Artificial Intelligence in the Knowledge Economy*

Enrique Ide & Eduard Talamàs

February 12, 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" (i.e., execute routine tasks), while more knowledgeable agents become "managers" (i.e., specialize in problem-solving). When AI's knowledge is that of a pre-AI worker, the introduction of AI (i) displaces humans from routine to managerial work, (ii) decreases the size of the largest and smallest firms, (iii) decreases the productivity of the least productive firms, (iv) increases the size of the firms managed by all but the best managers, and (v) decreases the productivity of all workers, as it worsens the quality of their managers. The first three results are reversed when AI has the knowledge of a pre-AI manager, but not the last two: In that case, all managers supervise larger firms, and the productivity of all but the least knowledgeable workers decreases after AI's introduction. AI creates winners and losers, and the winners are always at the extremes of the knowledge distribution. Only the most knowledgeable humans win when AI's knowledge is low. In contrast, only the least knowledgeable humans benefit when AI's knowledge and communication costs are high. By explicitly incorporating the endogenous organization of knowledge, our results provide new insights into the effects of AI on the future of work.

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).

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

*IESE Business School, C/ de Arnús i de Garí 3-7, 08034 Barcelona, Spain (eide@iese.edu and etalamas@iese.edu). We thank Andrea Galeotti for useful comments on an earlier draft. All errors are our own. We declare we have no relevant or material financial interests that relate to the research described in this paper.

tasks previously reserved for highly skilled workers with sophisticated knowledge (Webb, 2020). Consequently, it is not immediately apparent that experience with previous waves of automation will seamlessly extrapolate to the case of AI.[1]

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 the effects of AI on the future of work, it is imperative that we consider how economic agents organize themselves in terms of knowledge and time to achieve common production goals.

As we show below, this approach leads to new insights. For example, while there is a growing interest in understanding the extent to which AI increases productivity at any given task (e.g., Brynjolfsson et al., 2023; Eloundou et al., 2023), we show how, once organizational effects are taken into account, AI may lead to a *reduction* in the productivity of a significant subset of the population.

Our starting point 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 such a setting—which we consider 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.[2]

Our contribution is to incorporate AI into this otherwise canonical setting. We model AI as an algorithm that can use processing capacity (i.e., "capital") to mimic the behavior of a human with a given knowledge. As standard in the automation literature, we assume that all firms have access to this technology. Firms that use AI are identical to those that do not, except that they rent processing

[1]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."

[2]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."

capacity instead of hiring labor to do either production or managerial work.[3] Moreover, we assume that AI can be used at scale, in the sense that processing capacity is large relative to human time.

We then characterize the competitive equilibrium with AI. We show that when AI is used as a worker, it is necessarily the *most* knowledgeable worker, so it is supervised by the *most* knowledgeable humans. In addition, when AI is used as a manager, it is necessarily the *least* knowledgeable manager, so it supervises the *least* knowledgeable humans.

A notable property of the equilibrium is that, as long as AI has unlimited potential applications, the rental rate of processing capacity 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. Intuitively, even though processing capacity is large relative to human time, it is still scarce relative to its potential uses. 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 since they are the least likely to succeed on their own.[4]

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, and (iii) 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, firms have access to a relatively abundant substitute for workers with that knowledge. This reduces their 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. The same intuition, though in reverse, explains why AI displaces the least knowledgeable managers into 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) and productivity (defined as firm output divided by the time/processing capacity it uses). In the case of two-layer organizations, for example, the knowledge of the best worker determines the maximum firm size, while the knowledge of the worst manager determines the minimum firm size and the productivity of the least productive firm. Hence, when AI has the knowledge of a

[3] We normalize the units of processing capacity so that one unit of time is equivalent to one unit of processing capacity.

[4] In Section 5, we consider the case where processing capacity 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 processing capacity is zero and AI *does* lead 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.

pre-AI worker, the size of the biggest and the smallest firm and the productivity of the least productive firm decrease with AI (as the best pre-AI workers become managers post-AI). Similarly, when AI has the knowledge of a pre-AI manager, the size of the biggest and the smallest firm and the productivity of the least productive firm increase with AI (as the worst pre-AI managers become workers post-AI).[5]

After examining the repercussions of human displacement across occupations, we analyze the effects of AI on the workers and managers *who have not been displaced*. For workers, we focus on productivity (defined as their expected output). For managers, we focus on the size of the firm they supervise. We show that after the introduction of AI, all workers become less productive, except possibly the least knowledgeable ones, while all managers supervise larger firms, except possibly the very best ones. Interestingly, this result holds irrespective of whether AI has the knowledge of a pre-AI worker or a pre-AI manager.

To understand these results, note that a worker's productivity increases with her manager's knowledge, while the size of the firm that a manager supervises increases with her workers' knowledge. When AI has the knowledge of a pre-AI worker, the knowledge threshold to become a manager decreases, worsening the pool of managers available for non-displaced workers. At the same time, the pool of workers available for the non-displaced managers improves as the least knowledgeable agents match with the newly created class of managers. Similar forces arise when AI has the knowledge of a pre-AI manager: In such a case, the knowledge of the best worker increases, improving the pool of available workers for non-displaced managers and, therefore, leaving a worse pool of managers for non-displaced workers.

Finally, we study the effects of AI on labor income. The analysis is subtle 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, AI is a relatively abundant substitute for humans with a certain knowledge, so it also affects 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 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 related literature, Section 2 presents our model. After describing the pre-AI equilibrium in Section 2.3, we then characterize the equilibrium with AI in Section 3. In Section 4, we turn to our main endeavor: The comparison of the pre- and post-AI equilibria; this section contains our main results. Section 5 discusses how our results extend to the cases in which (i) processing capacity is small and (ii) processing capacity is large relative to its potential applications. Finally, Section 6 concludes.

[5] In Section 4.2, we also characterize the effects of AI on the size and productivity distribution of all firms. We show that the results are similar, except that the effects of AI on the size of the smallest firm are less clear-cut.

Related literature

This paper makes two contributions. First, it introduces automation and artificial intelligence to the literature on knowledge hierarchies. Second, it incorporates organizations and a different way of thinking about AI into 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.[6] 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 while offshoring gives firms access to a population of humans with different knowledge, artificial intelligence gives firms access to an algorithm that can solve problems *at scale* (up to a given difficulty). Consequently, the effects of AI are qualitatively different than those of offshoring.[7]

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; Acemoglu and Loebbing, 2022; Azar et al., 2023). 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 automation is to decrease the labor share, the economy might still exhibit a balanced growth path in which the labor share stays constant in 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.

[6] 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.

[7] Indeed, many of our results hinge on the fact that, in equilibrium, AI acts as the best worker and/or the worst manager, and there is no analog to this result in Antràs et al. (2006).

We depart from this literature in two key respects. First, we explicitly integrate organizations into the analysis by employing knowledge hierarchies. This is a distinct specification within the task-based framework where tasks are hierarchical, and the relationship between the different tasks is based on an explicit organizational problem. By incorporating organizations into our analysis, we present novel insights into the impact of AI on the size and productivity distributions of firms. Moreover, our approach offers fresh perspectives on how AI influences labor outcomes through its effects on 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 an exogenously determined subset of 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 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.[8]

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 associated density function $g(z)$. 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—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.

[8] Acemoglu and Loebbing (2022) provide an alternative assignment model where the tasks being automated are endogenously determined. Their model, however, is different than ours in that tasks are not hierarchical and there are no organizations.

The difficulty of each problem is ex-ante unknown and distributed uniformly on $[0, 1]$, independently across problems.[9] If the worker's knowledge exceeds the problem's difficulty, she solves the problem and produces one unit of output. Otherwise, the worker produces no output. We normalize the value of each unit of output to one.

Additionally, the firm can hire another human to act as its 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 $h \in [\underline{h}, 1]$ units of the manager's time, where $\underline{h} > 0$. If the manager knows the solution to a problem, she communicates the solution to the corresponding worker, who then produces a unit of output. Otherwise, no production takes place.

Artificial Intelligence.— We model AI as an algorithm that can use processing capacity to mimic the behavior of a human with knowledge $z_{\text{AI}} \in [0, \bar{z}_{\text{AI}}] \subseteq [0, 1)$. Hence, AI is an automation technology. We refer to z_{AI} as "AI's knowledge."

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 AI. Firms that use AI are identical to those that do not, except that they use AI instead of humans in either production work or problem-solving (i.e., managerial roles). To do this, they must rent one unit of processing capacity per production worker or manager that they replace. The amount of processing capacity in the economy—which is exogenously given—is denoted by $\mu \geq 0$.

We mainly focus on the case in which processing capacity is abundant relative to human time:

$$\frac{1}{\underline{h}(1 - \bar{z}_{\text{AI}})} \leq \mu \tag{1}$$

This inequality guarantees that for all parameter values, the binding constraint in human-AI interactions is human time, not processing capacity.[10]

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

As mentioned above, 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).

[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 most intensive (though not necessarily optimal) way to use processing capacity in organizations comprising humans and AI is to use it for production work. Since there is a unit mass of humans and the maximum number of production workers with knowledge z_{AI} that a human can manage is $[h(1 - z_{\text{AI}})]^{-1}$, condition (1) implies that, for all parameter values, there are not enough humans to exhaust all processing capacity.

The profit of a single-layer organization is simply:

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

The profit of a two-layer organization, in turn, depends on whether it does not use AI (i.e., it does not use automation, 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 since the equilibrium exhibits positive assortative matching.

Consequently, in order 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:

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

where z and m denote the knowledge of a human worker and manager, respectively.

For brevity, we will 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.

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. Moreover, 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 mass of processing capacity rented for independent production, production in two-layer firms, and supervision of humans, respectively.

Definition (Competitive Equilibrium). An equilibrium consists of non-negative masses (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 processing capacity $r \in \mathbb{R}_+$, 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.

[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.

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.

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 use $S \preceq S'$ to denote 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).[12] 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. Moreover, for any function $h(x)$ with domain H, we use $h(X)$ to denote $\lim_{x \to X} h(x)$ whenever X is a limit point of the domain (even if $X \notin H$).

2.2 Discussion of the Model

Before moving on to the main analysis, we briefly comment on some assumptions underlying our model. First, we focus on the case where processing capacity is large relative to time as the hallmark of AI lies in its ability to operate *at scale*. This scalability is a defining feature of automation technologies, distinguishing them from other input supply shocks, such as immigration and offshoring. In Section 5.1, we explore the role of this assumption by characterizing the opposite extreme in which processing capacity is small.

Second, our model assumes that AI can perform the same roles in the knowledge economy as humans and that there are unlimited uses for AI (i.e., processing capacity is abundant relative to human time but scarce relative to its potential applications). In Section 5.2, we discuss how our results change when processing capacity is abundant relative to its uses, while in Section 5.4, we provide examples that illustrate how AI can be a worker, a manager, or an independent producer.

Third, our main goal with this paper is to analyze how AI affects human labor outcomes (i.e., its effects on job displacement, productivity, and wages, among others). For this reason, we take AI technology and the economy's processing capacity as given and omit the implications of AI for its developers and for capital markets more generally. Studying firms' incentives to develop AI or increase the economy's processing capacity is an intriguing avenue for future research.

Fourth, we assume that AI cannot mimic the most knowledgeable human, i.e., $\bar{z}_{\text{AI}} < 1$. Moreover, 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 assump-

[12] Defining $S \preceq S'$ even if S or S' may be empty allows us to be more concise. For example, we will write $W_a \preceq W_p$ and $M_p \preceq M_a$ even if W_a and/or M_a are empty (in which case $\sup W_a = \infty$ and $\inf M_a = -\infty$).

tions primarily for the sake of simplicity, although we believe they offer a good first approximation to the problem at hand. Understanding the effects of artificial superintelligence and the impact of AI when human retraining or more complex organizations are feasible is left for future research.

Fifth 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.[13] However, introducing this large measure of anonymous firms allows us to present the model in the most concise and cleanest possible way.

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 processing capacity μ is zero and was originally described by Fuchs et al. (2015).[14] 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 a $\hat{z} \in (0,1)$ 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 (that is, the equilibrium matching function $f : W \to M$ 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 $z \in \text{int} W$.*
- *$w(z) = z$ 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. 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 one for managers (as they can leverage their knowledge by applying it to more than one problem), exactly equal to one for independent producers

[13] Fuchs et al. (2015) refer to the case where workers are the owners of firms as a "consultant market arrangement," while they refer to the case where managers are the owners of firms as a "referral market arrangement."

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

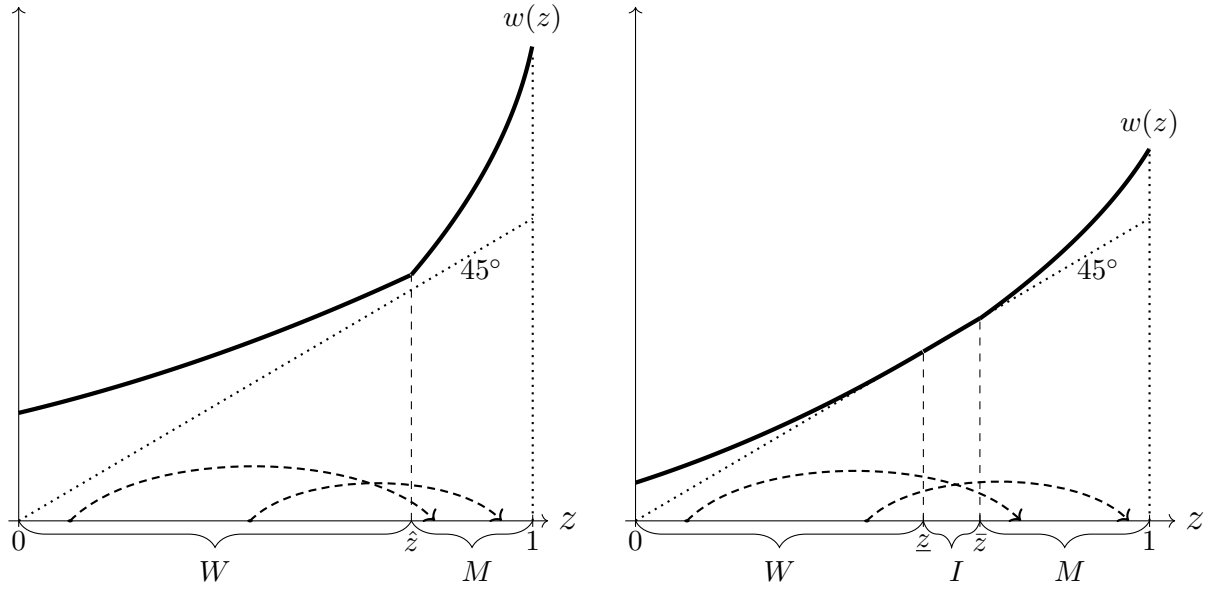


Figure 1: Illustration of the pre AI equilibrium when $h \leq h_0$ (left) and $h > h_0$ (right).

Notes. $G(z) = z$, which implies that $h_0 = 3/4$. The left panel corresponds to $h = 1/2$, while the right panel to $h = 0.8125$. The thick line depicts the wage function $w(z)$. The dashed arrows illustrate the matching function $f(z)$ at two arbitrary points. See Section 5 of the online Appendix for a detailed characterization of the pre-AI equilibrium when human knowledge is uniformly distributed.

(as expected output equals knowledge for them), and strictly lower than one for workers (as their knowledge is used to free up managerial time).[15]

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 one, 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 than one, if workers earned their expected output as independent producers, all two-layer firms would only want to hire the agents with the least knowledge.[16]

[15] A marginal increase in the knowledge of a worker with knowledge z liberates $h < 1$ units of her manager's time. This allows her firm to hire $1/(1-z)$ extra workers, with expected net output gain of $(f(z) - w(z))/(1-z) < 1$ (as the wage $w(z)$ of that worker is strictly greater than z in the interior of W).

[16] To be more precise, suppose there is an interval $[z', z'']$ such that, agents with knowledge $z \in [z', z'']$ earn their expected output as independent producers. Note that no two-layer firm is willing to hire agents with knowledge $z \in (z', z'')$. Indeed,

Fourth and finally, the income function $w(z)$ is continuous, strictly increasing, and weakly convex. Monotonicity follows 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.[17]

The identity of the best worker and that of the worst manager in the pre-AI equilibrium will be important in the development that follows. The next corollary describes how these identities evolve with changes in the communication cost h:

Corollary 1. *Let $\hat{z}(h)$, $\underline{z}(h)$, and $\bar{z}(h)$ be the cutoffs of the pre-AI equilibrium as a function of h. Then,*

- *$\hat{z}(h)$ is strictly decreasing in $h \in [\underline{h}, h_0]$ with $\hat{z}(h_0) = z_0$, where $z_0 \in (0,1)$. Moreover, $\hat{z}(\underline{h}) \to 1$ as $\underline{h} \to 0$.*
- *$\underline{z}(h)$ is strictly decreasing in $h \in (h_0, 1]$, with $\underline{z}(h_0) = z_0$, and $\underline{z}(1) = 0$.*
- *$\bar{z}(h_0) = z_0$, and $\bar{z}(1) = 1$.*

Proof. See Appendix A. □

3 The AI Equilibrium

We now characterize the equilibrium with AI. For future reference, we index this equilibrium using the superscript "$*$" (note that the pre-AI equilibrium had no superscript).

3.1 Overview of the Equilibrium

Prior to presenting the formal characterization, we offer an overview of the AI equilibrium. Given that processing capacity is abundant relative to human time, there are not enough humans to interact with AI inside two-layer organizations. Thus, a positive measure of processing capacity must necessarily be used for independent production. The question is whether AI will be used in any other capacity. This leads us to define four potential configurations for the AI equilibrium.

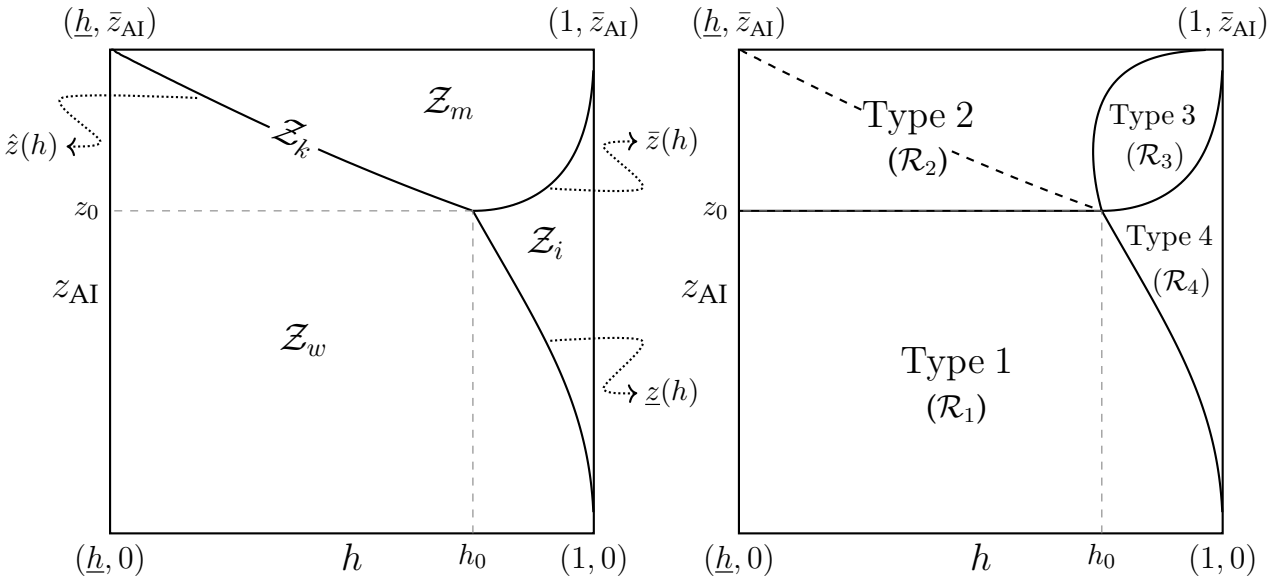
Definition. An AI equilibrium is:

- *Type 1* if AI is used as a worker but not as a manager: $W_a^* = \emptyset$, $M_a^* \neq \emptyset$.
- *Type 2* if AI is used both as a worker and as a manager: $W_a^* \neq \emptyset$, $M_a^* \neq \emptyset$.
- *Type 3* if AI is used as a manager but not as a worker: $W_a^* \neq \emptyset$, $M_a^* = \emptyset$.
- *Type 4* if AI is neither used as a worker nor as a manager: $W_a^* = M_a^* = \emptyset$.

---

if a firm hired a manager with knowledge $z \in (z', z'')$, it could strictly increase its profits by hiring the same workers and a manager with knowledge z''. Similarly, if a two-layer firm hired workers with knowledge in $z \in (z', z'')$, it could strictly increase its profits by hiring the same manager and workers with knowledge z'.

[17] If the income function discontinuously jumped up (down) at some knowledge level, there would be no demand for agents with knowledge just above (below) that level.



(a) AI's knowledge Equivalence

(b) The Four Types of Equilibrium

Figure 2: An Overview of the Equilibrium

Notes. $G(z) = z$, $\underline{h} = 0.01$, $\bar{z}_{\text{AI}} = 0.99$ (so $\mu > 10.000$ is the uniform bound over processing capacity required in this case). When human knowledge is uniformly distributed, then (i) $(h_0, z_0) = (3/4, 2/3)$, (ii) $\bar{z}(h)$ is strictly increasing in h, and (iii) the boundary between a Type 1 and a Type 2 equilibrium is a straight line at $z_0 = 2/3$ (though these results are not true in general). See Section 5 of the online Appendix for a detailed characterization of the AI equilibrium when human knowledge is uniformly distributed.

Figure 2 then provides an overview of the equilibrium. In particular, panel (a) illustrates the regions of the parameter space where AI's knowledge is that of (i) a pre-AI worker (denoted by $\mathcal{Z}_w$), (ii) a pre-AI manager (denoted $\mathcal{Z}_m$), (iii) a pre-AI worker and simultaneously a pre-AI manager (denoted $\mathcal{Z}_k$, where "k" stands for "knife-edge"), and (iv) a pre-AI independent producer (denoted $\mathcal{Z}_i$). Formally:[18]

$$\mathcal{Z}_w := \{(h, z_{\text{AI}}) \mid z_{\text{AI}} < \hat{z}(h) \text{ if } h \leq h_0 \text{ and } z_{\text{AI}} < \underline{z}(h) \text{ if } h > h_0\}$$

$$\mathcal{Z}_m := \{(h, z_{\text{AI}}) \mid z_{\text{AI}} > \hat{z}(h) \text{ if } h \leq h_0 \text{ and } z_{\text{AI}} > \bar{z}(h) \text{ if } h > h_0\}$$

$$\mathcal{Z}_k := \{(h, z_{\text{AI}}) \mid z_{\text{AI}} = \hat{z}(h) \text{ and } h \leq h_0\}$$

$$\mathcal{Z}_i := \{(h, z_{\text{AI}}) \mid \underline{z}(h) \leq z_{\text{AI}} \leq \bar{z}(h) \text{ and } h > h_0\}$$

Figure 2 panel (b), in turn, provides insights into the nature of the AI equilibrium within each of these regions. This figure illustrates several noteworthy results. First, each of the four different configurations (Type 1, 2, 3, and 4) can arise as an equilibrium depending on the parameter values.

[18] Strictly speaking, to be fully consistent, we should be separating the case where AI has the knowledge of an independent producer into three different cases: (a) AI has the knowledge of a "strict" independent producer, (b) AI has the knowledge of both an independent producer and a worker, and (c) AI has the knowledge of both an independent producer and a manager. However, this finer separation does not bring any new insights, so we have combined these three cases together into a single one.

However, when AI possesses the knowledge of a pre-AI worker, the resulting AI equilibrium is either Type 1 or Type 2. That is, AI is used as a worker and may also be used as a manager.

Similarly, if AI's knowledge is that of a pre-AI manager, the AI equilibrium is either Type 2 or Type 3. Hence, AI is used as a manager and may also be used as a worker. Lastly, when AI possesses the knife-edge knowledge of a pre-AI worker and pre-AI manager, the equilibrium is necessarily of Type 2 (AI is used both as worker and manager), and when it has the knowledge of a pre-AI independent producer, the AI equilibrium is necessarily of Type 4 (AI is only used for independent production).

3.2 Formal Characterization

We now formally characterize the equilibrium. The starting point is the following set of results:

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

1. ***The Rental Rate of Processing Capacity is equal to AI's Knowledge:*** $r^* = z_{AI}$.
2. ***Occupational Stratification:*** $W^* \preceq I^* \preceq M^*$.
3. ***AI is the Best Worker and/or the Worst Manager:*** $W^* \preceq \{z_{AI}\} \preceq M^*$.
4. ***Positive Assortative Matching:*** *The matching function* $f^* : W_p^* \to M_p^*$ *is strictly increasing and* $W_a^* \preceq W_p^*$ *and* $M_p^* \preceq M_a^*$.

Proof. See Appendix B. □

Lemma 1 states several important features of the AI equilibrium when processing capacity is abundant. First, the equilibrium rental rate of processing capacity is equal to AI's knowledge, as some processing capacity must necessarily be used for independent production, and the single-layer firms using AI must obtain zero profits.[19] 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).

Finally, if AI is used as a worker in equilibrium, then it is the *most* knowledgeable worker, while if it is used as a manager, then it is the *least* knowledgeable manager. This property follows from occupational stratification plus the fact that some processing capacity must necessarily be used for independent production. Positive assortative matching then implies that AI, when used as a worker, is managed by the most knowledgeable humans, while, when used as a manager, it manages the least knowledgeable humans.

A direct implication of Lemma 1 is that the four potential configurations for the AI equilibrium—if they arise—must take a very particular form:

[19] Note that the price of processing capacity does not go to zero despite its abundance. This is because processing capacity is abundant relative to human time but not relative to its potential applications. In Section 5.2, we discuss the case where processing capacity is larger than AI's potential uses.

Corollary 2.

- *If the equilibrium is Type 1, then there exist cutoffs $z_{AI} < \underline{z}_1^* \leq \bar{z}_1^* \leq 1$ such that:*

$$W_a^* = \emptyset,\ W_p^* = [0, z_{AI}],\ I^* = (z_{AI}, \underline{z}_1^*),\ M_p^* = [\underline{z}_1^*, \bar{z}_1^*],\ M_a^* = [\bar{z}_1^*, 1]$$
$$\textit{Hence } \mu_w^* = \textstyle\int_{\bar{z}_1^*}^1 n(z_{AI})dG(z),\ \mu_m^* = 0,\ \mu_i^* = \mu - \mu_w^*$$

- *If the equilibrium is Type 2, then there exist cutoffs $0 \leq \underline{z}_2^* \leq z_{AI} \leq \bar{z}_2^* \leq 1$ such that:*

$$W_a^* = [0, \underline{z}_2^*],\ W_p^* = [\underline{z}_2^*, z_{AI}],\ I^* = \emptyset,\ M_p^* = [z_{AI}, \bar{z}_2^*],\ M_a^* = [\bar{z}_2^*, 1]$$
$$\textit{Hence } \mu_w^* = \textstyle\int_{\bar{z}_2^*}^1 n(z_{AI})dG(z),\ \mu_m^* = \int_0^{\underline{z}_2^*} n(z)^{-1}dG(z),\ \mu_i^* = \mu - \mu_w^* - \mu_m^*$$

- *If the equilibrium is Type 3, then there exist cutoffs $0 \leq \underline{z}_3^* \leq \bar{z}_3^* < z_{AI}$ such that:*

$$W_a^* = [0, \underline{z}_3^*],\ W_p^* = [\underline{z}_3^*, \bar{z}_3^*],\ I^* = (\bar{z}_3^*, z_{AI}),\ M_p^* = [z_{AI}, 1],\ M_a^* = \emptyset$$
$$\textit{Hence } \mu_w^* = 0,\ \mu_m^* = \textstyle\int_0^{\bar{z}_3^*} n(z)^{-1}dG(z),\ \mu_i^* = \mu - \mu_m^*$$

- *If the equilibrium is Type 4, then there exist cutoffs $0 \leq \underline{z}_4^* < \bar{z}_4^* \leq 1$ such that:*

$$W_a^* = \emptyset,\ W_p^* = [0, \underline{z}_4^*],\ I^* = (\underline{z}_4^*, \bar{z}_4^*) \ni z_{AI},\ M_p^* = [\bar{z}_4^*, 1],\ M_a^* = \emptyset$$
$$\textit{Hence } \mu_w^* = 0,\ \mu_m^* = 0,\ \mu_i^* = \mu$$

Proof. This corollary is a direct implication of Lemma 1. □

With this in mind, consider then the following partition of the (h, z_{AI}) parameter space:

$$\mathcal{R}_1 \equiv \mathcal{Z}_w \cap \{(h, z_{\text{AI}}) : \Gamma_w(h, z_{\text{AI}}) < 0\} \qquad \mathcal{R}_3 \equiv \mathcal{Z}_m \cap \{(h, z_{\text{AI}}) : \Gamma_m(h, z_{\text{AI}}) < 0\}$$
$$\mathcal{R}_2 = (\mathcal{R}_1 \cup \mathcal{R}_3 \cup \mathcal{R}_4)^C \qquad \mathcal{R}_4 \equiv \mathcal{Z}_i$$

where $\Gamma_w(h, z_{\text{AI}})$ is defined for $(h, z_{\text{AI}}) \in \mathcal{Z}_w \cup \mathcal{Z}_k$ as follows:

$$\Gamma_w(h, z_{\text{AI}}) \equiv n(z_{\text{AI}})(f_w(z_{\text{AI}}; z_{\text{AI}}) - z_{\text{AI}}) - z_{\text{AI}} - \textstyle\int_{z_{\text{AI}}}^{f_w(z_{\text{AI}}; z_{\text{AI}})} n(e_w(z; z_{\text{AI}}))dz \tag{2}$$
$$\text{where } G(f_w(z; z_{\text{AI}})) = G(z_{\text{AI}}) + \textstyle\int_0^z h(1-u)g(u)du \ \text{ for } z \in [0, z_{\text{AI}}]$$

While $\Gamma_m(h, z_{\text{AI}})$, which is defined for $(h, z_{\text{AI}}) \in \mathcal{Z}_m \cup \mathcal{Z}_k$, is given by:

$$\Gamma_m(h, z_{\text{AI}}) \equiv \tfrac{1}{h} - z_{\text{AI}} - \textstyle\int_{z_{\text{AI}}}^1 n(e_m(z; z_{\text{AI}}))dz \tag{3}$$
$$\text{where } G(z) = 1 - \textstyle\int_{e_m(z; z_{\text{AI}})}^{z_{\text{AI}}} h(1-u)g(u)du \ \text{ for } z \in [e_m(z_{\text{AI}}; z_{\text{AI}}), z_{\text{AI}}]$$

In Section 1 of the online Appendix, we provide a more detailed characterization of the partition $(\mathcal{R}_j)_{j=1,\ldots,4}$, though the reader should envision something akin to Figure 2(b). We can now state the main result of this section:

Proposition 2. *In the presence of AI, there is a unique competitive equilibrium. It is given as follows:*

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

$$\bar{z}_1^* = f_1^*(z_{AI}; \underline{z}_1^*) \quad \textit{and} \quad n(z_{AI})(f_1^*(z_{AI}; \underline{z}_1^*) - z_{AI}) = \underline{z}_1^* + \textstyle\int_{\underline{z}_1^*}^{f_1^*(z_{AI}; \underline{z}_1^*)} n\big(e_1^*(z; \underline{z}_1^*)\big)dz$$
$$\textit{where } G(f_1^*(z; \underline{z}_1^*)) = G(\underline{z}_1^*) + \textstyle\int_0^z h(1-u)dG(u) \textit{ for } z \in [0, z_{AI}]$$

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

$$\bar{z}_2^* = f_2^*(z_{AI}; \underline{z}_2^*) \quad \textit{and} \quad n(z_{AI})(f_2^*(z_{AI}; \underline{z}_2^*) - z_{AI}) = z_{AI} + \int_{z_{AI}}^{f_2^*(z_{AI}; \underline{z}_2^*)} n\big(e_2^*(z; \underline{z}_2^*)\big) dz$$
$$\textit{where } G(f_2^*(z; \underline{z}_2^*)) = G(z_{AI}) + \int_{\underline{z}_2^*}^{z} h(1-u) dG(u) \textit{ for } z \in [\underline{z}_2^*, z_{AI}]$$

- *If $(h, z_{AI}) \in \mathcal{R}_3$, then the equilibrium is Type 3. The cutoffs $\underline{z}_3^*$ and $\bar{z}_3^*$ then satisfy:*

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

- *If $(h, z_{AI}) \in \mathcal{R}_4$, then the equilibrium is Type 4. The cutoffs are then $\underline{z}_4^* = \underline{z}(h)$ and $\bar{z}_4^* = \bar{z}(h)$, where $\underline{z}(h)$ and $\bar{z}(h)$ are the equilibrium cutoffs of the pre-AI equilibrium when $h > h_0$. This implies that AI has no effect on human labor outcomes, as it does not directly interact with humans in the workplace.*[20]

However, irrespective of the region, the wage function $w(z)$ is continuous, strictly increasing, and (weakly) convex, and satisfies:

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

where $f_j^*(z; \cdot)$ and $e_j^*(z; \cdot)$ are the matching and employee functions in a Type j equilibrium.

Proof. See Appendix B. □

A notable implication of Proposition 2 is that AI does not completely destroy worker positions for humans unless $z_{\text{AI}} = 0$. This result arises despite the fact that processing capacity is abundant and AI is more knowledgeable than a fraction of the human population.[21]

Intuitively, even though processing capacity 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 since they are the least likely to succeed on their own.

[20] Note that AI does increase total output. However, it does not increase total *labor* income, nor does it affect wages and the organization of work.

[21] In fact, it is not difficult to prove that the mass of human workers in the economy (i.e., the measure of W^*) always increases with z_{AI}, unless $(h, z_{\text{AI}}) \in \mathcal{Z}_i$. That is, the higher the mass of humans that are less knowledgeable than AI, the larger the pool of workers doing production work in two-layer organizations post-AI.

Let us construct one of the equilibria defined above to provide further insights regarding the AI equilibrium of Proposition 2. Consider, for example, the Type 2 equilibrium of region $\mathcal{R}_2$. By Corollary 2, such equilibrium must lead to the following partition of the human population:

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

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

Things are more interesting in the case of 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 $G(f_2^*(z)) = G(z_{\text{AI}}) + \int_{\underline{z}_2^*}^{z} h(1-u)dG(u)$ for all $z \in [\underline{z}_2^*, z_{\text{AI}}]$. 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 following differential equation (which comes from differentiating both sides of the resource constraint):

$$f_2^{*\prime}(z) = \frac{h(1-z)g(z)}{g(f_2^*(z))} \text{ for } z \in [\underline{z}_2^*, z_{\text{AI}}], \text{ plus the border condition } f_2^*(z_{\text{AI}}) = \bar{z}_2^*$$

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, we can derive the equilibrium wages for the rest of the population. To do this, 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:

$$\max_{m \in M_p^*} \Pi_2^{n\text{A}}(z, m) = n(z)[m - w(z)] - w(m)$$

The corresponding first-order condition evaluated at $m = f_2^*(z; \underline{z}_2^*)$ implies that $w^{*\prime}(f_2^*(z; \underline{z}_2^*)) = n(z)$, or, equivalently, $w^{*\prime}(z) = n(e_2^*(z; \underline{z}_2^*))$ for any $z \in M_p^*$. Thus,

$$w^*(z) = C + \int_{z_{\text{AI}}}^{z} n(e_2^*(u; \underline{z}_2^*))du, \text{ for any } z \in M_p^*$$

The wages of the workers of such firms then come from the zero profit condition of nA firms:

$$w^*(z) = f_2^*(z; \underline{z}_2^*) - \frac{w^*(f_2^*(z; \underline{z}_2^*))}{n(z)}, \text{ for any } z \in W_p^*$$

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 processing capacity, so $C^* = z_{\text{AI}}$.

On the other hand, 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:

$$n(z_{\text{AI}})(f_2^*(z_{\text{AI}}; \underline{z}_2^*) - z_{\text{AI}}) = z_{\text{AI}} + \int_{z_{\text{AI}}}^{f_2^*(z_{\text{AI}}; \underline{z}_2^*)} n(e_2^*(z; \underline{z}_2^*))dz \tag{4}$$

which is the equilibrium condition in the statement of Proposition 2. 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 $(h, z_{\text{AI}}) \in \mathcal{R}_2$. 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 (4) then implies that:

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

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

So the wage function 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.

4 The Effects of AI

In this section, we study the effects of AI by comparing the pre-AI and post-AI equilibrium. As mentioned in Proposition 2, when $(h, z_{\text{AI}}) \in \mathcal{R}_4 = \mathcal{Z}_i$, AI has no effect on human labor outcomes. Hence, from hereon, we focus on the (interesting) case where (h, z_{AI}) belongs either to $\mathcal{Z}_w$, $\mathcal{Z}_m$, or $\mathcal{Z}_k$.

4.1 Occupational Displacement

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

Proposition 3.

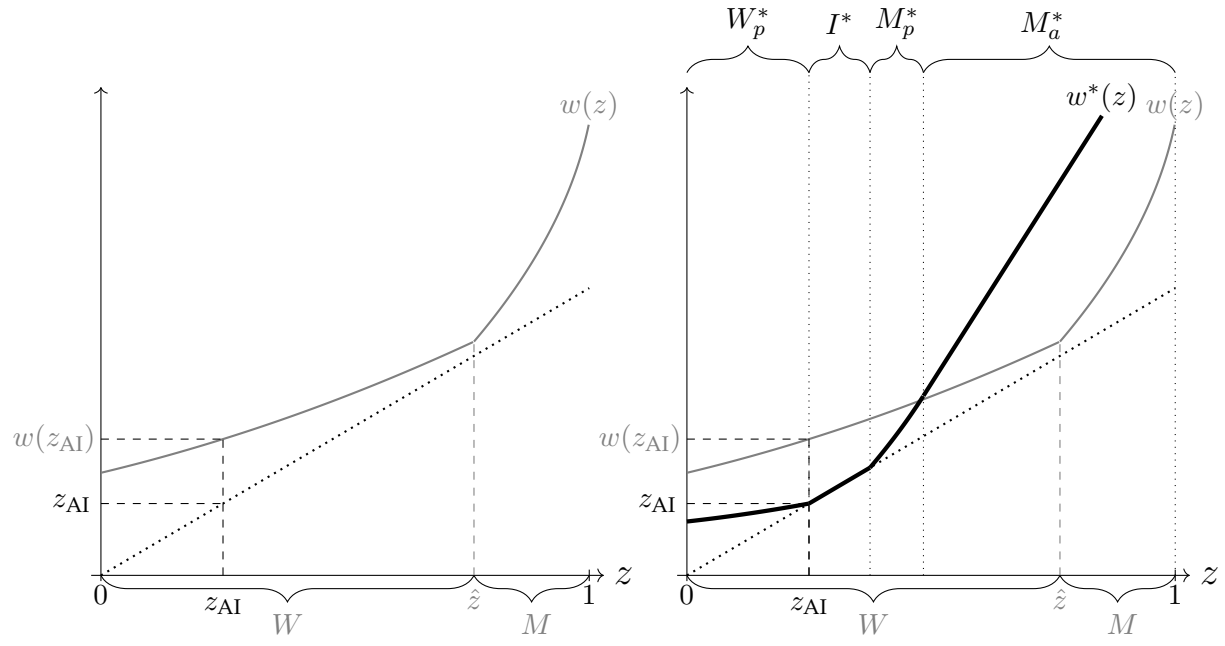
- *If $(h, z_{AI}) \in \mathcal{Z}_w$, then AI displaces humans from routine to managerial work, i.e., $W^* \subset W$ and $M^* \supset M$.*
- *If $(h, z_{AI}) \in \mathcal{Z}_m$, then AI displaces humans from managerial to routine work, i.e., $W^* \supset W$ and $M^* \subset M$.*
- *If $(h, z_{AI}) \in \mathcal{Z}_k$, 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 destruction of routine 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 into managerial positions. In contrast, the displacement goes in the opposite direction when AI has the knowledge of a pre-AI manager.

To understand why this is so, let us consider the effects of introducing AI with the same knowledge as a pre-AI worker when starting from the pre-AI scenario. Figure 3 panel (a), for instance, depicts a pre-AI situation when $h \leq h_0$ and $z_{\text{AI}} < \hat{z}$. Notice that the pre-AI wage of a worker with knowledge z_{AI} is strictly higher than z_{AI}. This is because workers always earn strictly more than their expected output as independent producers.

Starting from the situation described in 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



(a) Pre-AI Equilibrium when $h \leq h_0$

(b) The Effects of AI when $h \leq h_0$ and $z_{\text{AI}} < \hat{z}$

Figure 3: Human Displacement from Routine to Managerial Jobs

Notes. Distribution of knowledge: $G(z) = z$. Parameter values: $h = 1/2$ (so $\hat{z} = 3 - \sqrt{5}$) and $z_{\text{AI}} = 0.25$. See Section 5 of the online Appendix for a detailed characterization of the pre-AI and post-AI equilibrium when human knowledge is uniformly distributed.

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 both to the left and to the right of AI's knowledge level, as AI provides firms with cheaper labor 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, increasing the mass of humans specializing in problem-solving (i.e., the cutoff of the least knowledgeable manager starts decreasing beginning from $\hat{z}$).[22]

Consider, in contrast, the effects of an AI equipped with the knowledge of a pre-AI manager. We analyze this situation in Figure 4, which depicts the same pre-AI equilibrium as in Figure 3 but considers $z_{\text{AI}} > \hat{z}$.

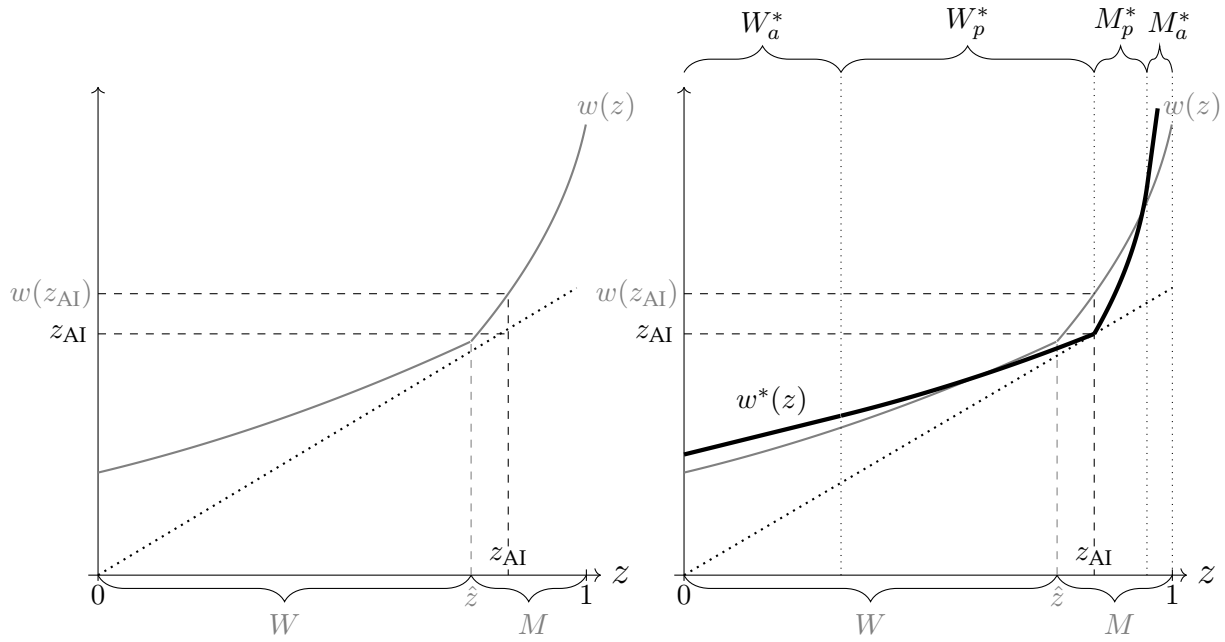
In this case, the introduction of AI originally forces those humans with knowledge $z \in [\hat{z}, z_{\text{AI}}]$ to become independent producers and, at least initially, lowers the wages of all the managers to the right of z_{AI}. The decline in managers' wages, coupled with the newly created class of independent producers and the relative abundance of processing capacity that can be used in managerial positions, increase the attractiveness of creating two-layer firms. The result is a surge in routine job opportunities within two-layer firms, leading to an increase in the fraction of humans executing routine work.[23]

Lastly, consider the knife-edge case where $z_{\text{AI}} = \hat{z}(h)$ when $h < h_0$. The introduction of AI reduces wages around z_{AI}, which increases the incentives to create two-layer firms. However, because AI has the knowledge of both a pre-AI worker and a pre-AI manager, and AI is always the best worker and the worst manager in the economy, there cannot be any human displacement between routine and managerial positions.

This implies that the managerial positions created by the newly established two-layer firms must

[22]This intuition also explains why when $(h, z_{\text{AI}}) \in \mathcal{Z}_w$ the equilibrium can only be Type 1 or Type 2. Indeed, 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 solely used as a worker (in addition to being used for independent production). Thus, the economy converges to an AI equilibrium Type 1 (this is the AI equilibrium depicted in Figure 3(b), where $W_a^* = \emptyset$). Otherwise, the class of 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. This leads to a Type 2 equilibrium, where AI is used both as a worker and as a manager.

[23]By similar reasoning as the one found in footnote 22, this intuition explains why when $(h, z_{\text{AI}}) \in \mathcal{Z}_w$ the equilibrium can only be Type 2 or Type 3. Note also that even though the introduction of AI originally decreases the wages of all human managers, once equilibrium effects are accounted for, the most knowledgeable human managers might still end up strictly better off (as can be seen from the Figure 4(b)). This is because they are now supervising production work executed by AI, which is cheaper and more knowledgeable than the best human workers pre-AI.



(a) Pre-AI Equilibrium when $h \leq h_0$

(b) The Effects of AI when $h \leq h_0$ and $z_{\text{AI}} > \hat{z}$

Figure 4: Human Displacement from Managerial to Routine Work

Notes. Distribution of knowledge: $G(z) = z$. Parameter values: $h = 1/2$ (so $\hat{z} = 3 - \sqrt{5}$) and $z_{\text{AI}} = 0.84$. See Section 5 of the online Appendix for a detailed characterization of the pre- and post-AI equilibrium when human knowledge is uniformly distributed.

be completely offset by the routine work positions created by the same firms. Thus, AI only generates within-worker and within-manager displacement in this case: Due to positive assortative matching, the least knowledgeable humans move into tA firms, while the most knowledgeable humans move into bA firms.

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. Recall from Corollary 1 that, in a pre-AI equilibrium without independent production, the knowledge cutoff to become a manager decreases in h. Moreover, as shown in Appendix A.3, the same 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 work into managerial roles 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.[24]

[24] An implicit assumption underlying this argument is that there are no cross-national teams. Understanding the impact

4.2 Team Formation and the Size and Productivity Distribution of Firms

We take *firm size* to be its output and *firm productivity* to be its output divided by the units of time and/or processing capacity it uses.[25] This implies that (i) the size and productivity of a single-layer firm are equal to the knowledge of its associated producer, and (ii) the size and productivity of a two-layer firm whose workers have knowledge z and whose manager has knowledge m are $mn(z)$ and m, respectively.

The following lemma—which follows because the equilibrium exhibits occupational stratification and positive assortative matching—describes the connection between the number of layers of a firm, its size, and its productivity:

Lemma 2. *Take two arbitrary firms f and f'. Irrespective of the presence of AI, in equilibrium:*

- *If f has two layers and f' has one layer, then f is strictly larger and strictly more productive than f'.*
- *Firm f is strictly larger than f' if and only if f is strictly more productive than f'.*

Proof. See Appendix C. □

Lemma 2 has two important implications. First, the largest firm is always the most productive, while the smallest firm is the least productive. Second, the biggest firm is always a two-layer firm, which combined with positive assortative matching, implies that its productivity is 1 (as it is managed by the most knowledgeable human).

We now turn to analyzing the implication of AI in the size and productivity distribution of firms. We begin by describing its effect on two-layer firms:

Proposition 4. *Regarding two-layer firms:*

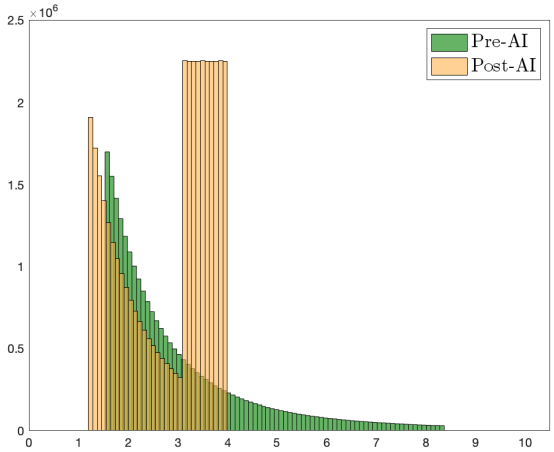
- *If $(h, z_{AI}) \in \mathcal{Z}_w$, then the size of the biggest firm and the size and productivity of the smallest/least productive firm decreases with AI.*
- *If $(h, z_{AI}) \in \mathcal{Z}_m$, then the size of the biggest firm and the size and productivity of the smallest/least productive firm increases with AI.*
- *If $(h, z_{AI}) \in \mathcal{Z}_k$, then the size of the biggest firm and the size and productivity of the smallest/least productive firm do not change with AI.*

Proof. See Appendix C. □

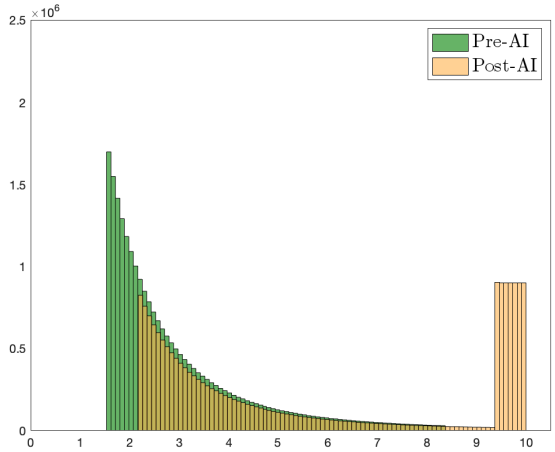
Proposition 4 is illustrated in Figure 5, and its intuition is closely connected to the displacement result described in Proposition 3. In the case of two-layer organizations, the knowledge of the best

of AI on offshoring is an interesting and relevant area for future research.

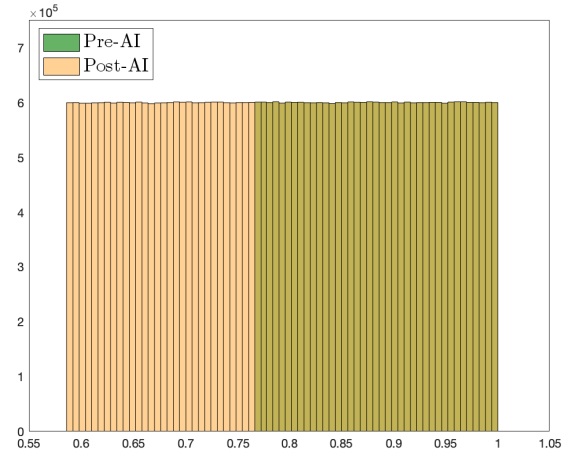
[25] An alternative would be to define firm size as the number of humans engaged in production. We decided against this definition because, in our setting, it would lead to some firms having size zero.



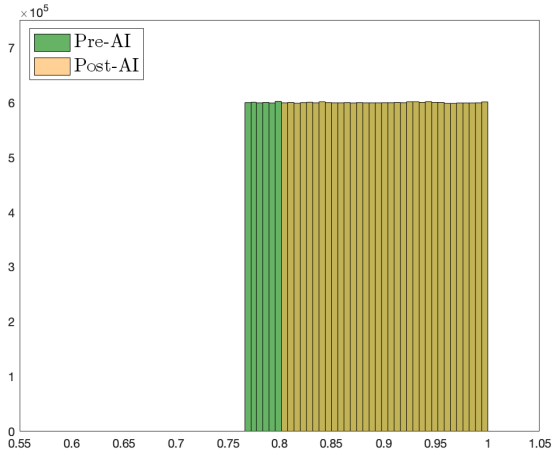
(a) Firm Size - $(h, z_{\text{AI}}) \in \mathcal{Z}_w$



(b) Firm Size - $(h, z_{\text{AI}}) \in \mathcal{Z}_m$



(c) Firm Productivity - $(h, z_{\text{AI}}) \in \mathcal{Z}_w$



(d) Firm Productivity - $(h, z_{\text{AI}}) \in \mathcal{Z}_m$

Figure 5: The Effects of AI on the Size and Productivity Distribution of Two-Layer Firms

Notes. Distribution of knowledge: $G(z) = z$. Histograms based on $N = 10.000.000$ random draws from this distribution. Parameter values: both panels have $h = 1/2$ (so $\hat{z} = 3 - \sqrt{5}$) and $\mu = 12.5 \times N$. Moreover, panels (a) and (c) assume $z_{\text{AI}} = 0.5$, while panels (b) and (d) assume $z_{\text{AI}} = 0.84$. See Section 5 of the online Appendix for a detailed characterization of the pre- and post-AI equilibrium when human knowledge is uniformly distributed.

worker determines the maximum firm size (as the best manager always has knowledge 1), while the knowledge of the worst manager determines the minimum firm size and the productivity of the least productive firm (as the worst worker always has knowledge 0). Hence, when AI has the knowledge of a pre-AI worker, the size of the biggest firm and the size and productivity of the smallest/least productive firm decreases with AI, as the best pre-AI workers become managers post-AI. A similar intuition explains the other two cases.

The effects of AI on the size and productivity distributions of *all firms* are similar:

Proposition 5. *The results regarding the productivity and size distribution of all firms mimic those of two-layer firms described in Proposition 4, except that:*

- *If $(h, z_{AI}) \in \mathcal{Z}_k$, the size of the smallest/least productive firm decreases with AI.*
- *If $(h, z_{AI}) \in \mathcal{Z}_m$ and $h \leq h_0$, the size of the smallest/least productive firm can increase or decrease with AI.*

Proof. See Appendix C. □

To understand these results, consider productivity first. In contrast to the case of two-layer firms, the productivity of the smallest/least productive firm is now determined by the knowledge of the best worker (rather than by the knowledge of the worst manager). Nevertheless, even though these knowledge levels need not coincide, they move in the same direction with the introduction of AI.

Second, consider size. Since the biggest firms are always two-layer firms, the results concerning the biggest firms remain unchanged. The difference arises in the results regarding the smallest/least productive firm: Intuitively, even in cases in which the size of the smallest two-layer firm increases with AI, the size of the smallest firm might decrease with AI because AI also leads to the creation of new, potentially smaller one-layer firms.

Propositions 4 and 5 are driven by the displacement of humans across occupations. We now analyze the effects of AI on the productivity of workers and the size of firms being managed by those *who have not been displaced*. In line with our definition of firm productivity, we define worker productivity as her expected output. For the proposition that follows, recall that $e(z)$ denotes the employee function of the pre-AI equilibrium:

Proposition 6.

- *If $(h, z_{AI}) \in \mathcal{Z}_w$, then:*
  - *Each $z \in W^* \subset W$ is strictly less productive post-AI than pre-AI.*
  - *Consider a $z \in M \subset M^*$. If $e(z) < z_{AI}$, then z manages a strictly larger firm post-AI than pre-AI. If $e(z) > z_{AI}$, then z manages a strictly smaller firm post-AI than pre-AI. Finally, if $e(z) = z_{AI}$, then z manages a firm of equal size post-AI and pre-AI.*
- *If $(h, z_{AI}) \in \mathcal{Z}_m$, then:*
  - *Consider a $z \in W \subset W^*$. If $z < e(z_{AI})$, then z is strictly more productive post-AI than pre-AI. If $z > e(z_{AI})$, then z is strictly less productive post-AI than pre-AI. Finally, if $z = e(z_{AI})$, then z is equally productive post-AI and pre-AI.*
  - *Each $z \in M^* \subset M$ manages strictly larger firm post-AI than pre-AI.*
- *If $(h, z_{AI}) \in \mathcal{Z}_k$, then:*
  - *Each $z \in W^* = W$ is less productive post-AI than pre-AI (strictly so for all $z \neq 0$).*

– *Each $z \in M^* = M$ manages a larger firm post-AI than pre-AI (strictly so for all $z \neq 1$).*

Proof. See Appendix C. □

For intuition, note that a worker's productivity is the knowledge of her manager, so a worker's productivity increases if and only if her managerial match improves. Similarly, the size of the firm of a given manager increases if and only the knowledge of her workers increases.

With this in mind, 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 created class of less knowledgeable managers.

Similarly, when AI has the knowledge of a pre-AI manager, the match of the worst workers improves since 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 (explaining why all managers are in charge of larger firms), while the second result implies that the pool of available managers for the remaining non-displaced workers worsens with AI (explaining why all workers become less productive).

4.3 Labor Income

In this section, we analyze the consequences of AI on labor income. The analysis 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 reduces the latter's wages and affects how each firm's output is divided between workers and managers.

We begin our analysis by showing that *if* a subset of humans with knowledge below z_{AI} win after the introduction of AI, then they must be the least knowledgeable humans in the population. Similarly, *if* a subset of humans above z_{AI} are better off after the advent of AI, then they must be the most knowledgeable humans:

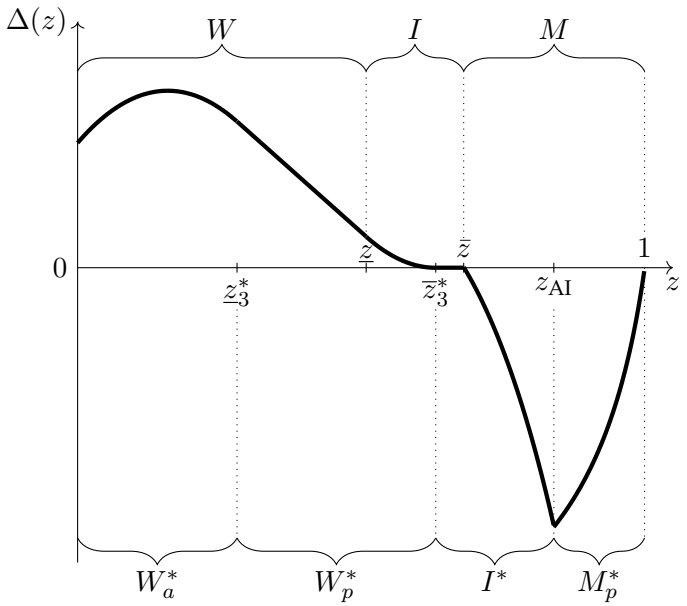


Figure 6: An Illustration of $\Delta(z)$ Function when $h > h_0$ and $(h, z_{\text{AI}}) \in \mathcal{Z}_m$

Notes. Distribution of knowledge: $G(z) = z$. This implies that $(h_0, z_0) = (3/4, 2/3)$. Parameter values: $h = z_{\text{AI}} = 0.84$. Thus, $\underline{z}_3^* = 0.28$, $\underline{z} = 0.51$, $\bar{z}_3^* = 0.63$, $\bar{z} = 0.68$. Moreover, $\Delta(0) = 0.025$, $\Delta(1) = 0$, $\max_z \Delta(z) = 0.036$, and $\min_z \Delta(z) = -0.053$. See Section 5 of the online Appendix for a detailed characterization of the pre- and post-AI equilibrium when human knowledge is uniformly distributed.

Lemma 3. *Define $\Delta(z) \equiv w^*(z) - w(z)$. Then $\Delta(z_{AI}) < 0$ and:*

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

Proof. See Section 3 of the online Appendix. □

Lemma 3 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, as stated in the lemma, the $\Delta(z)$ function can only cross zero from above in $z \in [0, z_{\text{AI}}]$, while it can only cross zero from below in $z \in [z_{\text{AI}}, 1]$.

Intuitively, the lemma is explained 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) there are winners below or above z_{AI}:

Proposition 7.

(i) There exists $z < z_{AI}$ such that $\Delta(z) > 0$ if and only if $z_{AI} > \mathrm{Z}(h)$, where:

$$\begin{cases} \mathrm{Z}(h) \in (0, \hat{z}(h)) & \text{if } h \leq h_0 \\ \mathrm{Z}(h) = \bar{z}(h) & \text{if } h > h_0 \end{cases}$$

(ii) $\Delta(z) < 0$ for all $z \in (z_{AI}, 1)$ if and only if $h > h_0$ and $(h, z_{AI}) \in \mathcal{Z}_m$.

Proof. See Appendix C. □

The intuition for the first part of the proposition is as follows. On the one hand, the introduction of AI reduces the wages around z_{AI}, as mentioned above. This leads to a reduction in the wages of the best workers and worst managers. Since these are the managers of the worst workers, this reduction in wages allows the worst workers to appropriate a larger share of the firm's output. On the other hand, when z_{AI} is sufficiently low, AI also worsens the matches of workers, reducing their firms' output. However, the quality of the worst managers increases with z_{AI}, so when z_{AI} is large enough, the worst workers have access to relatively good managers whose wages have decreased with AI's introduction. The result is an increase in these workers' wages increase post-AI.

The second part follows a similar logic. On the one hand, AI lowers the wage of those close to z_{AI}. On the other hand, those agents whose knowledge is immediately above z_{AI} lead larger teams post-AI than pre-AI, which implies that the slope of the wage function above z_{AI} increases with AI immediately above z_{AI} (as the marginal return to managerial knowledge is equal to team size). Hence, if this slope effect is large enough (which is always the case unless $h > h_0$), everyone whose knowledge is sufficiently higher than z_{AI} wins from the introduction of AI.

So far, we have analyzed the distributional consequences of AI based on the location of humans in the knowledge distribution relative to AI, i.e., whether those above or below z_{AI} potentially win. An alternative approach is to characterize whether increases or decreases in wages are related to whether a specific individual was displaced by AI to a different occupation or not. In particular, a reasonable conjecture is that all displaced humans are necessarily worse off. If that were the case, taxing the winners to compensate losers would be easier in practice, as displacement is more likely to be observable for an authority than an individual's knowledge.

Unfortunately, however, the conjecture turns out to be incorrect. For a counterexample, suppose that $h \leq h_0$ and consider the outcome of an agent with knowledge $\hat{z} + \delta$ after the introduction of AI with knowledge $z_{\mathrm{AI}} = 1 - \epsilon$, where δ and ϵ are strictly positive but small. Pre-AI, this human is a manager and her wage is $w(\hat{z} + \delta) \approx 1 - h(1 - \hat{z})w(1)$. Post-AI, the equilibrium is Type 2 and $\hat{z} + \delta \in W_a^*$, so her wage is $w^*(\hat{z} + \delta) \approx (1 - \epsilon)(1 - h(1 - \hat{z}))$. Since $w(1) > 1$, then $w^*(\hat{z} + \delta) > w(\hat{z} + \delta)$ for ϵ and δ small enough, even though this agent was displaced from a managerial position into a routine job.[26]

[26] Figure 3(b), in turn, provides an example where some pre-AI workers are strictly better off after being displaced into managerial positions due to AI.

5 Extensions and Discussion

In this section, we provide two extensions: First, we consider the case where processing capacity is small. Second, we discuss the other extreme in which processing capacity is large not only relative to human time but also relative to its potential applications. We also discuss the effects of AI on the labor share, as well as real-world examples in which AI acts as a "worker," a "manager," or an "independent producer" in the sense of our framework.

5.1 The Role of Abundance of Processing Capacity Relative to Human Time

Until now, we have considered the case in which processing capacity is large relative to human time. To understand the role of this assumption, in Section 6 of the online Appendix, we characterize the equilibrium in the opposite extreme in which processing capacity 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 work into problem-solving. In contrast, humans are displaced in the opposite direction when AI has the knowledge of a pre-AI manager. This also implies that our results regarding the size and productivity distribution of firms 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 $(h, z_{\text{AI}}) \in \mathcal{Z}_w$, or that all managers supervise larger firms when $(h, z_{\text{AI}}) \in \mathcal{Z}_m$. 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 $(h, z_{\text{AI}}) \in \mathcal{Z}_w$. Similarly, there is a set of managers—those with knowledge below z_{AI}—who manage smaller firms when $(h, z_{\text{AI}}) \in \mathcal{Z}_m$. Intuitively, these differences arise because, when processing capacity is abundant, there are no workers with knowledge above z_{AI} nor managers with knowledge below z_{AI}.

With regards to labor income, when processing capacity 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).[27] This intuition highlights the key role that abundant processing capacity plays in our analysis of the effects of AI on labor income.

[27] 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).

5.2 What if Processing Capacity Exceeds its Potential Uses?

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

In the Online Appendix, we relax the assumption that processing capacity is scarce relative to its uses. 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 processing capacity is abundant not only relative to human time but also relative to Q.[28]

In this case, and in contrast to our baseline setting, the equilibrium price of processing capacity 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 most knowledgeable humans 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 size and productivity distribution of firms continue to hold with minor changes. In particular, the size of the smallest two-layer firm can increase or decrease with AI when $(h, z_{\text{AI}}) \in \mathcal{Z}_w$ (while in the baseline, it always decreases), and necessarily increases when $(h, z_{\text{AI}}) \in \mathcal{Z}_k$ (in the baseline, its size remains unchanged). Intuitively, 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 the least knowledgeable worker, so AI increases the knowledge of the workers of the smallest/least productive firm.[29]

Regarding the quality of matches, the results on 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.

[28] 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.

[29] As a result, in this extension, the smallest firm has a worse manager but better workers post-AI than pre-AI when $(h, z_{\text{AI}}) \in \mathcal{Z}_w$ (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 $(h, z_{\text{AI}}) \in \mathcal{Z}_k$ (while in the baseline it has the same workers and manager).

5.3 The Effects of AI on the Labor Share

Finally, we discuss the effects of AI on the labor share. Pre-AI, the labor share is equal to one (since labor is the sole input in production). Interestingly, whether AI lowers the labor share or not depends on whether the set of potential uses for AI is larger or smaller than the amount of processing capacity in the economy.

In the baseline setting—where processing capacity has unlimited applications—the labor share decreases with AI (as long as $z_{\text{AI}} > 0$). This is because the equilibrium price of processing capacity is strictly positive. Moreover, the labor share converges to zero as $\mu \to \infty$ since, when processing capacity is abundant relative to human time, any additional unit of processing capacity 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 processing capacity 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 non-monotone in μ: it decreases with μ while processing capacity is abundant relative to time but not to problems, and then increases with μ as processing capacity catches up with the number of problems. We leave a more detailed study of this race between processing capacity and its potential uses for future research.

5.4 Examples of AI Worker, Independent Producer, and Manager

To understand how AI can play the role of a "worker," consider legal services, which is a human-capital intensive industry. 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 (Garicano and Hubbard, 2007).

These are arguably the tasks that are 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:[30]

> 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

[30] "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)

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").[31]

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

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.[32] 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 4.3: 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.

[31] 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).

[32] 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.

Our framework provides novel insights into 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.

APPENDIX

A Proof Omitted from Section 2.3

In this Appendix, we provide the full characterization of the equilibrium without AI, including the proofs of Proposition 1 and Corollary 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 $G(f(z;\bar{z})) = G(\bar{z}) + \int_0^z h(1-u)dG(u)$, and the cutoffs $0 < \underline{z} < \bar{z} < 1$ satisfy:*

$$\frac{1}{h} - \bar{z} = \int_{\bar{z}}^1 n\big(e(u;\bar{z})\big)du \quad \text{and} \quad f(\underline{z};\bar{z}) = 1 \quad \text{where } e(z;\bar{z}) = f^{-1}(z;\bar{z}) \tag{5}$$

- *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 $G(f(z;\hat{z})) = G(\hat{z}) + \int_0^z h(1-u)dG(u)$ and the cutoff $\hat{z} \in (0,1)$ satisfies $f(0;\hat{z}) = \hat{z}$, $f(\hat{z};\hat{z}) = 1$, and:*

$$\begin{array}{l} \frac{1}{h} - \hat{z} > \int_{\hat{z}}^1 n(e(z;\hat{z}))dz \text{ for all } h < h_0 \\ \frac{1}{h} - \hat{z} = \int_{\hat{z}}^1 n(e(z;\hat{z}))dz \text{ when } h = h_0 \end{array} \tag{6}$$

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 $h \gtrless h_0$, 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 (5)). 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:

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

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 $h > h_0$. □

A.2 Proof of Corollary 1

Consider first $\hat{z}(h)$. The proof that $\hat{z}(h)$ is strictly decreasing in h is in Fuchs et al. (2015, Lemma 2). That $\hat{z}(h_0) \in (0,1)$ follows, in turn, because $h_0 \in (0,1)$ and $\hat{z}(h_0)$ is given by $\int_{\hat{z}(h_0)}^{1} dG(u) = \int_{0}^{\hat{z}(h_0)} h_0(1-u)dG(u)$. Finally, that $\hat{z}(\underline{h}) \to 0$ as $\underline{h} \to 0$ follows from the fact that $\hat{z}(\underline{h})$ satisfies $\int_{\hat{z}(\underline{h})}^{1} dG(u) = \int_{0}^{\hat{z}(\underline{h})} \underline{h}(1-u)dG(u)$ and that this last expression becomes $\int_{\hat{z}(\underline{h})}^{1} dG(u) = 0$ when $\underline{h} \to 0$.

Now consider $\underline{z}(h)$ and $\bar{z}(h)$. The proof that $\underline{z}(h)$ is strictly decreasing in h is in Fuchs et al. (2015, Lemma 2). Thus, we focus on showing the other properties of these cutoffs. That $\bar{z}(h_0) = z_0$ follows because $\bar{z}(h)$, and $\hat{z}(h)$ are uniquely determined and $\hat{z}(h_0)$ satisfies the exact same condition as $\bar{z}(h)$

when $h \to h_0$. Moreover, that $\underline{z}(h_0) = \hat{z}(h_0)$ follows because $\underline{z}(h)$ is uniquely determined and satisfies $f(\underline{z}(h); \bar{z}(h), h) = 1$. Consequently:

$$f(\underline{z}(h_0); \bar{z}(h_0), h_0) = f(\underline{z}(h_0); \hat{z}(h_0), h_0) = 1 \Longrightarrow \underline{z}(h_0) = \hat{z}(h_0)$$

where the last implication follows because $f(\hat{z}(h_0); \hat{z}(h_0), h_0) = 1$ by Lemma A.1. Finally, that $\bar{z}(1) = 1$ and $\underline{z}(1) = 0$ follow because $\bar{z} = 1$ is a solution (and, therefore, the unique solution) to (5) when $h = 1$, and because the only way to satisfy $f(\underline{z}(1); \bar{z}(1), 1) = 1$ given that $\bar{z}(1) = 1$ is for $\underline{z}(1) = 0$. □

A.3 The Effects of Changes in the Distribution of Knowledge

We finish this appendix with the following result, which describes how the identity of the best worker and that of the worst manager in the pre-AI equilibrium changes as the population becomes more knowledgeable:

Lemma A.3. *Let $\hat{z}^k(h)$ be the pre-AI cutoff in the equilibrium with no independent production under the knowledge distribution $G^k(z)$ (so $\hat{z}^k(h)$ is defined for $h \in [\underline{h}, h_0^k)$). If $G^A \succcurlyeq_{\text{FOSD}} G^B$, then $\hat{z}^A(h) \geq \hat{z}^B(h)$ for any given $h \in [\underline{h}, h_0^A) \cap [\underline{h}, h_0^B)$.*

Proof. We want to show that if $G^A \succcurlyeq_{\text{FOSD}} G^B$, then $\hat{z}^A(h) \geq \hat{z}^B(h)$ for any given $h \in [\underline{h}, h_0^A) \cap [\underline{h}, h_0^B)$. To simplify notation, we omit the dependence on h of the equilibrium variables. Recall that $\hat{z}^k$ is given by $f^k(\hat{z}^k; \hat{z}^k) = 1$, where $f^k(z; \hat{z}^k)$ is the equilibrium matching function in the pre-AI equilibrium. This implies that $\hat{z}^k$ must satisfy $1 = G(\hat{z}^k) + \int_0^{\hat{z}^k} h(1-u) dG^k(u)$. Integrating by parts the integral on the right-hand side of this last expression and rearranging terms yields that

$$1 = \underbrace{G^k(\hat{z}^k)[1 + h(1 - \hat{z}^k)] + h \int_0^{\hat{z}^k} G^k(z) dz}_{\equiv \alpha^k(\hat{z}^k)}$$

Note then that $\alpha^k(x)$ is strictly increasing in x, so $\hat{z}^k$ is the unique solution to $\alpha^k(x) = 1$. We claim that $G^A \succcurlyeq_{\text{FOSD}} G^B$ implies that $\alpha^A(x) \leq \alpha^B(x)$ for any given x, which immediately implies that $\hat{z}^A \geq \hat{z}^B$. Indeed, note that:

$$\alpha^B(x) - \alpha^A(x) = [1 + h(1-x)][G^B(x) - G^A(x)] + h \int_0^x [G^B(u) - G^A(u)] du \geq 0$$

as $G^B(z) \geq G^A(z)$ given that $G^A \succcurlyeq_{\text{FOSD}} G^B$. □

B Proofs Omitted from Section 3

B.1 Proof of Lemma 1

• *The Rental Rate of Processing Capacity is equal to AI's Knowledge.*— This follows because some processing capacity must necessarily be used for independent production, and 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).

• *AI is the Best Worker and/or the Worst Manager.*— That AI is the best worker and/or the worst manager follows from occupational stratification and the fact that some processing capacity 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 AI is the best worker and/or the worst manager). □

B.2 Proof of Proposition 2

To prove that the proposed candidate in each of the four regions is an equilibrium, one can follow the same reasoning as the one developed in the main text for the Type 2 equilibrium of region $\mathcal{R}_2$. The same applies to the continuity of the wage function. The proof that the wage function is strictly increasing and weakly convex, on the other hand, follows a similar logic as the proof of Corollary 1 (which shows that the pre-AI wage function satisfies these two properties). Hence, here, we focus on proving the uniqueness of the equilibrium.

Lemma B.1. *If $(h, z_{AI}) \in \mathcal{R}_1$, then the unique equilibrium is the Type 1 equilibrium described in the Proposition 2.*

Proof. We will show that if $(h, z_{\text{AI}}) \in \mathcal{R}_1 = \mathcal{Z}_w \cap \{(h, z_{\text{AI}}) : \Gamma_w(h, z_{\text{AI}}) < 0\}$, there cannot be any other type of equilibrium. Suppose first by contradiction that there is an equilibrium Type 3. By Corollary 2, we know that such equilibrium must lead to the following partition of the human population:

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

Now suppose that $h \leq h_0$ (the case $h > h_0$ follows the exact same logic). As we show in the section 1 of online Appendix (specifically in Corollary 1.1), if $(h, z_{\text{AI}}) \in \mathcal{R}_1$, then $z_{\text{AI}} < \hat{z}$, where $\hat{z}$ satisfies $\int_{\hat{z}}^1 dG(u) = \int_0^{\hat{z}} h(1-u)dG(u)$ (by Lemma A.1). 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 says that $\int_{\underline{z}_3^*}^{\bar{z}_3^*} h(1-u)dG(u) = \int_{z_{\text{AI}}}^1 dG(u)$. Hence a Type 3 configuration cannot arise in this region.

Now consider a Type 2 configuration. As explained in the main text, for this to be an equilibrium, there must exist a cutoff $\underline{z}_2^* \in [0, z_{\text{AI}}]$ satisfying the following condition:

$$\overbrace{n(z_{\text{AI}})(f_2^*(z_{\text{AI}};\underline{z}_2^*) - z_{\text{AI}}) - z_{\text{AI}} - \int_{z_{\text{AI}}}^{f_2^*(z_{\text{AI}};\underline{z}_2^*)} n(e_2^*(z;\underline{z}_2^*))dz}^{\equiv \Lambda(\underline{z}_2^*)} = 0$$
$$\text{where } G(f_2^*(z;\underline{z}_2^*)) = G(z_{\text{AI}}) + \textstyle\int_{\underline{z}_2^*}^{z} h(1-u)dG(u) \text{ for } z \in [\underline{z}_2^*, z_{\text{AI}}]$$

The derivative of $\Lambda(\underline{z}_2)$ with respect to $\underline{z}_2$ is:

$$\Lambda'(\underline{z}_2) = -h \textstyle\int_{z_{\text{AI}}}^{f_2^*(z_{\text{AI}};\underline{z}_2)} n(e_2^*(z;\underline{z}_2))^2 \frac{\partial}{\partial \underline{z}_2} e_2^*(z;\underline{z}_2)dz = -h^2 \int_{z_{\text{AI}}}^{f_2^*(z_{\text{AI}};\underline{z}_2)} \frac{(1-z)g(z)}{g(e_2^*(z;\underline{z}_2))} n(e_2^*(z;\underline{z}_2))^3 dz < 0$$

Hence, there exists at most one $\underline{z}_2^*$ that satisfies $\Lambda(\underline{z}_2^*) = 0$, and for $\underline{z}_2^* \geq 0$, it must be that:

$$\Lambda(0) = n(z_{\text{AI}})(f_2^*(z_{\text{AI}};0) - z_{\text{AI}}) - z_{\text{AI}} - \textstyle\int_{z_{\text{AI}}}^{f_2^*(z_{\text{AI}};0)} n(e_2^*(z;0))dz > 0$$
$$\text{where } G(f_2^*(z;0)) = G(z_{\text{AI}}) + \textstyle\int_0^z h(1-u)dG(u) \text{ for } z \in [0, z_{\text{AI}}]$$

But this last condition implies that $f_2^*(z;0) = f_w(z; z_{\text{AI}})$, so $\Lambda(0) = \Gamma_w(h, z_{\text{AI}})$ (see equations (3) and (2) of the main text). Consequently, for $\underline{z}_2^* \geq 0$, we need that $\Gamma_w(h, z_{\text{AI}}) \geq 0$, which contradicts the fact that $(h, z_{\text{AI}}) \in \mathcal{R}_1 = \mathcal{Z}_w \cap \{(h, z_{\text{AI}}) : \Gamma_w(h, z_{\text{AI}}) < 0\}$.

Finally, consider a Type 4 configuration. By Corollary 2, 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 in the main text for a Type 2 equilibrium, for this to be an equilibrium, it must be that $\bar{z}_4^* = f_4^*(0; \underline{z}_4^*)$, where $f_4^*(0; \underline{z}_4^*)$ is given by:

$$f_4^*(\underline{z}_4^*; \underline{z}_4^*) = 1 \text{ and } G(f_4^*(z; \underline{z}_4^*)) = G(f_4^*(0; \underline{z}_4^*)) + \textstyle\int_{\underline{z}_2^*}^{z} h(1-u)dG(u) \text{ for } z \in [0, \underline{z}_4^*]$$

And there must exist a cutoff $\underline{z}_4^* < f_4^*(0; \underline{z}_4^*)$ such that:

$$1/h - f_4^*(0; \underline{z}_4^*) - \textstyle\int_{f_4^*(0;\underline{z}_4^*)}^{1} n(e_4^*(z; \underline{z}_4^*))dz = 0$$

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 $(h, z_{\text{AI}}) \in \mathcal{R}_1$. □

The proofs for the uniqueness of the Type 2 equilibrium in region $\mathcal{R}_2$ and the uniqueness of the Type 3 equilibrium in region $\mathcal{R}_3$ follow a similar logic. For this reason, we omit their proofs and simply state the result in the following lemma:

Lemma B.2.

- *If $(h, z_{AI}) \in \mathcal{R}_2$, then the unique equilibrium is the Type 2 equilibrium described in the Proposition 2.*
- *If $(h, z_{AI}) \in \mathcal{R}_3$, then the unique equilibrium is the Type 3 equilibrium described in the Proposition 2.*

We finish the proof by showing that the Type 4 equilibrium described in the proposition is the unique equilibrium in region $\mathcal{R}_4$:

Lemma B.3. *If $(h, z_{AI}) \in \mathcal{Z}_i = \mathcal{R}_4$, then the unique equilibrium is the Type 4 equilibrium described in the Proposition 2.*

Proof. Suppose $(h, z_{\text{AI}}) \in \mathcal{Z}_i = \mathcal{R}_4$ and consider a competitive equilibrium. By the First Welfare Theorem, such a competitive equilibrium must be efficient. Now, by Lemma A.1, we know that pre-AI, there is a unique efficient allocation. In it, everyone whose knowledge is between $\underline{z}$ and $\bar{z}$ is hired as an independent producer.

Consequently, it follows that if $z_{\text{AI}} \in (\underline{z}, \bar{z})$, then it is efficient to allocate the unit of processing capacity to independent production as well. Moreover, since the introduction of the first unit of processing capacity does not affect two-layer firms in any way, the same argument holds for the second unit, the third unit, and so on. Consequently, when $(h, z_{\text{AI}}) \in \mathcal{Z}_i = \mathcal{R}_4$, the unique, efficient allocation involves all processing capacity being used for independent production, so the unique competitive equilibrium is the Type 4 equilibrium described in the proposition. □

C Proofs Omitted from Section 4

C.1 Proof of Proposition 3

• $(h, z_{AI}) \in \mathcal{Z}_w$.— By Proposition 2 the AI equilibrium is either Type 1 or Type 2. Suppose first that it is Type 2. Then $\sup W^* = \inf M^* = z_{\text{AI}}$. Hence, $W^* \subset W$ and $M^* \supset M$, since $z_{\text{AI}} < \hat{z}$ and $z_{\text{AI}} < \underline{z}$ as $(h, z_{\text{AI}}) \in \mathcal{Z}_w$.

Now suppose the equilibrium is Type 1. Then $\sup W^* = z_{\text{AI}}$, which immediately implies that $W^* \subset W$ as $z_{\text{AI}} < \hat{z}$ and $z_{\text{AI}} < \underline{z}$. To show, on the other hand, that $M^* \supset M$, we will show that $\inf M^* = \underline{z}_1^*$ is strictly less than $\hat{z}$ when $h \leq h_0$, and strictly less than $\bar{z}$ when $h > h_0$.

Recall first that $\underline{z}_1^*$ satisfies:[33]

$$n(z_{\text{AI}})(f_1^*(z_{\text{AI}}; \underline{z}_1^*, z_{\text{AI}}) - z_{\text{AI}}) = \underline{z}_1^* + \int_{\underline{z}_1^*}^{f_1^*(z_{\text{AI}}; \underline{z}_1^*, z_{\text{AI}})} n\big(e_1^*(z; \underline{z}_1^*, z_{\text{AI}})\big) dz$$
$$\text{with } G(f_1^*(z; \underline{z}_1^*, z_{\text{AI}})) = G(\underline{z}_1^*) + \int_0^z h(1-u) dG(u) \text{ for } z \in [0, z_{\text{AI}}]$$

Define $\Gamma_1(x; z_{\text{AI}}) \equiv n(z_{\text{AI}})(f_1^*(z_{\text{AI}}; x, z_{\text{AI}}) - z_{\text{AI}}) - x - \int_x^{f_1^*(z_{\text{AI}}; x, z_{\text{AI}})} n\big(e_1^*(z; x, z_{\text{AI}})\big) dz$. It is not difficult to prove that $\Gamma_1(x; z_{\text{AI}})$ is strictly increasing in x and strictly decreasing in z_{AI}. Moreover, $\underline{z}_1^*$ is given by the unique solution to $\Gamma_1(x; z_{\text{AI}}) = 0$. We then claim that $\Gamma_1(\hat{z}; z_{\text{AI}}) > 0$ when $h \leq h_0$ and that

[33] To avoid any type of confusion, we are making explicit that $f_1^*(z; \underline{z}_1^*, z_{\text{AI}})$ and $e_1^*(z; \underline{z}_1^*, z_{\text{AI}})$ depend on both $\underline{z}_1^*$ and z_{AI}.

$\Gamma_1(\bar{z}; z_{\text{AI}}) > 0$ when $h > h_0$, which immediately implies that $\underline{z}_1^* < \hat{z}$ and that $\underline{z}_1^* < \bar{z}$. Indeed, note that $\Gamma_1(\hat{z}; z_{\text{AI}}) > \Gamma_1(\hat{z}; \hat{z}) = \frac{1}{h} - \hat{z} - \int_{\hat{z}}^1 n(e(z; \hat{z}, \hat{z}))dz$, where the first inequality follows because $z_{\text{AI}} < \hat{z}$ and the last equality because $e_1^*(z; \hat{z}, \hat{z}) = e(z; \hat{z})$ for all $z \in [\hat{z}, 1]$ (where $e(z; \hat{z})$ is the employee function of the pre-AI equilibrium). 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 $h \in [\underline{h}, h_0]$. Hence, $\Gamma_1(\hat{z}; z_{\text{AI}}) > 0$.

On the other hand, given that $z_{\text{AI}} < \underline{z}$ when $h \in (h_0, 1]$, then:

$$\Gamma_1(\bar{z}; z_{\text{AI}}) > \Gamma_1(\bar{z}; \underline{z}) = \tfrac{1}{h} - \bar{z} - \textstyle\int_{\bar{z}}^1 n\big(e(z; \bar{z})\big)dz = 0$$

where the second-to-last equality follows because $e_1^*(z; \bar{z}, \underline{z}) = e(z; \bar{z})$ for all $z \in [\bar{z}, 1]$, and the last equality due to the pre-AI equilibrium condition that determines $\bar{z}$. □

• $(h, z_{AI}) \in \mathcal{Z}_m$.— By Proposition 2, in this case, the AI equilibrium is either Type 1 or Type 2. Suppose first that it is Type 2. Then $\sup W^* = \inf M^* = z_{\text{AI}}$. Hence, $W^* \supset W$ and $M^* \subset M$ as $z_{\text{AI}} > \hat{z}$ and $z_{\text{AI}} > \bar{z}$ given that $(h, z_{\text{AI}}) \in \mathcal{Z}_m$.

Now suppose the equilibrium is Type 3. Then, $\inf M^* = z_{\text{AI}}$, so $M^* \subset M$ as $z_{\text{AI}} > \hat{z}$ and $z_{\text{AI}} > \bar{z}$. To prove that $W^* \supset W$, we will show that $\sup W^* = \bar{z}_3^* > \hat{z}$ (when $h \leq h_0$) and that $\sup W^* = \bar{z}_3^* > \underline{z}$ (when $h > h_0$). Recall that $\bar{z}_3^*$ is given by:

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

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:

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

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 to $\Gamma_3(x; z_{\text{AI}}) = 0$. 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.

On the other hand, to prove that $\bar{z}_3^* > \underline{z}$, we will show that $\Gamma_3(\underline{z}; z_{\text{AI}}) > 0$ when $h \in (h_0, 1]$. Indeed, because $z_{\text{AI}} > \bar{z}$ when $h \in (h_0, 1]$, then $\Gamma_3(\underline{z}; z_{\text{AI}}) > \Gamma_3(\underline{z}; \bar{z}) = \frac{1}{h} - \bar{z} - \int_{z_{\text{AI}}}^1 n\big(e(z; \bar{z})\big)dz = 0$, where the second-to-last equality follows because $\tilde{e}_3^*(z; \underline{z}, \bar{z}) = e(z; \bar{z})$, and the last equality due to the pre-AI equilibrium condition that determines $\bar{z}$. □

• $(h, z_{AI}) \in \mathcal{Z}_k$.— In Section 1 of the online Appendix—specifically, in Corollary 1.1—we show that when $h < h_0$ and $z_{\text{AI}} = \hat{z}(h)$, then the equilibrium is necessarily Type 2. Moreover, in the same

corollary, we show that $\underline{z}_2^* > 0$ and $\bar{z}_2^* < 1$. Consequently, if $(h, z_{\text{AI}}) \in \mathcal{Z}_k$, then $\sup W^* = \inf M^* = z_{\text{AI}} = \hat{z}(h)$ (so $W^* = W$ and that $M^* = M$), but that there is within-worker and within-manager displacement as $W_a^* \neq \emptyset$ and $M_a^* \neq \emptyset$.

C.2 Proof of Lemma 2

The size of a two-layer firm with a manager with knowledge m and workers with knowledge z is $mn(z)$. The size of a single-layer firm is the knowledge of its independent producer. The first part of the lemma then follows because $n(z) \geq 1$ for all $z \in [0, 1]$, and the fact that independent producers are always less knowledgeable than managers (due to occupational stratification).

Now the second part. Take two arbitrary firms f and f'. Without loss of generality, there are three cases to consider: (i) f and f' are two-layer firms, (ii) f and f' are single-layer firms, and (iii) f is a two-layer firm and f' is a single layer firm. In case (i), the statement follows directly from positive assortative matching. In case (ii), the statement follows because size equals productivity in one-layer firms. Finally, in case (iii), f is necessarily strictly more productive (due to occupational stratification) and bigger (because f has more layers than f'). □

C.3 Proof of Proposition 4

As noted in the main text, in the case of two-layer organizations, the best worker determines the maximum firm size (as the best manager always has knowledge 1), while the knowledge of the worst manager determines the minimum firm size and the productivity of the least productive firm (as the worst worker always has knowledge 0).

With this in mind, consider the case in which $(h, z_{\text{AI}}) \in \mathcal{Z}_w$. Then $\sup W^* < \sup W$ and $\inf M^* < \inf M$ (as $W^* \subset W$ and $M^* \supset M$), so the size of the biggest firm and the size and productivity of the smallest/least productive firm decreases with AI. Analogously, when $(h, z_{\text{AI}}) \in \mathcal{Z}_m$, then $\sup W^* > \sup W$ and $\inf M^* > \inf M$ (as $W^* \supset W$ and $M^* \subset M$). Hence, in this case, the size of the biggest firm and the size and productivity of the smallest/least productive firm increases with AI. Finally, if $(h, z_{\text{AI}}) \in \mathcal{Z}_k$, then $\sup W^* = \sup W$ and $\inf M^* = \inf M$, so the size of the biggest firm and the size and productivity of the smallest/least productive firm do not change with AI. □

C.4 Proof of Proposition 5

First, since two-layer firms are always larger than one-layer firms, the results regarding the size of the largest firm among all firms mimic those of two-layer firms described in Proposition 4. Consider now the smallest/least productive firm. For ease of exposition, we refer to this firm as the "smallest firm" and divide the proof in two cases:

• $h > h_0$.— In this case, the smallest firm is a single-layer firm in both the pre-AI and post-AI equilibrium. Hence, in both situations, the size and productivity of the smallest firm are equal to the knowledge of the worst independent producer, which is also the knowledge of the best worker. Consequently, the results in this case mimic those of two-layer firms, as $\sup W^* < \sup W$ when $(h, z_{\text{AI}}) \in \mathcal{Z}_w$, $\sup W^* > \sup W$ when $(h, z_{\text{AI}}) \in \mathcal{Z}_m$, and $\sup W^* = \sup W$ when $(h, z_{\text{AI}}) \in \mathcal{Z}_k$ (as formally shown in the proof of Proposition 4).

• $h \leq h_0$.— In this case, the smallest firm pre-AI is a two-layer firm with productivity $\hat{z} = \sup W$, while the smallest firm post-AI is a single-layer firm with productivity $\sup W^*$. Thus, following the exact same reasoning as above, we conclude that the results regarding the productivity of the smallest firm mimic those of two-layer firms described in Proposition 4.

Now consider size. Pre-AI, the size of the smallest firm is $\hat{z}n(0)$, while post-AI, the size of the smallest firm is $\sup W^*$. First, take $(h, z_{\text{AI}}) \in \mathcal{Z}_w$. By Proposition 4, $\sup W^* < \sup W = \hat{z} < \hat{z}n(0)$, so, in this case, the result mimics that of two-layer firms described in Proposition 4. Second, take $(h, z_{\text{AI}}) \in \mathcal{Z}_k$. Then $\sup W^* = z_{\text{AI}} = \hat{z} < \hat{z}n(0)$. Hence, in contrast to the results regarding two-layer firms, now the size of the smallest firm decreases with AI. Finally, take $(h, z_{\text{AI}}) \in \mathcal{Z}_m$. Then $\sup W^*$ is either z_{AI} (in a Type 2 equilibrium) or a $\bar{z}_3^*$ (in a Type 3). Hence, in contrast to the results regarding two-layer firms, now the size of the smallest firm can increase or decrease with AI, as $\hat{z}n(0) \gtrless z_{\text{AI}}$ and $\hat{z}n(0) \gtrless \bar{z}_3^*$. □

C.5 Proof of Proposition 6

Due to space constraints, we provide the proof assuming that $h \leq h_0$. The proof for the case $h > h_0$ can be found in Section 2 of the online Appendix. As noted in the main text, a worker's productivity increases if and only if her managerial match improves. Similarly, the size of the firm of a given manager increases if and only the knowledge of her workers increases.

• $(h, z_{AI}) \in \mathcal{Z}_w$.— 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 $(h, z_{\text{AI}}) \in \mathcal{Z}_w$.

Proving that every $z \in W_p^*$ also has a worse manager is more involved. Since $(h, z_{\text{AI}}) \in \mathcal{Z}_w$, 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:

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

Consequently, if $z \in W_p^* \cap W = W_p^*$, then $\int_{\underline{z}_1^*}^{f_1^*(z;\underline{z}_1^*)} dG(u) = \int_{\hat{z}}^{f(z;\hat{z})} dG(u)$. Hence, it must be that $f_1^*(z; \underline{z}_1^*) < f(z; \hat{z})$ 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:

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

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 $z_{\text{AI}} \geq e_j^*(z;\cdot) > e(z;\hat{z})$, where the first inequality is due to the fact that AI is the best worker, and the second inequality follows because $e_j^*(z;\cdot) > e(z;\hat{z})$ if $f_j^*(z;\cdot) < f(z;\hat{z})$ (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 $e_j^*(z;\cdot) > e(z;\hat{z})$, 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 $e(z;\hat{z})$. □

• $(h, z_{AI}) \in \mathcal{Z}_m$.— 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 $(h, z_{\text{AI}}) \in \mathcal{Z}_m$.

We now show that the same holds for every $z \in M_p^*$. Since $(h, z_{\text{AI}}) \in \mathcal{Z}_m$, 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:

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

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 $z_{\text{AI}} \leq f_j^*(z;\cdot) < f(z;\hat{z})$, where the first inequality is due to the fact that AI is the worst manager, and the second inequality follows because $f_j^*(z;\cdot) < f(z;\hat{z})$ if $e_j^*(z;\cdot) > e(z;\hat{z})$ (which we already showed it 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 $f_j^*(z;\cdot) < f(z;\hat{z})$ for $j = 2, 3$. □

• $(h, z_{AI}) \in \mathcal{Z}_k$.— We first show that each $z \in W^* = W$ is managed by a worse manager post-AI compared to pre-AI (strictly so for all $z \neq 0$).

As shown in Corollary 1.1 of the online Appendix, when $(h, z_{\text{AI}}) \in \mathcal{Z}_k$, then the equilibrium is necessarily Type 2 with $\underline{z}_2^* > 0$ 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 inequality 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.6 Proof of Proposition 7

We begin with the following preliminary claim:

Claim 1. *(i)* $\Delta(z_{AI}) < 0$*, (ii)* $\Delta(0) > 0$ *if* $(h, z_{AI}) \in \mathcal{Z}_m \cup \mathcal{Z}_k$*, (iii)* $\Delta(0) < 0$ *if* $h > h_0$ *and* $(h, z_{AI}) \in \mathcal{Z}_w$*, and (iv)* $\Delta(1) \geq 0$*, with strict inequality if and only if* $h < h_0$*.*

Proof. See Section 4 of the online Appendix. □

We now provide the proof of the proposition.

Part (i).— Suppose first that $h > h_0$. We want to show that there exists $z < z_{\text{AI}}$ such that $\Delta(z) > 0$ if and only if $z_{\text{AI}} > \text{Z}(h) = \bar{z}(h)$. We begin by showing that if there exists a $z < z_{\text{AI}}$ such that $\Delta(z) > 0$, then $z_{\text{AI}} > \text{Z}(h)$. To do this, we prove the contrapositive statement: If $z_{\text{AI}} \leq \text{Z}(h) = \bar{z}(h)$, then there is no such z. Indeed, note that if $z_{\text{AI}} \in [\underline{z}(h), \bar{z}(h)]$, then we know that AI has no effects on wages, so $\Delta(z) = 0$ in this case. And if $z_{\text{AI}} < \underline{z}(h)$, then $\Delta(0) < 0$ according to Claim 1 (as AI has the knowledge of a pre-AI worker), which immediately implies—due to Lemma 3— that $\Delta(z) < 0$ for all $z < z_{\text{AI}}$.

We now show that if $z_{\text{AI}} > \text{Z}(h) = \bar{z}(h)$, then there exists a $z < z_{\text{AI}}$ such that $\Delta(z) > 0$. Indeed, in this case, AI has the knowledge of a pre-AI manager, so $\Delta(0) > 0$ according to Claim 1. Consequently, Lemma 3 immediately implies that—irrespective of the value of z_{AI}—there exists $\zeta \in [0, z_{\text{AI}})$ such that $\Delta(z) > 0$ for $z \in [0, \zeta)$ and $\Delta(z) < 0$ for $z \in (\zeta, z_{\text{AI}}]$.

Now suppose $h \leq h_0$. We want to show that there exists $z < z_{\text{AI}}$ such that $\Delta(z) > 0$ if and only if $z_{\text{AI}} > \text{Z}(h)$, where $\text{Z}(h) \in (0, \hat{z}(h))$. We begin by constructing $\text{Z}(h)$, then we show that the statement is true. Let $\Delta(0; z_{\text{AI}}, h) \equiv w^*(0; z_{\text{AI}}, h) - w(0; h)$, and define $\text{Z}(h)$ as the solution to $\Delta(0; \text{Z}(h), h) = 0$. We first show that $\text{Z}(h)$ exists and is unique, and that $\text{Z}(h) \in (0, \hat{z}(h))$. We do this by showing that for a given h, $\Delta(0; z_{\text{AI}}, h)$ crosses zero once as we go from $z_{\text{AI}} = 0$ to $z_{\text{AI}} = 1$, and that this crossing point is at a $z_{\text{AI}} < \hat{z}(h)$.

Indeed, if $z_{\text{AI}} \geq \hat{z}(h)$, then Claim 1 states that $\Delta(0; z_{\text{AI}}, h) > 0$ irrespective of h (as $(h, z_{\text{AI}}) \in \mathcal{Z}_m \cup \mathcal{Z}_k$). Moreover, as shown in Section 1 of the online Appendix (specifically, see Corollary 1.1), when $z_{\text{AI}} \to 0$, the equilibrium is Type 1 irrespective of the value of h, so:

$$\lim_{z_{\text{AI}} \to 0} \Delta(0; z_{\text{AI}}, h) = \lim_{z_{\text{AI}} \to 0} [\underline{z}_1^*(z_{\text{AI}})(1 - h) - w(0; h)] = -w(0; h) < 0$$

where the second-to-last equality follows because $\underline{z}_1^*(z_{\text{AI}}) \to 0$ when $z_{\text{AI}} \to 0$, as can be easily be proven from the condition that determines $\underline{z}_1^*(z_{\text{AI}})$ (see the statement of Proposition 2).

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

Having constructed $\mathrm{Z}(h)$, we prove that there exists $z < z_{\mathrm{AI}}$ such that $\Delta(z; z_{\mathrm{AI}}, h) > 0$ if and only if $z_{\mathrm{AI}} > \mathrm{Z}(h)$. First we show that if there exists a $z < z_{\mathrm{AI}}$ such that $\Delta(z; z_{\mathrm{AI}}, h) > 0$, then $z_{\mathrm{AI}} > \mathrm{Z}(h)$. To do this, we prove the contrapositive statement: If $z_{\mathrm{AI}} \leq \mathrm{Z}(h)$, then there is no such z. Indeed, $\Delta(0; z_{\mathrm{AI}}, h) \leq 0$ for all $z_{\mathrm{AI}} \leq \mathrm{Z}(h)$ as shown above. Hence, Lemma 3 immediately implies that $\Delta(z; z_{\mathrm{AI}}, h) \leq 0$ for all $z < z_{\mathrm{AI}}$.

Now we prove that if $z_{\mathrm{AI}} > \mathrm{Z}(h)$, then there exists a $z < z_{\mathrm{AI}}$ such that $\Delta(z; z_{\mathrm{AI}}, h) > 0$. Indeed, as shown above, $z_{\mathrm{AI}} > \mathrm{Z}(h)$ then $\Delta(0; z_{\mathrm{AI}}, h) > 0$. Consequently, Lemma 3 immediately implies that there exists $\zeta \in [0, z_{\mathrm{AI}})$ such that $\Delta(z; z_{\mathrm{AI}}, h) > 0$ for $z \in [0, \zeta)$ and $\Delta(z; z_{\mathrm{AI}}, h) < 0$ for $z \in (\zeta, z_{\mathrm{AI}}]$. □

Part (ii).— We first show that if $\Delta(z) < 0$ for all $z \in (z_{\mathrm{AI}}, 1)$ then $h > h_0$ and $(h, z_{\mathrm{AI}}) \in \mathcal{Z}_m$. Suppose by contradiction that this is not the case. Note that if $h < h_0$, then $\Delta(1) > 0$. Hence, Lemma 3 implies that there exists a $\zeta \in (z_{\mathrm{AI}}, 1)$, such that $\Delta(z) \leq 0$ for all $z \in [z_{\mathrm{AI}}, \zeta]$ and $\Delta(z) > 0$ for all $z \in (\zeta, 1)$. This contradicts the premise that $\Delta(z) < 0$ for all $z \in (z_{\mathrm{AI}}, 1)$, so it must be that $h > h_0$ and that AI has the knowledge of a pre-AI worker (recall that AI cannot have the knife-edge knowledge of both when $h > h_0$).

We then claim that, in such a case, $\Delta(z) > 0$ for $z < 1$ but z sufficiently close to 1, so in this case, we also arrive at a contradiction. Indeed, since AI has the knowledge of a pre-AI worker, then $z \in M_a^*$ when $z < 1$ but $z \to 1$. Consequently, $\Delta'(z) = n(z_{\mathrm{AI}}) - n(e(z; \bar{z})) = n(z_{\mathrm{AI}}) - n(\underline{z}) < 0$, where the inequality follows because $z_{\mathrm{AI}} < \underline{z}$, as AI has the knowledge of a pre-AI worker. But if so, and given that $\Delta(1) = 0$ when $h \geq h_0$, then $\Delta(z) > 0$ for $z < 1$, but $z \to 1$.

Now we show that if $h > h_0$ and $(h, z_{\mathrm{AI}}) \in \mathcal{Z}_m$, then $\Delta(z) < 0$ for all $z \in (z_{\mathrm{AI}}, 1)$. We claim that $\Delta(z) < 0$ for $z < 1$ but z sufficiently close to 1, so Lemma 3 immediately implies that $\Delta(z) \leq 0$ for all $z \in [z_{\mathrm{AI}}, 1]$, with strict inequality if $z \in [z_{\mathrm{AI}}, 1)$. Indeed, in this case, the equilibrium can only be Type 2 or Type 3. If it is Type 2, then $z \in M_a^*$ when $z < 1$ but $z \to 1$. In such a case, $\Delta'(z) = n(z_{\mathrm{AI}}) - n(e(z; \bar{z})) = n(z_{\mathrm{AI}}) - n(\underline{z}) > 0$, as $z_{\mathrm{AI}} > \underline{z}$. On the other hand, if the equilibrium is Type 3, then $z \in M_p^*$ when $z < 1$ but $z \to 1$. In that case, $\Delta'(z) = n(\bar{z}_3^*) - n(e(z; \bar{z})) = n(\bar{z}_3^*) - n(\underline{z}) > 0$, as $\bar{z}_3^* > \underline{z}$ given that $W^* \supset W$ when $(h, z_{\mathrm{AI}}) \in \mathcal{Z}_m$. □

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 (2022): "Automation and Polarization," *National Bureau of Economic Research*.

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.

——— (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," Working Paper 23928, National Bureau of Economic Research.

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.

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," Tech. rep., National Bureau of Economic Research.

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.

ELOUNDOU, T., S. MANNING, P. MISHKIN, AND D. ROCK (2023): "Gpts are gpts: An early look at the labor market impact potential of large language models," *arXiv preprint arXiv:2303.10130*.

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," *The 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.

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.

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

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.