_{\rm AI}) < 0$. Furthermore, note that $f_1^*(z; 0, z_{\rm AI})$ satisfies $\int_0^{f_1^*(z;0,z_{\rm AI})} dG(u) = \int_0^z h(1-u)dG(u)$ for $z \in [0, z_{\rm AI}]$, which implies that $f_1^*(z; 0, z_{\rm AI}) = f_w(z; z_{\rm AI})$ (and, therefore, $e_1^*(z; 0, z_{\rm AI}) = e_w(z; z_{\rm AI})$). Hence, $\Gamma_1(0; z_{\rm AI}) = \Gamma_w(z_{\rm AI}) < 0$ as $z_{\rm AI} \in \mathcal{R}_1$.

Finally, we show that $\bar{z}_1^* \le 1$. This is more involved. To prove it, we show that $\underline{z}_1^* < \hat{z}$ and then use this result to conclude that $\bar{z}_1^* \le 1$. As a first step, note that $f_1^*(z; \hat{z}, \hat{z})$ satisfies $\int_{\hat{z}}^{f_1^*(z;\hat{z},\hat{z})} dG(u) = \int_0^z h(1-u)dG(u)$ for $z \in [0, \hat{z}]$, so $f_1^*(z; \hat{z}, \hat{z}) = f(z; \hat{z})$ where $f(z; \hat{z})$ is the pre-AI matching function. Recall then that $\underline{z}_1^*$ is the unique solution $\Gamma_1(\underline{z}_1^*; z_{\rm AI}) = 0$, where $\Gamma_1(x; z_{\rm AI})$ is strictly increasing in x. It is not difficult to prove that $\Gamma_1(x; z_{\rm AI})$ is also strictly decreasing in $z_{\rm AI}$ for any given x. We then claim that $\Gamma_1(\hat{z}; z_{\rm AI}) > 0$, which immediately implies that $\underline{z}_1^* < \hat{z}$. Indeed, note that $\Gamma_1(\hat{z}; z_{\rm AI}) \ge \Gamma_1(\hat{z}; \hat{z}) = \frac{1}{h} - \hat{z} - \int_{\hat{z}}^1 n(e(z; \hat{z}))dz$, where the first inequality follows because $\Gamma_1(x; z_{\rm AI})$ is strictly decreasing in $z_{\rm AI}$ and $z_{\rm AI} \le \hat{z}$ (as $z_{\rm AI} \in W$) and the last equality because $e_1^*(z; \hat{z}, \hat{z}) = e(z; \hat{z})$ for all $z \in [\hat{z}, 1]$. However, by Lemma A.1, we know that $\frac{1}{h} - \hat{z} - \int_{\hat{z}}^1 n(e(z; \hat{z}))dz > 0$ when $h < h_0$, so $\Gamma_1(\hat{z}; z_{\rm AI}) > 0$.

Having proved that $\underline{z}_1^* < \hat{z}$, we now show that $\bar{z}_1^* \le 1$. By construction $\bar{z}_1^*$ satisfies $\int_{\underline{z}_1^*}^{\bar{z}_1^*} dG(u) = \int_0^{z_{\rm AI}} h(1-u)dG(u)$. Note, moreover, that $z \in W$ implies that $\int_0^{z_{\rm AI}} h(1-u)dG(u) \le \int_0^{\hat{z}} h(1-u)dG(u) = \int_{\hat{z}}^1 dG(u)$ (as $z_{\rm AI} \le \hat{z}$ and $\int_0^{\hat{z}} h(1-u)dG(u) = \int_{\hat{z}}^1 dG(u)$). Hence, $\int_{\underline{z}_1^*}^{\bar{z}_1^*} dG(u) \le \int_{\hat{z}}^1 dG(u)$. Given that $\underline{z}_1^* < \hat{z}$, it must be that $\bar{z}_1^* < 1$.

[29] In the statement of Lemma B.3, we simply wrote $f_1^*(z; \underline{z}_1^*)$ instead of $f_1^*(z; \underline{z}_1^*, z_{\rm AI})$ to avoid cluttering notation.

Having verified market clearing, we now show that in the candidate equilibrium firms maximize their profits while obtaining zero profits. Given how the wages are constructed (see Section 3.3 of the main text), it is clear that firms are optimizing their profits "locally" while obtaining zero profits. Thus, we only need to consider "global deviations." As discussed above, to discard such global deviations, it is sufficient to show that $w^*(z)$ is continuous, strictly increasing, and weakly convex. This is what we prove next.

To show continuity, it suffices to verify that $w^*(z)$ is continuous at the junctures of (i) M_p^* and M_a^*, (ii) I^* and M_p^*, and (iii) W_p^* and I^*. For (i), note that $\lim_{z\uparrow\bar{z}_1^*} w^*(z) = \bar{z}_1^* + \int_{\underline{z}_1^*}^{\bar{z}_1^*} n(e_1^*(z;\underline{z}_1^*))dz$ and $\lim_{z\downarrow\bar{z}_1^*} w^*(z) = n(z_{\text{AI}})(\bar{z}_1^* - z_{\text{AI}})$. Hence, from the conditions determining $\underline{z}_1^*$ and $\bar{z}_1^*$, we obtain that $\lim_{z\uparrow\bar{z}_1^*} w^*(z) = \lim_{z\downarrow\bar{z}_1^*} w^*(z)$. For (ii) note that by construction, $\lim_{z\uparrow\underline{z}_1^*} w^*(z) = \underline{z}_1^* = \lim_{z\downarrow\underline{z}_1^*} w^*(z)$. Finally, for (iii) note that $\lim_{z\uparrow z_{\text{AI}}} w^*(z) = \bar{z}_1^* - w^*(\bar{z}_1^*)/n(z_{\text{AI}}) = z_{\text{AI}} = \lim_{z\downarrow z_{\text{AI}}} w^*(z)$ given that $w^*(\bar{z}_1^*) = n(z_{\text{AI}})(\bar{z}_1^* - z_{\text{AI}})$.

With continuity at hand, proving that $w^*(z)$ is strictly increasing and weakly convex is straightforward: The logic is analogous to the proof of Corollary A.1 (which shows that the pre-AI wage function satisfies these two properties). Consequently, in the case $z_{\text{AI}} \in \mathcal{R}_1$, the outcome described in the statement is indeed a competitive equilibrium. □

Proof of Step 2. We will only show this part for $z_{\text{AI}} \in \mathcal{R}_1$, as the other two cases follow the same logic. In particular, we show that if $z_{\text{AI}} \in \mathcal{R}_1$, then there cannot be any other type of equilibrium.

Suppose first by contradiction that there is a Type 3 equilibrium. By Corollary B.1, we know that such equilibrium must lead to the following partition of the human population:

$$W_a^* = [0, \underline{z}_3^*],\ W_p^* = [\underline{z}_3^*, \bar{z}_3^*],\ I^* = (\bar{z}_3^*, z_{\text{AI}}),\ M_p^* = [z_{\text{AI}}, 1],\ M_a^* = \emptyset,\ \text{where } 0 < \underline{z}_3^* \leq \bar{z}_3^* < z_{\text{AI}}$$

By Lemma B.2, we have that if $z_{\text{AI}} \in \mathcal{R}_1$, then $z_{\text{AI}} < \hat{z}$. However, if so, then:

$$\int_{\underline{z}_3^*}^{\bar{z}_3^*} h(1-u)dG(u) < \int_0^{z_{\text{AI}}} h(1-u)dG(u) < \int_0^{\hat{z}} h(1-u)dG(u) = \int_{\hat{z}}^1 dG(u) < \int_{z_{\text{AI}}}^1 dG(u)$$

which violates the resource constraint that the total time required to consult on the problems left unsolved by the human workers in the interval W_p^* must equal to the total time available of human managers in the interval M_p^*, i.e., $\int_{\underline{z}_3^*}^{\bar{z}_3^*} h(1-u)dG(u) = \int_{z_{\text{AI}}}^1 dG(u)$. Hence, a Type 3 configuration cannot arise when $z_{\text{AI}} \in \mathcal{R}_1 \subset W$.

Now suppose for contradiction that there is a Type 2 equilibrium. As explained above (see "Informal Construction of the Equilibrium"), for this to be an equilibrium, there must exist a cutoff $\underline{z}_2^* \in [0, z_{\text{AI}}]$ such that $\Gamma_2(\underline{z}_2^*; z_{\text{AI}}) = 0$, where $\Gamma_2(x; z_{\text{AI}}) \equiv n(z_{\text{AI}})(f_2^*(z_{\text{AI}}; x) - z_{\text{AI}}) - z_{\text{AI}} - \int_{z_{\text{AI}}}^{f_2^*(z_{\text{AI}}; x)} n(e_2^*(z; x))dz$. It is then not difficult to prove that $\Gamma_2(x; z_{\text{AI}})$ is strictly decreasing in x. Hence, there exists at most one $\underline{z}_2^*$ that satisfies $\Gamma_2(\underline{z}_2^*; z_{\text{AI}}) = 0$, and for $\underline{z}_2^* \geq 0$, it must be that $\Gamma_2(0; z_{\text{AI}}) \geq 0$, where $f_2^*(z; 0)$ is given by $\int_{z_{\text{AI}}}^{f_2^*(z;0)} dG(u) = \int_0^z h(1-u)dG(u)$ for $z \in [0, z_{\text{AI}}]$. However, from this last condition we have that $f_2^*(z; 0) = f_w(z; z_{\text{AI}})$ (as $f_2^*(z; 0)$ and $f_w(z; z_{\text{AI}})$ satisfy the same condition), so $\Gamma_2(0; z_{\text{AI}}) = \Gamma_w(z_{\text{AI}})$. Consequently, for $\underline{z}_2^* \geq 0$, we need that $\Gamma_w(z_{\text{AI}}) \geq 0$, which contradicts the fact that $z_{\text{AI}} \in \mathcal{R}_1$.

Finally, suppose by contradiction that there is a Type 4 equilibrium. By Corollary B.1, we know that such equilibrium must lead to the following partition of the human population:

$$W_p^* = \emptyset,\ W_p^* = [0, \underline{z}_4^*],\ I^* = (\underline{z}_4^*, \bar{z}_4^*) \ni z_{\text{AI}},\ M_p^* = [\bar{z}_4^*, 1],\ M_a^* = \emptyset$$

Moreover, following similar reasoning as the one developed above for a Type 2 equilibrium, for this to be an equilibrium, (i) it must be that $\bar{z}_4^* = f_4^*(0; \underline{z}_4^*)$, where $f_4^*(0; \underline{z}_4^*)$ satisfies $f_4^*(\underline{z}_4^*; \underline{z}_4^*) = 1$ and $\int_{f_4^*(0;\underline{z}_4^*)}^{f_4^*(z;\underline{z}_4^*)} dG(u) = \int_{\underline{z}_4^*}^{z} h(1-u)dG(u)$ for $z \in [0, \underline{z}_4^*]$, and (ii) there must exist a cutoff $\underline{z}_4^* < f_4^*(0; \underline{z}_4^*)$ such that $1/h - f_4^*(0; \underline{z}_4^*) - \int_{f_4^*(0;\underline{z}_4^*)}^{1} n(e_4^*(z; \underline{z}_4^*))dz = 0$. But by Lemma A.1, we know that the unique solution to these equilibrium conditions is $\underline{z}_4^* = \underline{z}$ and $\bar{z}_4^* = f_4^*(0; \underline{z}_4^*) = \bar{z}$, where $\underline{z}$ and $\bar{z}$ are the equilibrium cutoffs of the pre-AI equilibrium when $h > h_0$. This, however, implies that this configuration is an equilibrium only if $z_{\text{AI}} \in I^* = (\underline{z}, \bar{z})$, which contradicts the assumption that $z_{\text{AI}} \in \mathcal{R}_1$. □

C Proofs Omitted from Section 4

C.1 Proof of Proposition 3

• $z_{\text{AI}} \in \text{int}W$.— By Lemma B.3 the AI equilibrium is either Type 1 or Type 2. If it is Type 2, then $\sup W^* = \inf M^* = z_{\text{AI}}$, so $W^* \subset W$ and $M^* \supset M$, since $z_{\text{AI}} < \hat{z}$ when $z_{\text{AI}} \in \text{int}W$. If it is Type 1, then $\sup W^* = z_{\text{AI}}$ and $\inf M^* = \underline{z}_1^*$. That $\sup W^* = z_{\text{AI}}$ immediately implies that $W^* \subset W$ as $z_{\text{AI}} < \hat{z}$. That $\inf M^* = \underline{z}_1^*$ immediately implies that $M^* \supset M$ given that $\underline{z}_1^* < \hat{z}$ (as shown in the proof of Lemma B.3). □

• $z_{\text{AI}} \in \text{int}M$.— By Lemma B.3, the AI equilibrium is either Type 2 or Type 3. Suppose first that it is Type 2. Then $\sup W^* = \inf M^* = z_{\text{AI}}$, so $W^* \supset W$ and $M^* \subset M$ given that $z_{\text{AI}} > \hat{z}$ when $z_{\text{AI}} \in \text{int}M$.

Now suppose the equilibrium is Type 3. Then, $\inf M^* = z_{\text{AI}} > \hat{z}$, immediately implying that $M^* \subset M$. To prove that $W^* \supset W$, we show that $\sup W^* = \bar{z}_3^* > \hat{z}$. Indeed, recall that $\bar{z}_3^*$ is given by:[30]

$$\bar{z}_3^* = e_3^*(1; \underline{z}_3^*, z_{\text{AI}}) \quad \text{and} \quad \tfrac{1}{h} = z_{\text{AI}} + \int_{z_{\text{AI}}}^{1} n\big(e_3^*(z; \underline{z}_3^*, z_{\text{AI}})\big)dz$$
$$\text{where } \int_{z_{\text{AI}}}^{f_3^*(z;\underline{z}_3^*,z_{\text{AI}})} dG(u) = \int_{\underline{z}_3^*}^{z} h(1-u)dG(u) \text{ for } z \in [\underline{z}_3^*, \bar{z}_3^*]$$

Using the fact that $f_3^*(\bar{z}_3^*; \underline{z}_3^*, z_{\text{AI}}) = 1$, the equilibrium conditions that determine $\underline{z}_3^*$ and $\bar{z}_3^*$ can be written as follows:[31]

$$\underline{z}_3^* = \tilde{e}_3^*(z_{\text{AI}}; \bar{z}_3^*, z_{\text{AI}}) \quad \text{and} \quad \tfrac{1}{h} = z_{\text{AI}} + \int_{z_{\text{AI}}}^{1} n\big(\tilde{e}_3^*(z; \bar{z}_3^*, z_{\text{AI}})\big)dz$$
$$\text{where } G(z) = 1 - \int_{\tilde{e}_3^*(z;\bar{z}_3^*,z_{\text{AI}})}^{\bar{z}_3^*} h(1-u)dG(u) \text{ for } z \in [z_{\text{AI}}, 1]$$

Define $\Gamma_3(x; z_{\text{AI}}) \equiv \frac{1}{h} - z_{\text{AI}} - \int_{z_{\text{AI}}}^{1} n\big(\tilde{e}_3^*(z; x, z_{\text{AI}})\big)dz$. It is not difficult to prove that $\Gamma_3(x; z_{\text{AI}})$ is strictly decreasing in x and strictly increasing in z_{AI}. Moreover, $\bar{z}_3^*$ is given by the unique solution

[30]To avoid any type of confusion, we are making explicit that $f_3^*(z; \underline{z}_3^*, z_{\text{AI}})$ and $e_3^*(1; \underline{z}_3^*, z_{\text{AI}})$ depend on both $\underline{z}_3^*$ and z_{AI} (in the statement of Lemma B.3, we simply wrote $f_3^*(z; \underline{z}_3^*)$ instead of $f_3^*(z; \underline{z}_3^*, z_{\text{AI}})$ to avoid cluttering notation).

[31]Note that $\tilde{e}_3^*(z; \bar{z}_3^*, z_{\text{AI}})$ is the equilibrium employee function indexed by $\bar{z}_3^*$ instead of $\underline{z}_3^*$.

to $\Gamma_3(\bar{z}_3^*; z_{\text{AI}}) = 0$. To prove that $\bar{z}_3^* > \hat{z}$, it suffices to show that $\Gamma_3(\hat{z}; z_{\text{AI}}) > 0$. Since $z_{\text{AI}} > \hat{z}$, we have that $\Gamma_3(\hat{z}; z_{\text{AI}}) > \Gamma_3(\hat{z}; \hat{z}) = \frac{1}{h} - \hat{z} - \int_{z_{\text{AI}}}^{1} n\big(e(z;\hat{z})\big)dz > 0$, where the second-to-last inequality follows because $\tilde{e}_3^*(z; \hat{z}, \hat{z}) = e(z; \hat{z})$ for all $z \in [\hat{z}, 1]$, and the last inequality comes from the pre-AI equilibrium characterized in Lemma A.1. □

• $z_{\text{AI}} = \hat{z}$.— By Lemma B.2, we know that $z_{\text{AI}} = \hat{z} \in \mathcal{R}_2$, so Lemma B.3 implies that the equilibrium is necessarily Type 2. This implies that $\sup W^* = \inf M^* = z_{\text{AI}} = \hat{z}$, so there is no occupational displacement.

Showing that $W_a^* \neq \emptyset$ and $M_a^* \neq \emptyset$ requires more work. First, it is not difficult to prove that $\underline{z}_2^* > 0$ if and only if $\bar{z}_2^* < 1$. Hence, to prove that $W_a^* \neq \emptyset$ and $M_a^* \neq \emptyset$ it suffices to show that $\underline{z}_2^* > 0$. To do the latter, suppose for contradiction that $\underline{z}_2^* = 0$ ($\underline{z}_2^* < 0$ immediately contradicts that we are in a Type 2 equilibrium). Then the equilibrium matching function is given by $\int_{\hat{z}}^{f_2^*(z;0)} dG(u) = \int_0^z h(1-u)dG(u)$ for $z \in [\hat{z}, 1]$, implying that $f_2^*(z; 0) = f(z; \hat{z})$ for all $z \in [\hat{z}, 1]$. However, if so, then:

$$n(z_{\text{AI}})(f_2^*(z_{\text{AI}}; 0) - z_{\text{AI}}) = \tfrac{1}{h} \neq \hat{z} + \int_{\hat{z}}^{1} n\big(e(z;\hat{z})\big)dz = z_{\text{AI}} + \int_{z_{\text{AI}}}^{f_2^*(z_{\text{AI}};0)} n\big(e_2^*(z;0)\big)dz$$

where the inequation follows because $1/h - \hat{z} > \int_{\hat{z}}^{1} n(e(z;\hat{z}))dz$ (by Lemma A.1). Thus, the equilibrium condition for $\underline{z}_2^* = 0$ is not satisfied (see the statement of Lemma B.3). Contradiction. □

C.2 Proof of Proposition 4

As noted in the main text, a worker's productivity increases if and only if her managerial match improves. Similarly, the span of control of a given manager increases if and only the knowledge of her workers increases.

• $z_{\text{AI}} \in \text{int}W$.— First, we show that every $z \in W^*$ has a worse manager post-AI. Note that if $z \in W_a^*$, then such a worker is matched with AI in the AI equilibrium. However, if so, then $f(z; \hat{z}) \geq \hat{z} > z_{\text{AI}}$, as $z_{\text{AI}} \in \text{int}W$.

Proving that every $z \in W_p^*$ also has a worse manager is more involved. Since $z_{\text{AI}} \in \text{int}W$, then the AI equilibrium is either Type 1 or Type 2. Suppose first it is Type 1. Then, the matching functions pre- and post-AI are given by:

$$\int_{\underline{z}_1^*}^{f_1^*(z;\underline{z}_1^*)} dG(u) = \int_0^z h(1-u)dG(u) \text{ for } z \in W_p^* = [0, z_{\text{AI}}]$$
$$\int_{\hat{z}}^{f(z;\hat{z})} dG(u) = \int_0^z h(1-u)dG(u) \text{ for } z \in W = [0, \hat{z}]$$

Thus, if $z \in W_p^* \cap W = W_p^*$, then $\int_{\underline{z}_1^*}^{f_1^*(z;\underline{z}_1^*)} dG(u) = \int_{\hat{z}}^{f(z;\hat{z})} dG(u)$, so $f_1^*(z; \underline{z}_1^*) < f(z; \hat{z})$ as $\underline{z}_1^* < \hat{z}$.

Suppose instead that the AI equilibrium is Type 2. Then, the matching functions pre- and post-AI are given by:

$$\int_{z_{\text{AI}}}^{f_2^*(z;\underline{z}_2^*)} dG(u) = \int_{\underline{z}_2^*}^{z} h(1-u)dG(u) \text{ for } z \in W_p^* = [\underline{z}_2^*, z_{\text{AI}}]$$
$$\int_{\hat{z}}^{f(z;\hat{z})} dG(u) = \int_0^z h(1-u)dG(u) \text{ for } z \in W = [0, \hat{z}] \tag{7}$$

Consequently, if $z \in W_p^* \cap W = W_p^*$, then $\int_{\hat{z}}^{f(z;\hat{z})} dG(u) - \int_{z_{\text{AI}}}^{f_2^*(z;\underline{z}_2^*)} dG(u) = \int_0^{\underline{z}_2^*} h(1-u)dG(u) > 0$, which implies that $f_2^*(z;\underline{z}_2^*) < f(z;\hat{z})$ as $z_{\text{AI}} < \hat{z}$.

We now turn to managers, i.e., those $z \in M \subset M^*$. We first claim that if $e(z;\hat{z}) = z_{\text{AI}}$, then $z \in M_a^* \cap M$. This immediately implies that if $e(z;\hat{z}) = z_{\text{AI}}$, then z manages a firm of equal size pre- and post-AI. The proof is via the contrapositive. Suppose that $z \notin M_a^* \cap M$ (but that z is a manager). Then $z \in M_p^* \cap M$. However, if so, then $z_{\text{AI}} \geq e_j^*(z;\underline{z}_j^*) > e(z;\hat{z})$, where the first inequality is because AI is the best worker, and the second inequality follows because $e_j^*(z';\underline{z}_j^*) > e(z';\hat{z})$ for all $z' \in M_p^* \cap M$ if $f_j^*(z'';\underline{z}_j^*) < f(z'';\hat{z})$ for all $z'' \in W_p^* \cap W = W_p^*$ (which we already showed is true for $j = 1, 2$). Hence, $e(z;\hat{z}) \neq z_{\text{AI}}$.

The previous claim then implies that if $e(z;\hat{z}) < z_{\text{AI}}$, then z manages a strictly larger firm post-AI, while if $e(z;\hat{z}) > z_{\text{AI}}$, then z manages a strictly smaller firm post-AI. Indeed, if $e(z;\hat{z}) < z_{\text{AI}}$, then either $z \in M_p^* \cap M$ or $z \in M_a^* \cap M$. If $z \in M_p^*$, we already know that $e_j^*(z;\underline{z}_j^*) > e(z;\hat{z})$, so z manages a strictly larger firm post-AI, while if $z \in M_a^*$, then the knowledge of z's workers also increases since she is now managing AI, while before, she was managing humans with knowledge $e(z;\hat{z})$. Similarly, if $e(z;\hat{z}) > z_{\text{AI}}$, then $z \in M_a^* \cap M$. Hence, the knowledge of z's workers decreases since she is now supervising the work of AI, while before, she was managing humans with knowledge $e(z;\hat{z})$. □

• $z_{\text{AI}} \in \text{int}M$.— First, we will show that every $z \in M^* \subset M$ manages a strictly larger firm post-AI than pre-AI. Note that if $z \in M_a^*$, then such a manager is matched with AI in the post-AI equilibrium. However, if so, then $e(z;\hat{z}) \leq \hat{z} < z_{\text{AI}}$, as $z_{\text{AI}} \in \text{int}M$.

We now show that the same holds for every $z \in M_p^*$. Since $z_{\text{AI}} \in \text{int}M$, then the AI equilibrium is either Type 2 or Type 3. Suppose first it is Type 2. Then, the employee functions pre- and post-AI are given by:

$$\begin{aligned} &\int_z^{\bar{z}_2^*} dG(u) = \int_{e_2^*(z;\underline{z}_2^*)}^{z_{\text{AI}}} h(1-u)dG(u) \text{ for } z \in M_p^* = [z_{\text{AI}}, \bar{z}_2^*] \\ &\int_z^1 dG(u) = \int_{e(z;\hat{z})}^{\hat{z}} h(1-u)dG(u) \text{ for } z \in M = [\hat{z}, 1] \end{aligned} \tag{8}$$

Consequently, if $z \in M_p^* \cap M = M_p^*$, then $0 < \int_{\bar{z}_2^*}^1 dG(u) = \int_{e(z;\hat{z})}^{\hat{z}} h(1-u)dG(u) - \int_{e_2^*(z;\underline{z}_2^*)}^{z_{\text{AI}}} h(1-u)dG(u)$, which implies that $e_2^*(z;\underline{z}_2^*) > e(z;\hat{z})$ (since $z_{\text{AI}} > \hat{z}$).

Suppose instead that the AI equilibrium is Type 3. Then, the employee functions pre- and post-AI are given by:

$$\begin{aligned} &\int_z^1 dG(u) = \int_{e_3^*(z;\underline{z}_3^*)}^{\bar{z}_3^*} h(1-u)dG(u) \text{ for } z \in M_p^* = [z_{\text{AI}}, 1] \\ &\int_z^1 dG(u) = \int_{e(z;\hat{z})}^{\hat{z}} h(1-u)dG(u) \text{ for } z \in M = [\hat{z}, 1] \end{aligned}$$

Consequently, for $z \in M_p^* \cap M = M_p^*$, then $\int_{e_3^*(z;\underline{z}_3^*)}^{\bar{z}_3^*} h(1-u)dG(u) = \int_{e(z;\hat{z})}^{\hat{z}} h(1-u)dG(u)$. However, if so, then $e(z;\hat{z}) < e_3^*(z;\underline{z}_3^*)$ since $\hat{z} < \bar{z}_3^*$.

We now turn to workers, i.e., those with $z \in W \subset W^*$. We first claim that if $z = e(z_{\text{AI}};\hat{z})$, then $z \in W_a^* \cap W$. This immediately implies that if $z = e(z_{\text{AI}};\hat{z})$, then z is equally productive pre- and post-AI. The proof is via the contrapositive. Suppose that $z \notin W_a^* \cap W$ (but that z is a worker). Then $z \in W_p^* \cap W$. However, if so, then $z_{\text{AI}} \leq f_j^*(z; \underline{z}_j^*) < f(z; \hat{z})$, where the first inequality is because AI is the worst manager, and the second inequality follows because $f_j^*(z'; \underline{z}_j^*) < f(z'; \hat{z})$ for all $z' \in W_p^* \cap W$ if $e_j^*(z''; \underline{z}_j^*) > e(z''; \hat{z})$ for all $z'' \in M_p^* \cap M = M_p^*$ (which we already showed is true for $j = 2, 3$). Hence, $z \neq e(z_{\text{AI}}; \hat{z})$.

The previous claim then implies that if $z < e(z_{\text{AI}}; \hat{z})$, then z is strictly more productive post-AI than pre-AI, while if $z > e(z_{\text{AI}}; \hat{z})$, then z is strictly less productive post-AI than pre-AI. Indeed, if $z < e(z_{\text{AI}}; \hat{z})$, then $z \in W_a^* \cap W$. Hence, the knowledge of z's manager increases since she is now being managed by AI, while before, she was being managed by a human with knowledge $f(z; \hat{z}) < z_{\text{AI}}$. Similarly, if $z > e(z_{\text{AI}}; \hat{z})$, then $z \in W_a^* \cap W$ or $z \in W_p^* \cap W$. If $z \in W_a^* \cap W$, the knowledge of z's manager decreases since she is now being managed by AI (while before, she was being managed by a human with knowledge $f(z; \hat{z}) > z_{\text{AI}}$), while if $z \in W_p^* \cap W$, the knowledge of her manager again decreases since we already established that $f_j^*(z; \underline{z}_j^*) < f(z; \hat{z})$ for $j = 2, 3$. □

• $z_{\text{AI}} = \hat{z}$.— We first show that each $z \in W^* = W$ is managed by a worse manager post-AI than pre-AI (strictly so for all $z \neq 0$). By Lemma B.2, we know that $z_{\text{AI}} = \hat{z} \in \mathcal{R}_2$, so Lemma B.3 implies that the equilibrium is necessarily Type 2. Moreover, as shown in the proof of Proposition 3, in this case we have that $\underline{z}_2^* > 0$ and $\bar{z}_2^* < 1$. Consequently, if $z \in W_a^*$, then $f(z; \hat{z}) \geq \hat{z} = z_{\text{AI}}$, where the first inequality is strict when $z > 0$. If $z \in W_p^*$ instead, then the matching functions pre- and post-AI are given by (7) evaluated at $z_{\text{AI}} = \hat{z}$. Consequently, for $z \in W^* = W$, $\int_{f_2^*(z; \underline{z}_2^*)}^{f(z;\hat{z})} dG(u) = \int_0^{\underline{z}_2^*} h(1-u)dG(u) > 0$, which implies that $f_2^*(z; \underline{z}_2^*) < f(z; \hat{z})$ given that $\underline{z}_2^* > 0$.

We now show that each $z \in M^* = M$ improves her match post-AI compared to pre-AI (strictly so for all $z \neq 1$). Indeed, if $z \in M_a^*$, then $e(z; \hat{z}) \leq \hat{z} = z_{\text{AI}}$, where the first inequality is strict inequality when $z < 1$. If $z \in M_p^*$ instead, then the employee functions pre- and post-AI are given by (8) evaluated at $z_{\text{AI}} = \hat{z}$. Consequently, for $z \in M^* = M$, $\int_{\bar{z}_2^*}^1 dG(u) = \int_{e(z;\hat{z})}^{e_2^*(z;\underline{z}_2^*)} h(1-u)dG(u) > 0$, which implies that $e_2^*(z; \underline{z}_2^*) > e(z; \hat{z})$ since $\bar{z}_2^* < 1$. □

C.3 Proof of Lemma 2

For ease of exposition, we have divided the proof of the lemma into three smaller claims:

Claim C.1. *(i)* $\Delta(z_{\text{AI}}) < 0$, *(ii)* $\Delta(1) > 0$, *and (iii)* $\Delta(0) > 0$ *if* $z_{\text{AI}} \in M$.

Proof. That $\Delta(z_{\text{AI}}) = z_{\text{AI}} - w(z_{\text{AI}}) < 0$ follows directly from the fact that $w(z) > z$ for all $z \in [0, 1]$ when $h < h_0$ (see Lemma A.1 and Corollary A.1). Consider next $\Delta(1) = w^*(1) - w(1)$. By Lemma B.3, we have that $w^*(1) = 1/h$, while by Lemma A.1 and Corollary A.1, we have that $w(1) < 1/h$. Hence, $\Delta(1) > 0$. Finally, consider $\Delta(0) = w^*(0) - w(0)$ and suppose that $z_{\text{AI}} \in M$. From Lemma B.3, we have that $w^*(0) = z_{\text{AI}}(1 - h)$ as the human with zero knowledge is managed by AI irrespective of whether $z_{\text{AI}} \in \mathcal{R}_2$ or $z_{\text{AI}} \in \mathcal{R}_3$. Moreover, from Lemma A.1 and Corollary A.1 we have that

$w(0) = \hat{z} - hw(\hat{z})$. Thus, $\Delta(0) = z_{\mathrm{AI}}(1-h) - (\hat{z} - hw(\hat{z})) > (1-h)(z_{\mathrm{AI}} - \hat{z}) \geq 0$, where the first inequality follows because $w(\hat{z}) > \hat{z}$, and the second inequality follows because $z_{\mathrm{AI}} \geq \hat{z}$. □

Claim C.2. *If $\Delta(z) > 0$ for some $z \in [z_{\mathrm{AI}}, 1)$, then $\Delta(z') > 0$ for all $z' \in [z, 1)$.*

Proof. Given that $\Delta(z_{\mathrm{AI}}) < 0$ (by Claim C.1), it suffices to show that if $\Delta(z)$ crosses zero at some $z > z_{\mathrm{AI}}$, then it always crosses zero from below. We first consider $z_{\mathrm{AI}} \in \mathrm{int}W$ and then $z_{\mathrm{AI}} \in M$.

• $z_{\mathrm{AI}} \in \mathrm{int}W$.— Lemma B.3 implies that $\sup W^* = z_{\mathrm{AI}}$, while Proposition 3 that $W^* \subset W$ and $M^* \supset M$. Hence, if $z > z_{\mathrm{AI}}$, then z can only belong to either $I^* \cap W$, $M^* \cap W$, or $M^* \cap M$.

Now, irrespective of the presence of AI, the marginal return to knowledge is higher for managers than for independent producers, and it is higher for independent producers than for workers. Hence, $\Delta'(z) = w^{*\prime}(z) - w'(z) \geq 0$ whenever z is in either $I^* \cap W$ or $M^* \cap W$. This implies that if $\Delta(z)$ crosses zero in either of these sets, then it crosses it necessarily from below.

Consider then $z \in M^* \cap M$. Since $M^* = M_p^* \cup M_a^*$, here we have two cases to consider: $z \in M_p^* \cap M$ and $z \in M_a^* \cap M$. If $z \in M_p^* \cap M$, then $\Delta'(z) = n(e_j^*(z; \underline{z}_j^*)) - n(e(z; \hat{z}))$, where $e_j^*(z; \underline{z}_j^*)$ is the employee function in a Type $j = 1, 2$ equilibrium and $e(z; \hat{z})$ employee function in the pre-AI equilibrium. However, by Proposition 4, we know that $e_j^*(z; \underline{z}_j^*) > e(z; \hat{z})$ as every $z \in M_p^* \cap M$ manages better workers post-AI. Hence, in this case, $\Delta'(z) > 0$ also. Consequently, if $\Delta(z)$ crosses zero when $z \in M_p^* \cap M$, then it can only cross it from below.

Finally, consider the possibility that $\Delta(z)$ crosses zero at a $z \in M_a^* \cap M$. In this case, $\Delta'(z) = n(z_{\mathrm{AI}}) - n(e(z; \hat{z})) \gtrless 0$, so $\Delta(z)$ is no longer monotone in z in this set. Note, however, that $\Delta''(z) = -n'(e(z; \hat{z}))e'(z; \hat{z}) < 0$, so $\Delta(z)$ is concave. Moreover, if $M_a^* \neq \emptyset$, then $1 \in M_a^*$. The fact that $\Delta(z)$ is concave and that $\Delta(1) > 0$ then immediately implies that if $\Delta(z)$ crosses zero in this set, then it can only cross once and from below (otherwise, $\Delta(1) \leq 0$ contradicting the fact $\Delta(1) > 0$).

• $z_{\mathrm{AI}} \in M$.— From Lemma B.3, we know that $\inf M^* = z_{\mathrm{AI}}$. Moreover, from Proposition 3, we have that $W^* \supseteq W$ and $M^* \subseteq M$. Consequently, if $z > z_{\mathrm{AI}}$, then $z \in M^* \cap M$ necessarily. Since $M^* = M_p^* \cup M_a^*$, here we have two cases to consider: $z \in M_p^* \cap M$ and $z \in M_a^* \cap M$

If $z \in M_p^* \cap M$, then $\Delta'(z) = n(e_j^*(z; \underline{z}_j^*)) - n(e(z; \hat{z}))$, while if $z \in M_a^* \cap M$, then $\Delta'(z) = n(z_{\mathrm{AI}}) - n(e(z; \hat{z}))$, where $e_j^*(z; \underline{z}_j^*)$ is the employee function in a Type $j = 2, 3$ equilibrium and $e(z; \hat{z})$ the employee function in the pre-AI equilibrium. In either case, $\Delta'(z) > 0$ since Proposition 4 states that $e_j^*(z; \underline{z}_j^*) > e(z; \hat{z})$ and $z_{\mathrm{AI}} > e(z; \hat{z})$ when AI has the knowledge of a pre-AI manager, or if AI has the knife-edge knowledge of a pre-AI worker and a pre-AI manager. Consequently, $\Delta'(z) > 0$ for all $z \geq z_{\mathrm{AI}}$, so if $\Delta(z)$ crosses zero at some $z > z_{\mathrm{AI}}$, then it always crosses it from below. □

Claim C.3. *If $\Delta(z) > 0$ for some $z \in [0, z_{\mathrm{AI}}]$, then $\Delta(z') > 0$ for all $z' \in [0, z]$.*

Proof. Given that $\Delta(z_{\mathrm{AI}}) < 0$, it suffices to show that if $\Delta(z)$ crosses zero at some $z < z_{\mathrm{AI}}$, then it always crosses zero from above. We first consider $z_{\mathrm{AI}} \in W$, and then $z_{\mathrm{AI}} \in \mathrm{int}M$.

• $z_{\text{AI}} \in W$.— Lemma B.3 implies that $\sup W^* = z_{\text{AI}}$, while Proposition 3 that $W^* \subseteq W$ and $M^* \supseteq M$. Consequently, if $z < z_{\text{AI}}$, then $z \in W^* \cap W$, where $W^* = W_a^* \cup W_p^*$.

Now, if $z \in W_p^* \cap W$, then $\Delta'(z) = hw^*(f_j^*(z; \underline{z}_j^*)) - hw(f(z;\hat{z}))$, where $f_j^*(z;\underline{z}_j^*)$ is the matching function in a Type $j = 1, 2$ equilibrium, and $f(z;\hat{z})$ the matching function of the pre-AI equilibrium. However, using the firms' zero-profit condition:

$$\Delta'(z) = hw^*(f_j^*(z;\underline{z}_j^*)) - hw(f(z;\hat{z})) = \frac{f_j^*(z;\underline{z}_j^*) - f(z;\hat{z}) - \Delta(z)}{1-z}$$

Consequently, if $\Delta(z) = 0$ at some z in this interval, say at $z = \zeta$, then $\Delta'(\zeta) = f_j^*(\zeta;\underline{z}_j^*) - f(\zeta;\hat{z}) \leq 0$, where the last inequality follows because every worker is managed by a worse manager post-AI in this case, as shown in Proposition 4.

On the other hand, if $z \in W_a^* \cap W$, then following the same reasoning as before, we have that:

$$\Delta'(z) = hw(z_{\text{AI}}) - hw(f(z;\hat{z})) = \frac{z_{\text{AI}} - f(z;\hat{z}) - \Delta(z)}{1-z}$$

Consequently, if $\Delta(z) = 0$ at some z in this interval, say at $z = \zeta$, then $\Delta'(\zeta) = z_{\text{AI}} - f(\zeta;\hat{z}) \leq 0$, where the inequality follows, again, from Proposition 4.

• $z_{\text{AI}} \in \text{int}M$.— In this case, Lemma B.3 implies that $\inf M^* = z_{\text{AI}}$, while Proposition 3 that $W^* \supset W$ and $M^* \subset M$. Consequently, if $z < z_{\text{AI}}$, then z can only belong to either $I^* \cap M$, $W^* \cap M$, or $W^* \cap W$.

Now, $\Delta'(z) \leq 0$ whenever z is in either $I^* \cap M$ or $W^* \cap M$, given that the marginal return to knowledge is higher for managers than for independent producers, and it is higher for independent producers than for workers. Hence, if $\Delta(z)$ crosses zero in either of these sets, then it crosses it necessarily from above.

Consider then $z \in W^* \cap W$. Since $W^* = W_p^* \cup W_a^*$, we have two cases to consider: $z \in W_p^* \cap W$ and $z \in W_a^* \cap W$. If $z \in W_p^* \cap W$, then:

$$\Delta'(z) = hw^*(f_j^*(z;\underline{z}_j^*)) - hw(f(z;\hat{z})) = \frac{f_j^*(z;\underline{z}_j^*) - f(z;\hat{z}) - \Delta(z)}{1-z}$$

where $f_j^*(z;\underline{z}_j^*)$ is the matching function in a Type $j = 2, 3$ equilibrium, and $f(z;\hat{z})$ the matching function of the pre-AI equilibrium. Consequently, if $\Delta(z) = 0$ at some z in this interval, say at $z = \zeta$, then $\Delta'(\zeta) = f_j^*(\zeta;\underline{z}_j^*) - f(\zeta;\hat{z}) \leq 0$, where the last inequality follows because every $z \in W^* \cap W$ is managed by a worse manager post-AI when AI has the knowledge of a pre-AI manager.

Finally, consider the possibility that $\Delta(z)$ crosses zero at a $z \in W_a^* \cap W$. In this case, $\Delta'(z) = hz_{\text{AI}} - hw(f(z;\hat{z}))$, so $\Delta''(z) = -hw'(f(z;\hat{z}))f'(z;\hat{z}) < 0$, implying that $\Delta(z)$ is concave. Moreover, if $W_a^* \neq \emptyset$, then $0 \in W_a^*$, and we know that $\Delta(0) > 0$ in this case. Consequently, if $\Delta(z)$ crosses zero in this set, then it can only cross once and from above (otherwise, $\Delta(0) \leq 0$ contradicting the fact $\Delta(0) > 0$). □

C.4 Proof of Proposition 5

Part (i) ("there always exists z strictly greater than z_{AI} such that $\Delta(z) > 0$") follows directly from Lemma 2 and the fact that $\Delta(1) > 0$ (see Claim C.1). Hence, it only remains to prove part (ii) ("there exists z strictly smaller than z_{AI} such that $\Delta(z) > 0$ if and only if $z_{\text{AI}} > \bar{z}_{\text{AI}}$, where $\bar{z}_{\text{AI}} \in (0, \hat{z})$"). To prove this part, we begin by constructing $\bar{z}_{\text{AI}}$ and then show that the statement is true.

Let $\Delta(0; z_{\text{AI}}) \equiv w^*(0; z_{\text{AI}}) - w(0)$, and define $\bar{z}_{\text{AI}}$ as the solution to $\Delta(0; \bar{z}_{\text{AI}}) = 0$. We first show that $\bar{z}_{\text{AI}}$ exists and is unique and that $\bar{z}_{\text{AI}} \in (0, \hat{z})$. We do this by showing that $\Delta(0; z_{\text{AI}})$ crosses zero once as we go from $z_{\text{AI}} = 0$ to $z_{\text{AI}} = 1$, and that this crossing point is at a $z_{\text{AI}} < \hat{z}$. Indeed, if $z_{\text{AI}} \geq \hat{z}$, then Claim C.1 states that $\Delta(0; z_{\text{AI}}) > 0$ (as $z_{\text{AI}} \in M$ in this case). Moreover, as shown in Lemma B.2, $0 \in \mathcal{R}_1$, so when $z_{\text{AI}} = 0$, the equilibrium is always Type 1. The latter implies that $\Delta(0; 0) = \underline{z}_1^*(0)(1 - h) - w(0) = -w(0) < 0$,[32] where the second-to-last equality follows because $\underline{z}_1^*(0) = 0$, as can be easily be proven from the condition that determines $\underline{z}_1^*(z_{\text{AI}})$ (see the statement of Lemma B.3).

Now, when $z \in (0, \hat{z}) = W$, the equilibrium is either Type 1, in which case $w^*(0; z_{\text{AI}}) = \underline{z}_1^*(z_{\text{AI}})(1 - h)$, or Type 2, in which case $w^*(0; z_{\text{AI}}) = z_{\text{AI}}(1 - h)$. Using the equilibrium condition that determines $\underline{z}_1^*(z_{\text{AI}})$, it is not difficult to prove that (i) $\underline{z}_1^*(z_{\text{AI}})$ is strictly increasing in z_{AI}, and that (ii) $\underline{z}_1^*(z_{\text{AI}}) = z_{\text{AI}}$ whenever we switch from a Type 1 into a Type 2 equilibrium (and vice versa). Consequently, irrespective of the equilibrium type in this region, $w^*(0; z_{\text{AI}})$ is continuous and strictly increasing in z_{AI}, which implies that so is $\Delta(0; z_{\text{AI}})$. This result, combined with the fact that $\Delta(0; 0) < 0$ and $\Delta(0; \hat{z}) > 0$, immediately yields the desired result.

Having constructed $\bar{z}_{\text{AI}}$, we prove that there exists $z < z_{\text{AI}}$ such that $\Delta(z; z_{\text{AI}}) > 0$ if and only if $z_{\text{AI}} > \bar{z}_{\text{AI}}$. First we show that if there exists a $z < z_{\text{AI}}$ such that $\Delta(z; z_{\text{AI}}) > 0$, then $z_{\text{AI}} > \bar{z}_{\text{AI}}$. To do this, we prove the contrapositive statement: If $z_{\text{AI}} \leq \bar{z}_{\text{AI}}$, there is no such z. Indeed, $\Delta(0; z_{\text{AI}}) \leq 0$ for all $z_{\text{AI}} \leq \bar{z}_{\text{AI}}$ as shown above. Hence, Lemma 2 immediately implies that $\Delta(z; z_{\text{AI}}) \leq 0$ for all $z < z_{\text{AI}}$.

Now we prove that if $z_{\text{AI}} > \bar{z}_{\text{AI}}$, then there exists a $z < z_{\text{AI}}$ such that $\Delta(z; z_{\text{AI}}) > 0$. Indeed, as shown above, $z_{\text{AI}} > \bar{z}_{\text{AI}}$ then $\Delta(0; z_{\text{AI}}) > 0$. Consequently, Lemma 2 immediately implies that there exists $\zeta \in [0, z_{\text{AI}})$ such that $\Delta(z; z_{\text{AI}}) > 0$ for $z \in [0, \zeta)$ and $\Delta(z; z_{\text{AI}}) < 0$ for $z \in (\zeta, z_{\text{AI}}]$. □

References

ACEMOGLU, D. AND D. AUTOR (2011): "Chapter 12 - Skills, Tasks and Technologies: Implications for Employment and Earnings," Elsevier, vol. 4 of *Handbook of Labor Economics*, 1043–1171.

ACEMOGLU, D. AND J. LOEBBING (2024): "Automation and Polarization," Working Paper.

[32] To avoid any confusion, here we are making explicit that $\underline{z}_1^*(z_{\text{AI}})$ depends on z_{AI} as can be seen from the condition that determines $\underline{z}_1^*(z_{\text{AI}})$ in the statement of Lemma B.3.

Acemoglu, D. and P. Restrepo (2018): "The Race Between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment," *American Economic Review*, 108, 1488–1542.

——— (2019): "8. Artificial Intelligence, Automation, and Work," in *The Economics of Artificial Intelligence: An Agenda*, ed. by A. Agrawal, J. Gans, and A. Goldfarb, University of Chicago Press, 197–236.

——— (2022): "Tasks, Automation, and the Rise in U.S. Wage Inequality," *Econometrica*, 90, 1973–2016.

Aghion, P., B. F. Jones, and C. I. Jones (2017): "Artificial Intelligence and Economic Growth," NBER Working Paper No. 23928.

Alfred Sloan (1924): "The Most Important Thing I Ever Learned about Management," *System: The Magazine of Business*, 194, 140–141.

Antràs, P., L. Garicano, and E. Rossi-Hansberg (2006): "Offshoring in a Knowledge Economy," *The Quarterly Journal of Economics*, 121, 31–77.

Autor, D. (2024): "Applying AI to Rebuild Middle Class Jobs," NBER Working Paper No. 32140.

Azar, J., M. Chugunova, K. Keller, and S. Samila (2023): "Monopsony and Automation," Max Planck Institute for Innovation & Competition Research Paper No. 23-21.

Brynjolfsson, E. (2022): "The Turing Trap: The Promise & Peril of Human-Like Artificial Intelligence," *Daedalus*, 151, 272–287.

Brynjolfsson, E., D. Li, and L. R. Raymond (2023): "Generative AI at Work," NBER Working Paper No. 31161.

Brynjolfsson, E. and A. McAfee (2016): *The Second Machine Age*, WW Norton.

Caicedo, S., J. Lucas, Robert E., and E. Rossi-Hansberg (2019): "Learning, Career Paths, and the Distribution of Wages," *American Economic Journal: Macroeconomics*, 11, 49–88.

Caliendo, L., G. Mion, L. D. Opromolla, and E. Rossi-Hansberg (2020): "Productivity and Organization in Portuguese Firms," *Journal of Political Economy*, 128, 4211–4257.

Caliendo, L., F. Monte, and E. Rossi-Hansberg (2015): "The Anatomy of French Production Hierarchies," *Journal of Political Economy*, 123, 809–852.

Caliendo, L. and E. Rossi-Hansberg (2012): "The Impact of Trade on Organization and Productivity," *The Quarterly Journal of Economics*, 127, 1393–1467.

Carmona, G. and K. Laohakunakorn (Forthcoming): "Improving the Organization of Knowledge in Production by Screening Problems," *Journal of Political Economy*.

Eeckhout, J. and P. Kircher (2018): "Assortative Matching with Large Firms," *Econometrica*, 86, 85–132.

Fuchs, W., L. Garicano, and L. Rayo (2015): "Optimal Contracting and the Organization of Knowledge," *The Review of Economic Studies*, 82, 632–658.

Garicano, L. (2000): "Hierarchies and the Organization of Knowledge in Production," *Journal of Political Economy*, 108, 874–904.

GARICANO, L. AND T. HUBBARD (2007): "Managerial Leverage Is Limited by the Extent of the Market: Hierarchies, Specialization, and the Utilization of Lawyers' Human Capital," *The Journal of Law & Economics*, 50, 1–43.

GARICANO, L. AND T. N. HUBBARD (2012): "Learning About the Nature of Production from Equilibrium Assignment Patterns," *Journal of Economic Behavior & Organization*, 84, 136–153.

GARICANO, L. AND E. ROSSI-HANSBERG (2004): "Inequality and the Organization of Knowledge," *American Economic Review Papers and Proceedings*, 94, 197–202.

——— (2006): "Organization and Inequality in a Knowledge Economy," *The Quarterly Journal of Economics*, 121, 1383–1435.

——— (2012): "Organizing growth," *Journal of Economic Theory*, 147, 623–656.

——— (2015): "Knowledge-Based Hierarchies: Using Organizations to Understand the Economy," *Annual Review of Economics*, 7, 1–30.

HAYEK, F. A. (1945): "The Use of Knowledge in Society," *American Economic Review*, 35, 519–530.

JOHNSON, S. AND D. ACEMOGLU (2023): *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity*, Hachette UK.

KORINEK, A. AND D. SUH (2024): "Scenarios for the Transition to AGI," NBER Working Paper No. 32255.

MESEROLE, C. O. (2018): "What is Machine Learning?" Tech. rep., Brookings Institution.

MURO, M., J. WHITON, AND R. MAXIM (2019): "What Jobs are Affected by AI?" Tech. rep., Brookings Institution.

NORDHAUS, W. D. (2007): "Two Centuries of Productivity Growth in Computing," *The Journal of Economic History*, 67, 128–159.

NOY, S. AND W. ZHANG (2023): "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence," *Science*, 381, 187–192.

THOMSON REUTERS (2021): "The Business Case for AI-Enabled Legal Technology," Tech. rep., Thomson Reuters.

WEBB, M. (2020): "The Impact of Artificial Intelligence on the Labor Market," Working Paper.

ZEIRA, J. (1998): "Workers, Machines, and Economic Growth," *The Quarterly Journal of Economics*, 113, 1091–1117.