

Artificial Intelligence in the Knowledge Economy*

Enrique Ide & Eduard Talamàs

May 31, 2024

Abstract

This paper provides a new framework for studying the impact of Artificial Intelligence (AI) on the organization of knowledge work. We incorporate AI into an economy where humans endogenously form hierarchical firms: Less knowledgeable agents become “workers” solving routine problems, while more knowledgeable agents become “solvers” handling exceptions. We model AI as an algorithm that uses compute to mimic humans. We compare the equilibrium before and after AI’s introduction, distinguishing between “basic” AI (with knowledge equivalent to pre-AI workers) and “advanced” AI (with knowledge equivalent to pre-AI solvers). We show that basic AI increases the knowledge content of human work, leading to smaller, less productive, and less decentralized firms. In contrast, advanced AI decreases the knowledge content of human work, resulting in larger, more productive, and more decentralized firms. In any case, the most knowledgeable humans benefit from AI, while the least knowledgeable benefit only when AI is sufficiently advanced. We discuss how these effects depend on AI’s autonomy and the availability of compute.

*IESE Business School, C/ de Arnús i de Garí 3-7, 08034 Barcelona, Spain (eide@iese.edu and etalamas@iese.edu). We have benefited from the comments of Nano Barahona, Nicolás Figueroa, Andrea Galeotti, Luis Garicano, Benjamin Golub, Jason Hartline, Mike Ostrovsky, Cristóbal Otero, Sebastián Otero, and participants of the Spring 2024 NBER Organizational Economics Working Group and the various other venues where this work has been presented. We declare we have no relevant or material financial interests that relate to the research described in this paper.

1 Introduction

Artificial Intelligence (AI) is a new and powerful form of automation based on machines that can perform sophisticated knowledge work, including coding, research, and complex problem solving. Although the potential of AI to reshape the landscape of work is undeniable, its precise implications have become the center of a growing controversy (Meserole, 2018; Brynjolfsson, 2022; Johnson and Acemoglu, 2023; Autor, 2024; Acemoglu, 2024). This controversy stems from two factors. First, there is uncertainty about whether lessons from previous automation waves—which led to the creation of tools proficient at handling repetitive tasks—apply to AI (Muro et al., 2019; Agrawal et al., 2019). Second, since AI is still in its infancy and individuals and firms are still experimenting with it (McElheran et al., 2023), current empirical evidence cannot account for the equilibrium effects of AI.¹

In this paper, we provide a new framework for studying the equilibrium effects of AI on the future of work. The novelty of our approach is that it explicitly incorporates the peculiarities of both AI and knowledge work by embedding algorithms and computing power (or “compute”) in a canonical model of a knowledge economy: The knowledge hierarchies first introduced by Garicano (2000).² Using this framework, we provide predictions about the equilibrium effects of AI on organizational and labor outcomes, including (i) occupational choice, (ii) firm productivity and decentralization, and (iii) the distribution of labor income. Our analysis highlights how the impact of AI depends on its capabilities and the amount of compute available in the economy.

Our starting point—the pre-AI economy—is the baseline model of Antràs et al. (2006) and Fuchs et al. (2015). Labor and knowledge are the sole inputs in production. Humans are endowed with one unit of time and are heterogeneous in terms of knowledge. Individuals use their time to pursue production opportunities but encounter problems of varying difficulty during the production process. Output is produced when an individual can successfully solve the problem she confronts, which occurs when her knowledge exceeds the problem’s difficulty. If a human cannot solve a problem on her own, she may seek help from another human. Help, however, is costly in terms of time.

The competitive equilibrium of the pre-AI economy involves humans either pursuing production opportunities on their own (becoming “independent producers”) or joining hierarchical firms. These firms have the following properties. First, they consist of many “workers” (who pursue production opportunities) and one “solver” (who is more knowledgeable than the workers and specializes in assisting them with unresolved problems). Second, they exhibit positive assortative matching: A firm with more knowledgeable workers has a more knowledgeable solver. Third, a firm’s productivity is increasing in the knowledge of its solver (as a more knowledgeable solver enables the resolution of more problems), and a firm’s decentralization is increasing in the knowledge of its workers (as more

¹For experimental evidence on the productivity effects of AI, see, for example, Noy and Zhang (2023), Brynjolfsson et al. (2023), Peng et al. (2023), and Otis et al. (2023).

²There 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).

knowledgeable workers solve more problems on their own). Fourth, solvers in more decentralized firms have a greater span of control because a solver can assist more workers when each worker is less likely to require help.

Our innovation is to incorporate AI into this otherwise canonical setting. We model AI as an algorithm with an exogenously fixed knowledge level that requires compute to run. This algorithm is comparable to human intelligence in that it can mimic the behavior of humans in all three possible roles in the economy—independent producer, worker, and solver. This assumption is motivated by the observation that there appear to be stronger incentives for developing human-like AI than for developing human-augmenting AI (Acemoglu and Restrepo, 2019; Brynjolfsson, 2022; Johnson and Acemoglu, 2023).³ The amount of compute in the economy is exogenous and, in the baseline setting, scarce relative to production opportunities. The price of compute is determined endogenously in equilibrium.

Human and artificial intelligence differ, however, in one key aspect: AI can be used *at scale* in the following two senses. First, the same algorithm can be leveraged across all units of compute (the use of human knowledge, in contrast, is constrained by the time of the human possessing it). Second, we assume that the binding constraint in human-AI interactions is human time, not compute. That is, even though compute is scarce relative to production opportunities, it is abundant relative to human time. Our motivation for these assumptions is the nonrival nature of digital information (Brynjolfsson and McAfee, 2016; Goldfarb and Tucker, 2019) and the exponential growth of computational capacity over the past two centuries (Nordhaus, 2007).

We start by characterizing the post-AI equilibrium and showing that the equilibrium price of compute is equal to the fraction of problems that AI can solve on its own. This follows because compute is abundant relative to human time, and thus, some compute must be allocated to independent production. Moreover, we show that if AI is used as a worker, then it is necessarily the *most* knowledgeable worker post-AI; thus, it is supervised by the *most* knowledgeable human solvers. In addition, if AI is used as a solver, then it is necessarily the *least* knowledgeable solver post-AI, and thus, it assists the *least* knowledgeable human workers.

A notable property of the equilibrium is that even when AI has an absolute advantage over a fraction of the population, its introduction does not lead to unemployment among humans. The key driver of this result is that compute is scarce relative to production opportunities, so it continues to be worthwhile for every human to be employed in some capacity. All individuals less knowledgeable than AI sort into worker positions because that is their comparative advantage: They are less likely than AI to succeed on their own, so it is more valuable for them to receive assistance. This result reflects that it may not be cost-effective to deploy AI in all tasks in which such technology is superior to humans (Svanberg et al., 2024; Acemoglu, 2024).

³Furthermore, as Kahneman (2019) emphasizes “[there might not be] very much that we can do that computers will not eventually be programmed to do.”

We then turn to our main endeavor: Analyzing the effects of AI on knowledge work by comparing the pre- and post-AI equilibrium. We show that if AI has the knowledge of a pre-AI worker (a relatively “basic” form of AI), then its introduction displaces humans from routine production work to specialized problem solving; i.e., it increases the knowledge content of human work. In the process, it creates less productive firms and destroys the most decentralized firms. Moreover, it decreases the productivity of all workers who remain workers, increases the span of control of the worst solvers who remain solvers, and decreases the span of control of the best solvers who remain solvers.

In contrast, if AI has the knowledge of a pre-AI solver (a relatively “advanced” form of AI), then its introduction displaces humans from specialized problem solving to routine production work; i.e., it decreases the knowledge content of human work. This leads to the destruction of the least productive firms and the creation of more decentralized firms. As a result, AI increases the span of control of all solvers who remain solvers, increases the productivity of the worst workers who remain workers, and decreases the productivity of the best workers who remain workers. Furthermore, in this case, AI may also lead to the creation of superstar firms that have “scale without mass” (Brynjolfsson et al., 2008; Autor et al., 2020). These are firms that (i) are at the top of the post-AI size distribution in terms of output, (ii) are larger than the largest pre-AI firms, and (iii) use a single human solver to supervise production work by AI.

For intuition, let us focus on the case in which AI has the knowledge of a pre-AI worker. In this case, AI serves as a relatively inexpensive technology for performing routine work, reducing workers’ wages and increasing the attractiveness of creating hierarchical firms. The result is a surge in the demand for solvers to match with the less expensive workers, which induces the most knowledgeable routine workers of the pre-AI equilibrium to switch to specialized problem solving. Hence, AI destroys the most decentralized firms (as the most knowledgeable pre-AI workers become solvers) and creates less productive firms (as the newly appointed solvers are less knowledgeable than the least knowledgeable pre-AI solvers).

Moreover, AI not only displaces humans across occupations but also affects those individuals who are not occupationally displaced. The reason for this is that its introduction induces a complete reorganization of all matches in the economy. In particular, AI’s introduction worsens the match (and hence the productivity) of every worker who remains a worker because the best solvers switch to working with AI. Relatedly, since the worst pre-AI solvers become average solvers post-AI, they switch from assisting the least knowledgeable workers to assisting average workers post-AI (increasing their span of control). Finally, the best pre-AI solvers (who continue to be the best solvers post-AI) see their span of control reduced because they switch from assisting the best pre-AI workers to assisting a less knowledgeable AI.

We then turn to studying the impact of AI on labor income. Even though AI increases total labor income, its introduction necessarily creates winners and losers in the labor market. More precisely, we show that the least knowledgeable humans benefit when AI’s knowledge is high because its intro-

duction gives them access to better and/or less expensive solvers in this case. The most knowledgeable humans, in contrast, always benefit from AI—irrespective of AI’s knowledge level—because it serves as a relatively inexpensive technology to leverage their knowledge.

Finally, we provide two extensions to analyze how our results depend on compute availability and AI’s autonomy. In the first extension, we study the case where compute is abundant not only relative to time but also relative to production opportunities.⁴ In this case, the equilibrium price of compute is zero, and AI leads to technological unemployment: All humans who are less knowledgeable than AI become unemployed. Organizations, however, still display a hierarchical structure in that the most knowledgeable humans specialize in tackling the problems that AI cannot solve. Moreover, while in the baseline setting, the income generated by AI accrues exclusively to labor and compute, in this extension, the owners of production opportunities obtain a fraction of these gains.

In the second extension, we return to the baseline where production opportunities are abundant relative to the available compute but introduce a second dimension of intelligence in which AI cannot mimic humans. In this case, as in the first extension, the equilibrium price of compute is zero, but for a different reason: AI is no longer autonomous. All individuals then specialize in assisting AI in the dimension where it lacks human-like intelligence. As a result, there is no technological unemployment, and all the gains from AI’s introduction accrue to labor.

Related Literature

This paper contributes to two different streams of literature. On the one hand, it introduces automation and AI to the literature on knowledge hierarchies. On the other hand, it incorporates the peculiarities of AI and knowledge work into the literature on automation.

The literature on knowledge hierarchies starts with [Garicano \(2000\)](#), who introduces the model and describes the circumstances under which knowledge hierarchies are optimal when agents are homogenous.⁵ [Garicano and Rossi-Hansberg \(2004, 2006\)](#) embed this model in a setting with heterogeneous agents to study inequality. [Fuchs et al. \(2015\)](#) characterize the equilibrium contractual arrangements when there is asymmetric information about knowledge.⁶ Our innovation with respect to this literature is to introduce and study the impact of a technology capable of automating knowledge work.

In the context of knowledge hierarchies, the most closely related paper is that of [Antràs et al.](#)

⁴In the Online Appendix, we also study the opposite extreme where the amount of compute is small.

⁵The literature on knowledge hierarchies is part of the larger literature on organizational economics. Some recent theoretical contributions to this literature include [Dessein and Santos \(2006\)](#), [Cremer et al. \(2007\)](#), [Alonso et al. \(2008\)](#), and [Dessein et al. \(2016\)](#). See [Gibbons and Roberts \(2013\)](#) for a survey.

⁶Other important contributions to the literature on knowledge hierarchies include [Garicano and Hubbard \(2007\)](#), [Garicano and Rossi-Hansberg \(2012\)](#), [Bloom et al. \(2014\)](#), [Caicedo et al. \(2019\)](#), [Gumpert et al. \(2022\)](#), and [Carmona and Lao-hakunakorn \(Forthcoming\)](#).

(2006). They study the effects of offshoring by comparing the equilibrium of a closed economy with one in which firms can form international teams. Our paper differs from theirs in two key respects. First, while offshoring gives firms access to a population of humans with different knowledge levels, AI gives firms access to an algorithm that can solve problems *at scale*. As we discuss in detail in Section 4.5, this implies that the effects of AI are qualitatively different than those of offshoring.

Second, because we explicitly incorporate algorithms and compute into the model, we can study how the impact of AI depends on its capabilities, such as its knowledge and autonomy, and the availability of compute. These dimensions are absent in Antràs et al. (2006) because they focus on a different phenomenon.

Our paper also contributes to the literature on automation, which 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) contend that, by depressing wages, automation also encourages the creation of new tasks in which labor has a comparative advantage. In a different vein, Autor et al. (2003) and Acemoglu and Autor (2011) argue that routine tasks, which are easier to automate than are other types of tasks, are typically handled by those workers in the middle of the skill and wage distribution. Hence, the advent of automation can explain the emergence of employment and wage polarization. Acemoglu and Loebbing (2024) add to this point by showing that automating middle-skill tasks becomes more profitable when the cost of capital decreases.⁷

Our innovation with respect to this literature is that to understand AI’s distinctive impact, we focus on a specific type of automation: The automation of knowledge work. This is an important part of the economy that has been partly shielded from previous automation waves. To do this, we take a stance on the nature of AI and embed this algorithm in a canonical model of the knowledge economy.⁸ This approach allows us to study the effects of AI on the endogenous organization of knowledge work, including how its impact depends on the knowledge of the existing population, the state of communication technologies, AI’s capabilities, and the availability of compute relative to human time and societal needs.

2 The Model

This section introduces the model, discusses its main assumptions, and characterizes the pre-AI equilibrium. This pre-AI benchmark is the baseline model of Antràs et al. (2006) and Fuchs et al. (2015).

⁷Other important contributions include Aghion et al. (2017), Acemoglu and Restrepo (2022), Moll et al. (2022), Azar et al. (2023), Korinek and Suh (2024), Acemoglu (2024), and Jones (2024).

⁸As explained by Garicano and Rossi-Hansberg (2015), knowledge hierarchies are a particular specification of the task-based framework where tasks are hierarchical and the relationship between them arises from an explicit organizational problem.

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 . The distribution of knowledge in the population is given by a continuous probability distribution with full support on $[0, 1]$, cumulative distribution function G , and density g . The knowledge of each individual is perfectly observable.

There is a large measure of identical competitive firms. Production occurs inside firms, which are the residual claimants of all output. Labor and knowledge are the sole inputs in production. Firms have no fixed costs and two layers at most.

Single-layer firms hire a single human to produce. This “independent producer” devotes her full unit of time to pursuing a single production opportunity. Each production opportunity is linked to a problem whose difficulty x is ex-ante unknown and distributed uniformly on $[0, 1]$, independently across problems. If the knowledge of the human engaging in production exceeds the problem’s difficulty, she solves the problem and produces one unit of output. Otherwise, no output is produced. Given that the distribution of knowledge in the population is arbitrary, assuming that $x \sim U[0, 1]$ is simply a normalization, and it is, therefore, without loss. Under this normalization, an individual’s knowledge z is interpreted as the fraction of problems she can solve on her own.

Two-layer firms hire one “solver” and multiple “workers,” where all workers have the same knowledge. This restriction is without loss because—as we show below—the equilibrium matching arrangement between workers and solvers exhibits strict positive assortative matching.

As in single-layer firms, each worker in a two-layer firm devotes her time to a single production opportunity. The difference is that if the worker cannot solve the problem on her own, she can ask the solver for help. If the solver’s knowledge exceeds the problem’s difficulty, she communicates the solution to the corresponding worker, who then produces a unit of output. Otherwise, no production takes place. However, communication is costly in that whenever a worker asks the solver for help, that exchange consumes $h \in (0, 1)$ units of the solver’s time. Hence, a two-layer organization optimally hires exactly $n(z) = [h \times (1 - z)]^{-1}$ workers of knowledge z to fully exploit its solver’s time.

Figure 1 depicts the two possible firm configurations of the pre-AI world, where the letter attached to each human corresponds to her knowledge.

Artificial Intelligence.— We model AI as an algorithm that requires compute to run and has an exogenously fixed level of knowledge $z_{AI} \in [0, 1]$.⁹ This algorithm can replicate the behavior of humans in all three possible roles in the economy: independent producer, worker, and solver.

All firms have access to AI. Thus, in contrast to the pre-AI economy, firms decide not only their organizational structure but also whether to use this technology. Firms that use AI are identical to those that do not use AI, except that they automate part of the production process. To do this, they

⁹We assume that $z_{AI} < 1$ because the equilibrium has a discontinuity at $z_{AI} = 1$. In Section 4 of the Online Appendix, we consider the case where $z_{AI} = 1$.




Figure 1: The Two Possible Firm Configurations in the Pre-AI World

must rent one unit of compute per human replaced. The amount of compute in the economy, which we denote by μ , is exogenous. We assume there are more production opportunities than the human time and compute necessary to pursue them all, i.e., production opportunities are abundant relative to compute and time.

Figure 2 illustrates the five possible post-AI firm configurations. In addition to single- and two-layer firms that hire only humans (the only possible pre-AI firms), there are three additional possible configurations: Single-layer automated firms (which use AI as an independent producer), bottom-automated firms (which use AI exclusively as a worker), and top-automated firms (which use AI exclusively as a solver). Note that a firm will never use AI in both layers of the organization because an AI solver knows the solution to the same set of problems as does an AI worker.

Wages, Prices, and Profits.—Let $w(z)$ be the wage of a human with knowledge z and denote by r the rental rate of one unit of compute. All agents in this economy are risk neutral and maximize their income. We normalize the value of each unit of output to one.

The problem of a firm is to decide (i) whether to use humans or AI in production (and the knowledge of the humans hired, 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 a solver (and the knowledge of the human hired, when appropriate).

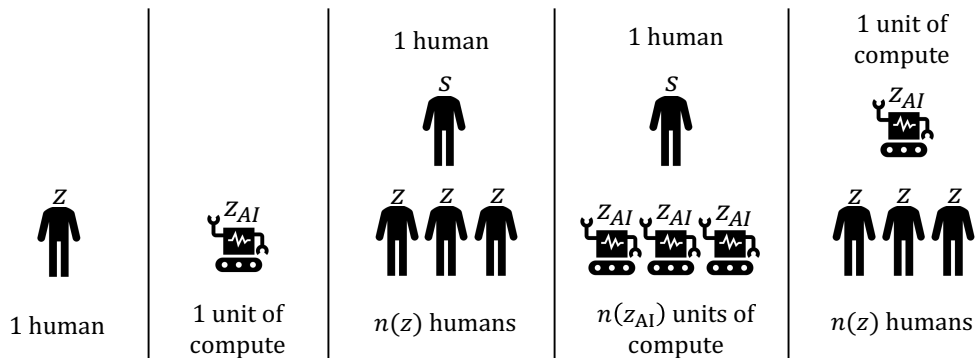


Figure 2: The Five Possible Firm Configurations in the Post-AI World

The profits of a single-layer organization are as follows:

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

In other words, the profit of a single-layer nonautomated firm is the expected output z of its independent producer net of her wage $w(z)$. The profit of a single-layer automated firm is the expected output of AI as an independent producer z_{AI} minus the cost of renting one unit of compute r .

The profit of a two-layer organization, in turn, depends on whether it uses AI as a solver (i.e., automates the top layer, a “ tA ” firm), uses AI as a worker (i.e., automates the bottom layer, a “ bA ” firm), or it does not use AI (an “ nA ” firm):

$$\Pi_2^{tA}(z) = n(z)[z_{\text{AI}} - w(z)] - r \quad (\text{where } z \leq z_{\text{AI}})$$

$$\Pi_2^{bA}(s) = n(z_{\text{AI}})[s - r] - w(s) \quad (\text{where } z_{\text{AI}} \leq s)$$

$$\Pi_2^{nA}(s, z) = n(z)[s - w(z)] - w(s) \quad (\text{where } z \leq s)$$

where z and s denote the knowledge of a human worker and a human solver, respectively, and we use the fact that no two-layer firm hires a solver who is less knowledgeable than its workers. In all three cases, the profit of a firm is its expected output minus the cost of the resources it uses. For instance, in the case of a tA firm that hires workers with knowledge z , its total expected output is $n(z)z_{\text{AI}}$, while the cost of resources is $n(z)w(z) + r$.

Competitive Equilibrium.— Let μ_i , μ_w , and μ_s be the amount of compute rented for independent production, production in two-layer firms, and the supervision of humans, respectively. We denote by I the set of humans hired as independent producers, and by W_p and W_a the set of human workers assisted by human solvers and AI, respectively.¹⁰ Similarly, we denote by S_p the set of humans who assist other humans, and by S_a the set of humans who assist the production work of AI. Finally, we let $m : W_p \rightarrow S_p$ be the function describing the pointwise matching arrangement generated by the hiring decisions of nA firms; i.e., $m(z)$ is the knowledge of the solver assisting workers with knowledge z .¹¹

Definition (Competitive Equilibrium). An equilibrium consists of nonnegative amounts (μ_i, μ_w, μ_s) , sets (W_p, W_a, I, S_p, S_a) , a matching function $m : W_p \rightarrow S_p$, a wage schedule $w : [0, 1] \rightarrow \mathbb{R}_{\geq 0}$ and a rental rate of compute $r \in \mathbb{R}_{\geq 0}$, such that:

1. Firms optimally choose their structure (while earning zero profits).
2. tA firms hiring $n(z)$ workers with knowledge $z \in W_a$ rent one unit of compute.
3. bA firms hiring a solver with knowledge $s \in S_a$ rent $n(z_{\text{AI}})$ units of compute.

¹⁰We use the subscript “ p ” (for people) instead of “ h ” (for humans), to avoid any confusion with the helping cost h .

¹¹As discussed in Fuchs et al. (2015), $m(z)$ is an approximation of a firm that hires a small interval of human workers $(z - \varepsilon_1, z + \varepsilon_1)$ and a small interval of human solvers $(m(z) - \varepsilon_2, m(z) + \varepsilon_2)$ with the requirement that the mass of solvers in the firm $\int_{m(z) - \varepsilon_2}^{m(z) + \varepsilon_2} dG(u)$ is equal to h times the mass of problems left unsolved by the workers in the firm $\int_{z - \varepsilon_1}^{z + \varepsilon_1} (1 - u) dG(u)$. The latter requirement captures that nA firms optimally hire $n(z) = [h \times (1 - z)]^{-1}$ workers of knowledge z .

4. nA firms hiring $n(z)$ workers with knowledge $z \in W_p$ hire a solver with knowledge $m(z) \in S_p$.
5. Markets clear: (i) $\mu_i + \mu_w + \mu_s = \mu$, and (ii) the union of the sets (W_p, W_a, I, S_p, S_a) is $[0, 1]$ and the intersection of any two of these sets has measure zero.

Note that the human workers in W_p are endogenously matched with the human solvers in S_p according to the pointwise matching function m . In contrast, all the humans in W_a are matched with an AI solver (which has knowledge z_{AI}), while all the humans in S_a are matched with AI workers (whose knowledge is z_{AI}).

Compute is “Abundant” Relative to Time.— We focus on the case in which compute is sufficiently abundant so that the binding constraint in human-AI interactions is human time, not compute. In other words, there is more compute available than the one demanded by tA and bA firms, so the leftover compute must be rented by single-layer automated firms. Hence, even though compute is scarce relative to production opportunities (as mentioned above), it is abundant relative to human time. In Section 8 of the Online Appendix, we also study the case where the amount of compute is small.

The following is a sufficient condition for compute to be abundant relative to time:¹²

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

To understand this condition, note that a firm will never hire a solver who is less knowledgeable than its workers. Consequently, the compute allocation that maximizes human-AI interactions subject to the aforementioned constraint involves (i) matching every human who is less knowledgeable than AI with an AI solver and (ii) matching every human who is more knowledgeable than AI with $n(z_{AI})$ units of compute using AI for production work. Since (i) uses $\int_0^{z_{AI}} n(z)^{-1} dG(z)$ units of compute, while (ii) uses $n(z_{AI})(1 - G(z_{AI}))$ units of compute, if condition (1) holds, then there are not enough humans to interact with all the available compute.

Some Notation.— For future reference, we define $W \equiv W_a \cup W_p$ and $S \equiv S_a \cup S_p$ as the overall set of human workers and solvers of the economy, respectively. We also denote by $e : S_p \rightarrow W_p$ the inverse of the matching function m . That is, e is the “employee matching function” denoting the knowledge of the human worker matched with a human solver with knowledge $s \in S_p$. This function always exists given that, as shown below, the equilibrium matching function is strictly increasing.

Finally, for any arbitrary set B , we denote by $\text{int}B$ and by $\text{cl}B$ the interior and closure of B , respectively. We also use $B \preceq B'$ to indicate that the set $B \subseteq [0, 1]$ “lies below” the set $B' \subseteq [0, 1]$. Formally, $B \preceq B'$ if $\sup B \leq \inf B'$. For example, $W_a \preceq W_p$ means that the best worker assisted by AI is weakly less knowledgeable than the worst worker assisted by a human.

¹²Note that for any distribution G and helping cost $h \in (0, 1)$, there exists a finite μ that satisfies this condition for all $z_{AI} \in [0, 1]$. The reason for this is that the left-hand side of (1) is continuous in $z_{AI} \in (0, 1)$ and is bounded as $z_{AI} \rightarrow 0$ and $z_{AI} \rightarrow 1$ (it converges to $1/h$ and $g(1)/h + \int_0^1 n(z)^{-1} dG(z)$, respectively).

2.2 Discussion of the Model

Before moving on to the analysis, we briefly comment on some assumptions underlying our model.

First, in contrast to human intelligence, AI can be used *at scale*. This manifests in the model in two distinct ways: Compute is abundant relative to human time, and AI can be leveraged across *all units* of compute (implying that all units of compute can solve problems up to the same difficulty). Our motivation for the first assumption is that compute has been growing exponentially over the past two centuries (Nordhaus, 2007). Our motivation for the second assumption is that digital information is nonrival and has a nearly zero marginal cost of reproduction (Brynjolfsson and McAfee, 2016; Goldfarb and Tucker, 2019).

Second, motivated by the idea that scalability is AI’s distinguishing feature, in our baseline setting, we take it to be the *only* difference between humans and AI. In particular, we assume that (i) AI can perfectly mimic the behavior of humans in all three possible roles in the economy (independent producer, worker, and solver), (ii) AI uses the same amount of compute irrespective of the difficulty of the problem it faces,¹³ and (iii) the communication cost of human-to-human interactions is the same as that of human-AI interactions. In Section 6, we relax (i) by considering the case where humans have some knowledge that AI cannot replicate on its own. Relaxing (ii) and (iii) is left for future research.

Third, in our baseline model we assume that there are more production opportunities than the economy’s capacity to pursue them. Hence, even though compute is abundant relative to time, it is still scarce relative to its potential uses. In Section 5, we study an extension where compute is not only abundant relative to time but also relative to production opportunities.

Fourth, our main goal in this paper is to analyze how AI affects human labor outcomes. For this reason, we take the AI technology and the economy’s compute as given and do not explicitly model the owners of compute or the developers of AI. In particular, this implies that compute is exclusively used for the deployment of AI systems rather than for their training.¹⁴ Studying firms’ incentives to develop AI or to increase the economy’s compute are intriguing avenues for future research.

Finally, we follow Antràs et al. (2006) and Fuchs et al. (2015) in assuming that the distribution of human knowledge is exogenous and that organizations have at most two layers. We opt for these assumptions primarily for the sake of simplicity, although we believe that they offer a good first

¹³This assumption is also in line with how current AI models operate. See Fridman, Lex. “Sam Altman: OpenAI, GPT-5, Sora, Board Saga, Elon Musk, Ilya, Power & AGI.” *The Lex Fridman Podcast #419*, March 18, 2024. <https://lexfridman.com/sam-altman-2-transcript> (accessed March 23, 2024).

¹⁴The industry separates AI’s use of compute between “training” (i.e., teaching AI systems how to respond) and “deployment” or “inferencing” (i.e., reacting to new bits of information). As of 2023, more than 40% of Nvidia’s data center business was dedicated to the deployment of AI systems, and that share is predicted to grow in the future. See Asa Fitch, “How a Shifting AI Chip Market Will Shape Nvidia’s Future,” *The Wall Street Journal*, February 25, 2024, <https://www.wsj.com/tech/ai/how-a-shifting-ai-chip-market-will-shape-nvidias-future-f0c256b1> (accessed February 26, 2024).

approximation of the problem at hand.

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 the amount of compute μ is zero and was originally described by Fuchs et al. (2015).¹⁵ Note that in this case $W_a = S_a = \emptyset$, so $W = W_p$ and $S = S_p$.

Proposition 1. *In the absence of AI, there is a unique equilibrium. The equilibrium has the following features:*

- *Occupational stratification: $W \preceq I \preceq S$.*
- *Positive assortative matching: The function $m : W \rightarrow S$ is strictly increasing.*
- *$W \neq \emptyset$ and $S \neq \emptyset$. However, $I \neq \emptyset$ if and only if $h > h_0 \in (0, 1)$.*

Moreover, the wage function w is continuous, strictly increasing, and convex (strictly so when $z \in W \cup S$), and is given by:

- *$w(z) = m(z) - w(m(z))/n(z)$ for all $z \in W$.*
- *$w(z) = z$ for all $z \in I$.*
- *$w(z) = C + \int_{\inf S}^z n(e(u))du > z$ for all $z \in S$, where $C > \inf S$ when $h < h_0$ (and $C = \inf S$ otherwise).*

In particular, $w(z) > z$ for all $z \notin \text{cl}I$ (so $w(z) > z$ for all $z \in [0, 1]$ when $h < h_0$).

Proof. See Appendix A. □

The equilibrium without AI—which we illustrate in Figure 3—has several salient features. First, it exhibits occupational stratification: Workers are less knowledgeable than independent producers, who are, in turn, less knowledgeable than solvers. Intuitively, more knowledgeable agents have a comparative advantage in specialized problem solving, as this allows them to leverage their knowledge by applying it to more than one problem. Hence, $(W \cup I) \preceq S$. Similarly, less knowledgeable agents have a comparative advantage in assisted rather than independent production work, as they are less likely to succeed on their own. Consequently, $W \preceq I$.

Second, there is strict positive assortative matching: Conditional on W and S , more knowledgeable workers in W match with more knowledgeable solvers in S . The reason is that worker and solver knowledge are complements: For a given team of workers, a more knowledgeable solver increases expected output, while for any given solver, more knowledgeable workers increase team size.

Third, the set of workers and the set of solvers are always nonempty, but the set of independent producers is empty when the communication cost h is below a threshold $h_0 \in (0, 1)$. Intuitively,

¹⁵Note that an economy with $\mu = 0$ is different than that with $\mu > 0$ but $z_{AI} = 0$. This is because, even if AI cannot solve any problems, it can still draw them, thus enlarging the production possibility frontier of the economy.




Figure 3: Illustration of the Pre-AI Equilibrium.

Notes. Distribution of knowledge: $G(z) = z$. Parameter values: For panel (a), $h = 1/2 (< h_0 = 3/4)$, while for panel (b), $h = 0.8125 (> h_0 = 3/4)$. The thick line depicts the equilibrium wage function. The dashed arrows illustrate the matching between workers and solvers. See Section 1 of the Online Appendix for a detailed characterization of the pre-AI equilibrium when human knowledge is uniformly distributed.

the lower the value of h is, the more attractive it is to form two-layer organizations compared to single-layer organizations.

Fourth, workers and solvers earn strictly more than their expected output as independent producers (except in the case of the most knowledgeable worker and the least knowledgeable solver when $h \geq h_0$). The reason for this is that the marginal value of knowledge is strictly lower than 1 for workers (as their knowledge is used to free up solver time),¹⁶ exactly equal to 1 for independent producers (as their expected output equals their knowledge), and strictly greater than 1 for solvers (as they can leverage their knowledge by applying it to more than one problem). Hence, if solvers were to earn their expected output as independent producers, then all two-layer firms would try to hire the most knowledgeable agents as solvers. Similarly, if workers were to earn their expected output as independent producers, then all two-layer firms would try to hire the least knowledgeable agents as workers.

Finally, the equilibrium wage function w is continuous, strictly increasing, and convex (strictly so when $z \in W \cup S$). The wages for the different occupations are obtained as follows. For independent

¹⁶A marginal increase in the knowledge of a worker with knowledge z liberates $h < 1$ units of her solver's time. This allows her firm to hire $1/(1-z)$ extra workers, with an expected net output gain of $(m(z) - w(z))/(1-z) < 1$ (as the wage $w(z)$ of that worker is strictly greater than z in the interior of W).

producers, $w(z) = z$ is an immediate implication of the zero-profit condition of single-layer firms. For workers and solvers, consider the problem of a two-layer organization that recruited $n(z)$ workers with knowledge $z \in W$ and is deciding which solver $s \in S$ to hire:

$$\max_{s \in S} \left\{ n(z)[s - w(z)] - w(s) \right\}$$

The corresponding first-order condition evaluated at $s = m(z)$ implies that $w'(m(z)) = n(z)$, or, equivalently, $w'(z) = n(e(z))$ for any $z \in S$. Thus $w(z) = C + \int_{\inf S}^z n(e(u))du$ for any $z \in S$, where the constant C is chosen so that the wage function is continuous, as the latter condition is necessary for market clearing. The wages of workers are then determined by the zero-profit condition of two-layer organizations: $w(z) = m(z) - w(m(z))/n(z)$.

For simplicity, in what follows, we restrict attention to $h < h_0$. This implies that there are no independent producers in the pre-AI equilibrium. In a previous version of this paper (Ide and Talamàs, 2024), we show that virtually all of our results extend to $h \geq h_0$.

3 The AI Equilibrium

Our objective is to understand the effects of AI by comparing the pre- and post-AI equilibrium. Toward this goal, in this section, we present a partial characterization of the post-AI equilibrium containing the essential information needed for our main results (which appear in Section 4). Appendix B provides the complete characterization of this equilibrium.

For future reference, we index the post-AI equilibrium using the superscript “*” (note that the pre-AI equilibrium has no superscript). Furthermore, recall that $W^* \equiv W_a^* \cup W_p^*$ is the overall set of human workers and that $S^* \equiv S_a^* \cup S_p^*$ is the overall set of human solvers.

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

- *Occupational stratification: $W^* \preceq I^* \preceq S^*$.*
- *No worker is better than AI; no solver is worse than AI: $W^* \preceq \{z_{\text{AI}}\} \preceq S^*$.*
- *Positive assortative matching: $m^* : W_p^* \rightarrow S_p^*$ is strictly increasing and $W_a^* \preceq W_p^*$ and $S_p^* \preceq S_a^*$.*

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

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

In any case, as long as $z_{\text{AI}} > 0$, AI does not lead to the complete destruction of routine human jobs; i.e., $W^ \neq \emptyset$. Finally, the rental rate of compute r^* is equal to z_{AI} , and the wage function w^* is continuous, strictly increasing, and convex (strictly so when $z \in W_p^* \cup S_p^*$), and is given by:*

- *$w^*(z) = z_{\text{AI}}(1 - 1/n(z)) > z$ for all $z \in W_a^*$.*

- $w^*(z) = m^*(z) - w^*(m^*(z))/n(z)$ for all $z \in W_p^*$.
- $w^*(z) = z$ for all $z \in I^*$.
- $w^*(z) = C^* + \int_{\inf S_p^*}^z n(e^*(u))du$ for all $z \in S_p^*$, where $C^* = \inf S_p^*$.
- $w^*(z) = n(z_{AI})(z - z_{AI}) > z$ for all $z \in S_a^*$.

In particular, $w^*(\sup W^*) = \sup W_p^*$, $w^*(z_{AI}) = z_{AI}$, and $w^*(\inf S^*) = \inf S_p^*$.

Proof. See Appendix B. □

Figure 4 illustrates the post-AI equilibrium in two different cases. Panel (a) depicts a situation of a relative “basic” AI that is used as a worker and an independent producer but not as a solver. Panel (b) depicts a relatively “advanced” AI that is used in all three possible roles. As both panels show, AI is the “best worker” in the economy and is therefore assisted by the most knowledgeable human solvers. In panel (b), AI is also the “worst solver” in the economy, and thus, it assists the production work of least knowledgeable human workers.

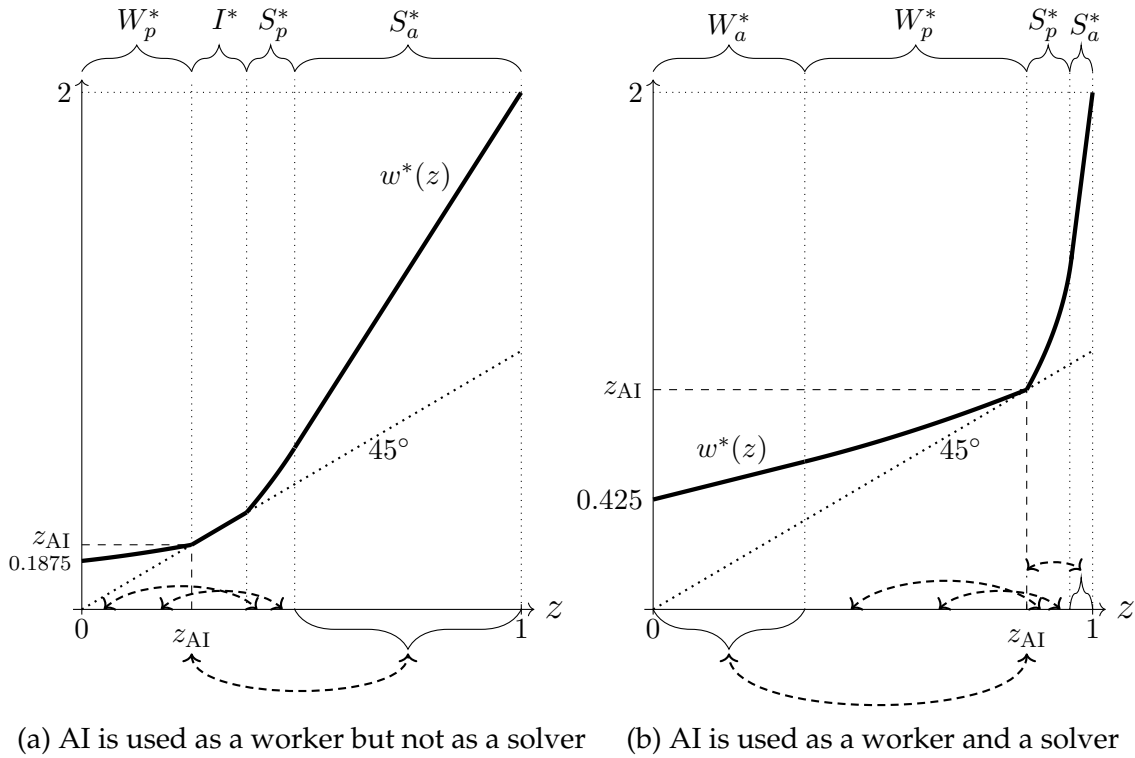


Figure 4: Illustration of the Post-AI Equilibrium for Two Different Values of z_{AI}

Notes. Distribution of knowledge: $G(z) = z$. Parameter values: Both panels have $h = 1/2$. For panel (a), $z_{AI} = 0.25$, while for panel (b), $z_{AI} = 0.85$. The thick line depicts the equilibrium wage function. The dashed arrows illustrate the matching between workers and solvers, both humans and AI. Human workers in W_p^* are endogenously matched with the human solvers in S_p^* according to m^* . All the humans in W_a^* are matched with an AI solver, while all humans in S_a^* are matched with AI workers. See Section 1 of the Online Appendix for a detailed characterization of the post-AI equilibrium when human knowledge is uniformly distributed.

Irrespective of the panel considered, the overall set of human workers $W^* \equiv W_a^* \cup W_p^*$ is comprised exclusively of those humans who are less knowledgeable than AI. Moreover, the human with knowledge z_{AI} earns exactly her output as an independent producer, as she is a perfect substitute for a unit of compute running with AI (whose price is $r^* = z_{AI}$).

For intuition, we now provide more details on Proposition 2. This proposition has three parts. The first part states the basic properties of the equilibrium. The second part describes how firms use humans and AI as a function of AI’s knowledge. Finally, the third part characterizes the equilibrium wages and the equilibrium price of compute.

Let us consider part one. First, the post-AI equilibrium continues to exhibit occupational stratification and positive assortative matching. This is because AI does not change the fact that (i) more knowledgeable agents have a comparative advantage in specialized problem solving, (ii) less knowledgeable agents have a comparative advantage in assisted rather than independent production work, and (iii) there are complementarities between worker and solver knowledge.

The remaining properties of the equilibrium follow from occupational stratification and positive assortative matching, in addition to the fact that AI can be used at scale. Indeed, because AI can be leveraged across all units of compute, and compute is abundant relative to time, a fraction of the available compute must be used for independent production. Hence, occupational stratification immediately implies that no worker can be better than AI and that no solver can be worse than AI (i.e., $W^* \preceq \{z_{AI}\} \preceq S^*$). Moreover, by positive assortative matching, if AI is used as a worker, then it is supervised by the most knowledgeable human solvers (i.e., $S_p^* \preceq S_a^*$) and if AI is used as a solver, then it assists the least knowledgeable human workers (i.e., $W_a^* \preceq W_p^*$).

The second part of the proposition—how firms use humans and AI as a function of z_{AI} —is driven by the differences in the comparative advantages of agents at different knowledge levels. First, recall that less knowledgeable agents have a comparative advantage in performing assisted production work, while more knowledgeable agents have a comparative advantage in specialized problem solving. Consequently, if AI has the knowledge of a pre-AI worker, then some compute receives assistance post-AI, as some humans who are more knowledgeable than AI were receiving assistance in the pre-AI equilibrium. Similarly, if AI has the knowledge of a pre-AI solver, then some compute is rented to provide assistance in the post-AI equilibrium, as humans who are less knowledgeable than AI provided assistance in the pre-AI equilibrium.

Given that pre-AI workers are always less knowledgeable than pre-AI solvers, and that some compute is always used for independent production, the previous results imply that when z_{AI} is relatively low, AI is necessarily used for assisted and independent production work. Likewise, when z_{AI} is relatively high, AI must be used for independent production and specialized problem solving. Determining whether in any of these circumstances, AI is simultaneously used in all three roles is subtle and not fundamental for our main results. For this reason, we relegate such details to Appendix B.

Agents’ comparative advantages also explain why AI does not lead to the complete destruction

of routine human jobs when $z_{AI} > 0$. Indeed, in the post-AI world, it continues to be worthwhile for every human to be employed in some capacity because compute is scarce relative to production opportunities (despite being abundant relative to human time). Hence, humans who are less knowledgeable than AI are hired as workers, as they are less likely than AI to succeed on their own, and it is, therefore, more valuable to provide them with assistance.

Finally, consider the third part of the proposition—equilibrium prices. The fact that $r^* = z_{AI}$ follows because some compute must be used for independent production and single-layer automated firms must obtain zero profits. The wages of those humans being hired by tA and bA firms, i.e., those in W_a^* and S_a^* , come from the zero-profit conditions of these firms in addition to the fact that $r^* = z_{AI}$:

$$\begin{aligned} n(z)[z_{AI} - w^*(z)] - r^* = 0 &\implies w^*(z) = z_{AI}(1 - 1/n(z)), \text{ for any } z \in W_a^* \\ n(z_{AI})[z - r^*] - w^*(z) = 0 &\implies w^*(z) = n(z_{AI})(z - z_{AI}), \text{ for any } z \in S_a^* \end{aligned}$$

The wages of the human workers and solvers hired by nA firms, i.e., those in W_p^* and S_p^* , are derived following a similar logic as in the pre-AI equilibrium. The only difference is that the best workers and worst solvers earn their expected output as independent producers (while they earn strictly more than that in the pre-AI equilibrium). The reason for this is that the equilibrium wage function must be continuous, and there is always some independent production in the post-AI equilibrium (done by AI and possibly also humans).

4 The Impact of AI

We now turn to our main endeavor: Analyzing the impact of AI by comparing the pre- and post-AI equilibrium. In essence, we compare the pre-AI equilibrium of Figure 3 with the post-AI equilibrium of Figure 4, as depicted in Figure 5. This comparison reveals substantial changes in (i) occupational choices, (ii) the matching between workers and solvers, and (iii) total labor income and its distribution. These changes also affect the distribution of firm size, productivity, and decentralization.

For the analysis, we use the following terminology:

- $z_{AI} \in \text{int}W$: “AI has the knowledge of a pre-AI worker.”
- $z_{AI} \in \text{int}S$: “AI has the knowledge of a pre-AI solver.”
- $z_{AI} \in W \cap S$: “AI has the knife-edge knowledge of both a pre-AI worker and a pre-AI solver.”

For brevity, we focus on the cases $z_{AI} \in \text{int}W$ and $z_{AI} \in \text{int}S$. The results for the knife-edge case $z_{AI} \in W \cap S$ are in Section 3 of the Online Appendix.




Figure 5: Comparison of the Pre- and Post-AI Equilibrium

Notes. Distribution of knowledge: $G(z) = z$. Parameter values: For both panels, $h = 1/2$ ($< h_0 = 3/4$). For panel (a), $z_{AI} = 0.25$, while for panel (b), $z_{AI} = 0.85$. In the post-AI equilibrium depicted in panel (a), AI is used as a worker and independent producer. In panel (b), AI is used in all three possible roles. See Section 1 of the Online Appendix for a detailed characterization of the pre-AI and post-AI equilibrium when human knowledge is uniformly distributed.

4.1 Occupational Displacement

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

Proposition 3.

- If $z_{AI} \in \text{int}W$, AI displaces humans from routine to specialized problem solving, i.e., $W^* \subset W$ and $S^* \supset S$.
- If $z_{AI} \in \text{int}S$, AI displaces humans from specialized to routine problem solving, i.e., $W^* \supset W$ and $S^* \subset S$.

Proof. See Appendix C.1. □

Note that the set of humans doing production work in either one- or two-layer organizations is the complement of the set of human solvers. Hence, when $z_{AI} \in \text{int}W$, AI decreases not only the number of humans doing assisted production work (i.e., workers) but also the number of humans doing any type of production work (i.e., workers and independent producers). Similarly, when $z_{AI} \in \text{int}S$, AI increases not only the number of humans doing assisted production work but also the number of humans doing any type of production work.

According to Proposition 3, the impact of AI on occupational choice 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 production work to specialized problem solving. In contrast, when AI has the knowledge of a pre-AI solver, the displacement goes in the opposite direction.

Intuitively, when AI has the knowledge of a pre-AI worker, it serves as a relatively inexpensive technology for routine production work, thus reducing workers' wages and increasing the attractiveness of creating two-layer firms. The result is a surge in the demand for solvers to match with the less expensive, more abundant "workers" (both humans and AI), which induces the most knowledgeable routine workers of the pre-AI equilibrium to switch to specialized problem solving.

Similarly, when AI has the knowledge of a pre-AI solver, it serves as a relatively inexpensive technology for specialized problem solving. This leads to the creation of additional two-layer organizations, increasing the demand for workers to match with the less expensive, more abundant "solvers" (both humans and AI). As a result, the least knowledgeable pre-AI solvers switch to routine production work.

Assuming that advanced economies have better communication technologies and/or a more knowledgeable population than developing economies, Proposition 3 leads to the following implication: The same AI technology may displace humans from routine to specialized problem solving in advanced economies but displace humans in the opposite direction in developing economies. The reverse situation, however, cannot occur: If AI displaces humans to specialized problem solving in developing economies, then the same must be true in advanced economies. This result follows because—as shown in Section 7 of the Online Appendix—the knowledge cutoff to become a solver in the pre-AI equilibrium is greater in advanced economies than in developing economies.

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

This section analyzes the effects of AI on the distribution of firm size, productivity, and span of control. We focus on two-layer organizations and begin with the following preliminary result:

Corollary 1. *AI necessarily increases the number (i.e., the measure) of two-layer firms.*

Proof. See Appendix C.2. □

We take *firm size* as its output, and *firm productivity* as its output divided by the units of time or compute used in production. Hence, a firm's productivity is equal to the level of knowledge of its solver. We define a firm's *span of control* as the amount of resources under its solver's supervision. Hence, the span of control of a firm that hires human workers with knowledge z is $n(z)$, while that of a firm that uses AI for production is $n(z_{AI})$. Note that span of control is a measure of decentralization because a solver can supervise more resources only if its workers solve more problems on their own.¹⁷

¹⁷We focus on this measure (rather than the number of *humans* under the supervision of a given solver) because the spirit

Proposition 4.

- If $z_{AI} \in \text{int}W$, then AI decreases the maximum span of control, the minimum firm productivity, and both the minimum and maximum firm size.
- If $z_{AI} \in \text{int}S$, then AI increases the maximum span of control, the minimum firm productivity, and both the minimum and maximum firm size.

Proof. See Appendix C.3. □

Proposition 4 is illustrated in Figure 6, which depicts the pre- and post-AI distributions of firm decentralization, productivity, and size for the two cases illustrated in Figure 5 (i.e., $z_{AI} = 0.25$ and $z_{AI} = 0.85$). The dark-shaded and light-shaded bars correspond to the firms created and destroyed by AI's introduction, respectively. Hence, the pre-AI distribution is equal to the white bars plus the light-shaded bars, while the post-AI distribution is equal to the white bars plus the dark-shaded bars.

As the figure shows, when $z_{AI} \in \text{int}W$, AI destroys the largest and most decentralized firms and leads to the creation of firms that are smaller and less productive than the smallest and least productive pre-AI firms. Similarly, when $z_{AI} \in \text{int}S$, AI destroys the smallest and least productive firms and leads to the creation of firms that are larger and more decentralized than the largest and most decentralized pre-AI firms.


For intuition, let us first consider productivity and span of control. Recall that when $z_{AI} \in \text{int}W$, AI induces the most knowledgeable pre-AI workers to switch to specialized problem-solving (Proposition 3). In terms of productivity, this implies that AI leads to the creation of less productive firms, as the newly created class of solvers is less knowledgeable than the pre-AI solvers. In terms of span of control, this displacement of humans across occupations implies that AI destroys the most decentralized firms, as the most knowledgeable pre-AI workers become solvers post-AI. The intuition for the case $z_{AI} \in \text{int}S$ is analogous.

Now, consider the effects of AI on the distribution of firm size. Note that positive assortative matching implies that more productive firms are also more decentralized. Moreover, the least knowledgeable workers always have knowledge $z = 0$, and the most knowledgeable solvers always have knowledge $z = 1$. Hence, the size of the smallest firm is $n(0)$ times the minimum firm productivity, while the size of the largest firm is 1 times the maximum span of control.


Consequently, when $z_{AI} \in \text{int}W$, AI reduces the minimum and maximum firm size, as it decreases the maximum span of control and the minimum firm productivity. Similarly, when $z_{AI} \in \text{int}S$, AI increases the minimum and maximum firm size, as it increases the maximum span of control and the minimum firm productivity.

We end this subsection by showing that when AI has the knowledge of a pre-AI solver, and this knowledge is sufficiently high, AI also leads to the creation of superstar firms that have “scale without


of our baseline model is that a unit of compute using AI is indistinguishable from a human with knowledge z_{AI} .




(a) Span of Control - $z_{AI} \in \text{int}W$




(b) Span of Control - $z_{AI} \in \text{int}S$




(c) Productivity - $z_{AI} \in \text{int}W$



(d) Productivity - $z_{AI} \in \text{int}S$



(e) Size - $z_{AI} \in \text{int}W$



(f) Size - $z_{AI} \in \text{int}S$

Figure 6: The Effects of AI on the Distribution of Firms' Span of Control, Productivity, and Size

Notes. This histogram is based on a human population of $N = 100 \times 10^6$ individuals. Distribution of knowledge: $G(z) = z$. For both panels, $h = 1/2$. Moreover, for panel (a), $z_{AI} = 0.25$, while for panel (b), $z_{AI} = 0.85$. Pre-AI, there are 23.6×10^6 firms. Post-AI, the number of firms increases to 62.5×10^6 in panel (a) and to 29.3×10^6 in panel (b). See Section 2 of the Online Appendix for the exact expressions of the pre- and post-AI distributions of firm size, productivity, and span of control.

mass.” This term was coined by Brynjolfsson et al. (2008) and Autor et al. (2020) to refer to those firms that attain large market shares with relatively few employees.

More precisely, we say that AI creates superstar firms with scale but no mass if the following two conditions are satisfied: (i) the largest post-AI firms are larger than the largest pre-AI firms, and (ii) the largest post-AI firms are bottom-automated (that is, they use a single human solver to supervise production work by AI).

Corollary 2. *There exists a $\zeta \in \text{int}S$ such that for all $z_{\text{AI}} \geq \zeta$, AI leads to the creation of superstar firms with scale but no mass.*

Proof. See Appendix C.4. □

Intuitively, AI increases the maximum firm size only when $z_{\text{AI}} \in \text{int}S$. In such a case, we know that AI is necessarily used as a solver and possibly also as a worker (Proposition 2). The condition that $z_{\text{AI}} \geq \zeta$ is necessary and sufficient for AI to be used in both roles, in which case the largest post-AI firms are bA firms (as AI is the best worker in the economy). These superstar firms with scale but no mass can, in fact, be seen in panel (f) of Figure 6.

4.3 The Effects of AI on Workers and Solvers who Are Not Occupationally Displaced

We now analyze the impact of AI on the productivity of pre-AI workers who remain workers post-AI (i.e., workers who are not occupationally displaced) and the span of control of pre-AI solvers who remain solvers post-AI (i.e., solvers who are not occupationally displaced).

In line with our definition of firm productivity, we define a worker’s productivity as her expected output per unit of time. Thus, it is equal to the knowledge of the solver with whom she is matched. Similarly, a solver’s span of control is equal to the resources under her supervision. Hence, her span of control is increasing in the knowledge of the workers with whom she is matched.

As the next proposition shows, AI affects not only those humans who directly interact with it but also those who do not interact with it. The reason for this is that when a subset of humans switches from matching with other humans to matching with AI, it induces a complete reorganization of all matches in the economy. For the proposition that follows, let us recall that $e(z)$ is the employee function of the pre-AI equilibrium.

Proposition 5.

- If $z_{\text{AI}} \in \text{int}W$, then:
 - The productivity of $z \in W^* \subset W$ is strictly smaller post-AI than pre-AI.
 - The span of control of $z \in S \subset S^*$ is strictly greater post-AI than pre-AI if $e(z) < z_{\text{AI}}$, and strictly smaller post-AI than pre-AI if $e(z) > z_{\text{AI}}$.
- If $z_{\text{AI}} \in \text{int}S$, then:

- The productivity of $z \in W \subset W^*$ is strictly greater post-AI than pre-AI if $z < e(z_{\text{AI}})$, and strictly smaller post-AI than pre-AI if $z > e(z_{\text{AI}})$.
- The span of control of $z \in S^* \subset S$ is strictly greater post-AI than pre-AI.

Proof. See Appendix C.5. □

Proposition 5 is illustrated in Figure 7 when $z_{\text{AI}} \in \text{int}W$ and AI is used as a worker and as an independent producer (but not as a solver). As the figure shows, AI worsens the matches of all pre-AI workers who remain workers (i.e., those humans in W_p^*), reducing their productivity. Moreover, AI improves the match of the least knowledgeable solvers who are solvers both pre- and post-AI—increasing their span of control—but worsens the match of the most knowledgeable solvers who are solvers both pre- and post-AI.

For intuition, suppose first that AI has the knowledge of a pre-AI worker. In this case, (i) AI is the best worker in the post-AI economy and is therefore assisted by the best solvers (Proposition 2), and (ii) the knowledge of the worst solvers decreases with AI’s introduction (Proposition 3). Both

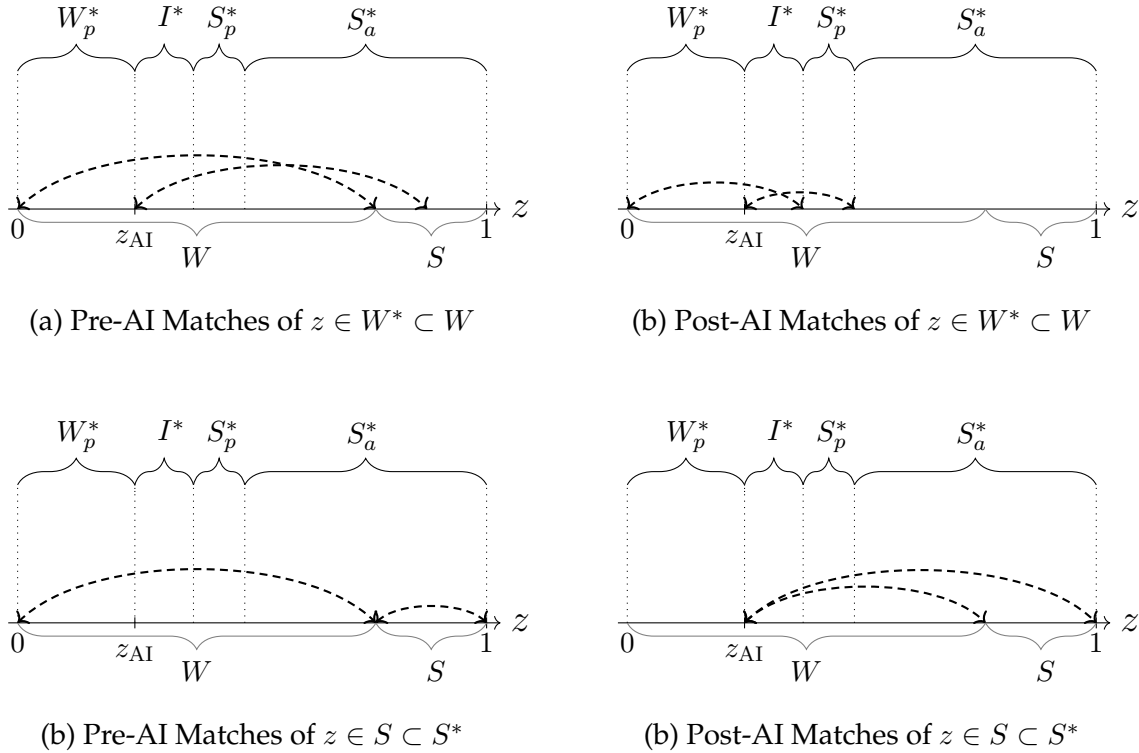


Figure 7: An Illustration of Proposition 5 - $z_{\text{AI}} \in \text{int}W$

Notes. Distribution of knowledge: $G(z) = z$. Parameter values: For all panels, $h = 1/2$ and $z_{\text{AI}} = 0.25$. In the post-AI equilibrium, AI is used as a worker and independent producer. The dashed arrows illustrate the matching between workers and solvers. See Section 1 of the Online Appendix for a detailed characterization of the pre-AI and post-AI equilibrium when human knowledge is uniformly distributed.

effects worsen the pool of solvers available for those pre-AI workers who remain workers post-AI, thus reducing their productivity.

Regarding the solvers who are not occupationally displaced, the match of the best solvers worsens with AI's introduction because post-AI, they assist production by AI, while pre-AI, they assist humans who are more knowledgeable than AI. However, the match of the worst solvers improves with AI's introduction because post-AI, the least knowledgeable workers are assisted by the newly appointed solvers (humans and possibly AI), who are less knowledgeable than the worst pre-AI solvers.

Similarly, when AI has the knowledge of a pre-AI solver, (i) AI is the worst solver in the post-AI economy and therefore assists the least knowledgeable workers (Proposition 2), and (ii) the knowledge of the best workers increases with AI's introduction (Proposition 3). Both effects improve the pool of workers available for pre-AI solvers who remain solvers post-AI, thus increasing their span of control.

Regarding the workers who are not occupationally displaced, the match of the worst workers improves with AI's introduction because post-AI, they are assisted by AI, while pre-AI, they are assisted by humans who are less knowledgeable than AI. However, the match of the best workers worsens with AI's introduction because post-AI, the best solvers assist the production work of the newly appointed workers (humans and possibly AI) who are more knowledgeable than the best pre-AI workers.

Together, Propositions 4 and 5 show that even when AI destroys the most decentralized firms (i.e., when $z_{AI} \in \text{int}W$), a fraction of the non-occupationally displaced solvers are reallocated to more decentralized firms. Similarly, even when AI destroys the least productive firms (i.e., when $z_{AI} \in \text{int}S$), a fraction of the non-occupationally displaced workers are reallocated to less productive firms.

4.4 Labor Income

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

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

The proof of this result is intuitive and relatively straightforward, so we provide it here as part of the main text. Moreover, this proof is instructive because it shows that not all gains from AI accrue to labor.

Proof. The fact that AI increases total output follows because AI expands the production possibility frontier, and the First Welfare Theorem holds in this setting. The fact that AI increases total labor income, in turn, follows from two observations. First, if all compute is assigned to independent production, then AI does not affect total labor income (as it does not interact with humans in the

workplace). Second, capital income is equal to μz_{AI} regardless of how compute is used. Hence:

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

Consequently, given that the post-AI equilibrium is efficient, unique, and does not allocate all compute to independent production, the following must hold:

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

Hence, total labor income must be strictly greater post-AI than pre-AI. \square

We now turn to analyzing the impact of AI on the distribution of labor income. Given that this is a competitive economy, the wage of an individual is her marginal product (defined as the output increase from introducing such an agent into the economy). Hence, understanding whose wage increases or decreases with the introduction of AI, i.e., who are the “winners” and “losers” from AI, amounts to understanding which humans are complemented by the technology (in the sense that AI increases their marginal product) and which humans are substituted by the technology (in the sense that AI decreases their marginal product). Note that an agent’s marginal product is not necessarily equal to her productivity (as defined in Section 4.3) because an agent’s introduction affects the output of other agents by changing worker-solver matches.

Disentangling the distributional effects of AI is not 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 (as shown by Proposition 5). On the other hand, by mimicking humans with knowledge z_{AI} , AI changes the relative scarcities of different knowledge levels, affecting how each firm’s output is divided between workers and solvers.

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

Lemma 2. Define $\Delta(z) \equiv w^*(z) - w(z)$. Then $\Delta(z_{AI}) < 0$ 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 $\Delta(z') > 0$ for all $z' \in [z, 1]$.

Proof. See Appendix C.6. \square

Next, we characterize when there are winners below or above z_{AI} :

Proposition 6.

- (i) *There always exists a set $[z, 1] \subseteq (z_{AI}, 1]$ of winners.¹⁸*
- (ii) *There exists a set $[0, z] \subseteq [0, z_{AI}]$ of winners if and only if $z_{AI} > \bar{z}_{AI}$, where $\bar{z}_{AI} \in \text{int}W$.*

Proof. See Appendix C.7. □

Figure 5 at the beginning of this section provides an illustration of Proposition 6. As shown in panel (a) of the figure, only the most knowledgeable humans win when z_{AI} is relatively low (more precisely, all $z > 0.48$ are winners). In contrast, as shown in panel (b), both the least and the most knowledgeable humans win when z_{AI} is relatively high (more precisely, all $z < 0.63$ and all $z > 0.95$ are winners in this case).

We now provide intuition for this proposition. In light of Lemma 2, to verify the existence of winners below or above z_{AI} , it is sufficient to analyze the effects of AI on the wages of the least knowledgeable humans (i.e., $z = 0$) and the most knowledgeable humans (i.e., $z = 1$).

With this in mind, let us consider part (i). In both the pre- and post-AI equilibrium, the most knowledgeable humans are always solvers. Moreover, their wages are given by the following:

$$\begin{aligned} \text{Wage of the solver} &= \# \text{ of problems solved by the solver} \\ &+ \# \text{ of problems independently solved by workers} - \text{worker wage bill} \end{aligned}$$

Hence, the result that the most knowledgeable humans are always winners follows because (i) the number of problems they solve is the same pre- and post-AI (it is equal to $1/h$) and (ii) the number of problems independently solved by the workers of those solvers minus their wage bill is strictly negative pre-AI but equal to 0 post-AI. This last result follows because the most knowledgeable workers (with whom they match) now earn exactly their expected output as independent producers (see Proposition 2), while pre-AI, these workers earned more (as shown in Proposition 1).

For part (ii), let $\alpha(z)$ and $\alpha^*(z)$ be the pre- and post-AI shares of output appropriated by the workers of a firm that hires workers of knowledge z ; i.e.,

$$\alpha(z) = \frac{\overbrace{n(z)w(z)}^{\text{Wage Bill}}}{\underbrace{n(z)m(z)}_{\text{Firm Output}}} = \frac{w(z)}{m(z)} \quad \text{and} \quad \alpha^*(z) = \frac{n(z)w^*(z)}{n(z)m^*(z)} = \frac{w^*(z)}{m^*(z)}$$

Consequently, $\Delta(0) = \alpha^*(0)m^*(0) - \alpha(0)m(0)$ since the least knowledgeable individuals are workers both pre- and post-AI.

¹⁸In a previous version of this paper (Ide and Talamàs, 2024) we show that this result relies on $h < h_0$. If $h \geq h_0$, in contrast, then all humans in $z \in [z_{AI}, 1]$ can be losers from AI when z_{AI} is sufficiently high (the humans with $z = 1$ are always indifferent in this case).

Consider then the impact of AI on $\Delta(0)$ when $z_{AI} \in \text{int}W$. This technology has two opposing effects on the wages of the least knowledgeable individuals. On the one hand, by Proposition 5, AI decreases the productivity of these individuals as it worsens their match, i.e., $m^*(0) < m(0)$. On the other hand, AI increases the share of firm output that they appropriate; i.e., $\alpha^*(0) > \alpha(0)$.¹⁹ This latter result follows because the least knowledgeable solvers (with whom they match) now earn exactly their expected output as independent producers (Proposition 2), while pre-AI, they earned strictly more (Proposition 1). We conclude that $\Delta(0)$ is strictly positive when the second effect dominates the first effect, which occurs when $z_{AI} > \bar{z}_{AI}$, where $\bar{z}_{AI} \in W$.²⁰

4.5 Discussion: AI versus Offshoring

We end this section by arguing that the effects of AI are different than those of offshoring. These differences are driven by AI's capacity to operate at scale (in the two senses discussed in Section 2.2).

The canonical model of offshoring in a knowledge economy is due to Antràs et al. (2006), who consider a two-country model in which countries differ only in their knowledge distributions. In particular, one country, the North, has a distribution of knowledge with a relatively high mean, while the other country, the South, has a distribution of knowledge with a relatively low mean. They show that allowing the formation of international teams shifts humans from routine work to specialized problem solving in the North, while it shifts humans in the opposite direction in the South.

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

As we mentioned above, the key difference between AI and offshoring is the capacity of AI to operate at scale. This implies that, although both offshoring and AI induce the best northern workers pre-offshoring/pre-AI to switch to specialized problem solving, the best northern solvers switch to supervising the best non-occupationally displaced northern workers in the case of offshoring, while they switch to supervising AI in the case of artificial intelligence.

Similarly, when $z_{AI} \in \text{int}S$, both offshoring and AI induce the worst southern solvers to switch to routine work, improving the overall pool of southern workers. However, in the case of offshoring,

¹⁹Formally, $\alpha^*(0) = 1 - h > 1 - hw(\inf S)/\inf S = \alpha(0)$, as $w(\inf S) > \inf S$ (see Proposition 1).

²⁰Note that if $z_{AI} \in \text{int}S$, then humans with knowledge $z = 0$ always win from AI because both effects go in the same direction: AI increases the share of output and improves the match of the least knowledgeable humans.

the best southern workers are matched with the best northern solvers, leaving a worse pool of workers for the non-occupationally displaced southern solvers. In contrast, in the case of AI, the worst workers end up being supervised by AI, leaving the best workers for the solvers who are not occupationally displaced.

5 Extension I: Compute Abundant Relative to Production Opportunities

In our baseline setting, compute is abundant relative to time but scarce relative to the existing production opportunities. This has two noteworthy implications. First, the equilibrium rental rate of compute is equal to AI’s knowledge. Second, all humans earn labor income in the post-AI equilibrium—even those who are less knowledgeable than AI.

In this section, we relax the assumption that compute is scarce relative to production opportunities. In this case, the equilibrium price of compute is zero, and AI leads to technological unemployment: All humans who are less knowledgeable than AI become unemployed. Organizations, however, still display a hierarchical structure in that the most knowledgeable humans specialize in tackling problems that AI cannot solve.

5.1 The Model

The model is exactly as described in Section 2, except that the number of production opportunities is smaller than the compute available to pursue them (but larger than such capacity pre-AI). More precisely, denoting by ϕ the number of production opportunities, we assume that $1 < \phi < \mu$.²¹

As in the baseline model, pursuing each opportunity requires one unit of time or compute and solving a problem whose difficulty is not known ex-ante. To produce, firms must purchase production opportunities—which, given their scarcity, will now command a strictly positive price—and hire the necessary labor or rent the necessary compute to pursue them. Moreover, just as with the owners of compute, we do not explicitly model the “entrepreneurs” who own these opportunities.

Wages, Prices, and Profits.—Let p be the price of production opportunities (note that all opportunities must have the same price since they are all ex-ante identical), and recall that $w(z)$ and r denote the wage of a human with knowledge z and the rental rate of compute, respectively. The profit of a single-layer organization is as follows:

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

²¹To stay as close as possible to the baseline model, we assume that the number of opportunities ϕ is exogenous. In particular, AI cannot be used to generate new production opportunities.

In other words, the profit of a single-layer nonautomated firm is the expected output z of its independent producer net of her wage $w(z)$ and the price of a production opportunity p . The profit of a single-layer automated firm is analogous.

The profit of a two-layer organization as a function of its configuration (tA , bA , or nA) is as follows:

$$\begin{aligned}\Pi_2^{tA}(z) &= n(z)[z_{AI} - w(z) - p] - r && \text{(where } z \leq z_{AI}\text{)} \\ \Pi_2^{bA}(s) &= n(z_{AI})[s - r - p] - w(s) && \text{(where } z_{AI} \leq s\text{)} \\ \Pi_2^{nA}(s, z) &= n(z)[s - w(z) - p] - w(s) && \text{(where } z \leq s\text{)}\end{aligned}$$

where z and s denote the knowledge of a human worker and a human solver, respectively, and we are using the fact that no two-layer firm hires a solver who is less knowledgeable than its workers.

As in the baseline case, the profit of a two-layer organization is its expected output minus the cost of its resources. For instance, in the case of a tA firm that hires workers with knowledge z , its expected output is $n(z)z_{AI}$, while the cost of resources is $n(z)(w(z) + p) + r$, as the firm hires $n(z)$ workers, rents one unit of compute, and purchases $n(z)$ production opportunities.

Competitive Equilibrium.— Define the sets (W_p, W_a, I, S_p, S_a) , the masses μ_i, μ_w , and μ_s , and the matching function $m : W_p \rightarrow S_p$ as in the baseline model. We denote by U the set of humans that are unemployed and by μ_u the mass of compute not being used. Note that the total mass of opportunities pursued is equal to $\mu_w + \mu_i + \int_{z \in (W \cup I)} dG(u)$.

Definition (Competitive Equilibrium). An equilibrium consists of nonnegative masses $(\mu_i, \mu_w, \mu_s, \mu_u)$, sets $(W_p, W_a, I, S_p, S_a, U)$, a matching function $m : W_p \rightarrow S_p$, a wage schedule $w : [0, 1] \rightarrow \mathbb{R}_{\geq 0}$, a rental rate of compute $r \in \mathbb{R}_{\geq 0}$, and a price of production opportunities $p \in \mathbb{R}_{\geq 0}$ such that:

1. Firms optimally choose their structure (while earning zero profits).
2. tA firms hiring $n(z)$ workers with knowledge $z \in W_a$ rent one unit of compute and purchase $n(z)$ production opportunities.
3. bA firms hiring a solver with knowledge $s \in S_a$ rent $n(z_{AI})$ units of compute and purchase $n(z_{AI})$ production opportunities.
4. nA firms hiring $n(z)$ workers with knowledge $z \in W_p$ hire a solver with knowledge $m(z) \in S_p$ and purchase $n(z)$ production opportunities.
5. Markets clear: (i) $\mu_i + \mu_w + \mu_s + \mu_u = \mu$, (ii) the union of the sets $(W_p, W_a, I, S_p, S_a, U)$ is $[0, 1]$ and the intersection of any two of these sets has measure zero, and (iii) $\mu_w + \mu_i + \int_{z \in (W \cup I)} dG(u) = \phi$.

5.2 The Impact of AI When Compute is Abundant Relative to Production Opportunities

The pre-AI equilibrium is the same as in the baseline, as the number of opportunities is greater than the economy's pre-AI capacity to pursue them (i.e., $\phi > 1$). The post-AI equilibrium, in turn, is characterized by the next proposition. We index this new post-AI equilibrium with the superscript " \star " instead of the superscript " $*$ " to distinguish it from that of the baseline.

Proposition 7. *In the presence of AI, there is a unique equilibrium. The equilibrium allocations are:*

$$U^* = [0, z_{AI}), S_a^* = [z_{AI}, 1], W_a^* = W_p^* = I^* = S_p^* = \emptyset$$

$$\mu_w^* = n(z_{AI})[1 - G(z_{AI})], \mu_s^* = 0, \mu_i^* = \max\{0, \phi - \mu_w^*\}, \mu_u^* = \mu - \mu_w^* - \mu_i^*$$

The equilibrium prices are $p^* = z_{AI}$, $r^* = 0$, and:

$$w^*(z) = \begin{cases} 0 & \text{if } z \in U^* \\ n(z_{AI})(z - z_{AI}) & \text{if } z \in S_a^* \end{cases}$$

Proof. See Section 5 of the Online Appendix. □

According to Proposition 7, which is illustrated in Figure 8, the equilibrium when compute exceeds the available production opportunities is significantly different from that of the baseline. In particular, the price of compute is now zero, and all agents who are less knowledgeable than AI are unemployed. Moreover, all those humans who are more knowledgeable than AI become solvers in bottom-automated firms.

Intuitively, given that compute is more abundant than production opportunities, firms use AI to pursue all available production opportunities. This displaces all those humans who are less knowl-




Figure 8: The Effects of AI when Compute is Greater than Production Opportunities

Notes. Distribution of knowledge: $G(z) = z$. Parameter values: For both panels, $h = 1/2$. For panel (a), $z_{AI} = 0.25$, while for panel (b), $z_{AI} = 0.85$. The thick gray line depicts the pre-AI equilibrium wage function w . The thick black line depicts the post-AI equilibrium wage function w^* .

edgeable than AI to unemployment. Furthermore, since some compute is left idle, the rental rate of compute is zero. Knowledge beyond that of AI, however, is still scarce because AI can attempt to solve but not actually solve all problems. Consequently, organizations are still hierarchical: All problems are initially attempted by AI, and humans who are more knowledgeable than AI specialize in solving problems that AI cannot solve. These humans are the only humans who are rewarded for their work.

The scarcity of opportunities also leads to a different distribution of income compared to the baseline. In particular, while in our baseline setting, “capitalists” (i.e., compute owners) earn z_{AI} per unit of compute and “entrepreneurs” (i.e., the owners of production opportunities) earn 0, in this extension, capitalists earn 0, and entrepreneurs earn z_{AI} per opportunity. Moreover, AI need not increase total labor income relative to the pre-AI benchmark, as a significant fraction of the population is displaced toward unemployment. The combination of these results suggests that as compute becomes abundant, there are incentives to reallocate labor and compute from pursuing production opportunities toward generating them.

Nevertheless, despite these differences, many of the results of our baseline continue to hold in this setting. Indeed, if $z_{AI} \in \text{int}W$, then AI still displaces humans from routine production work to specialized problem-solving, while if $z_{AI} \in \text{int}S$, then AI still reduces the number of humans doing specialized problem-solving (although the displacement, in this case, is toward unemployment). This immediately implies that our results regarding the distribution of firm productivity and span of control continue to hold without changes. The same applies to the effects of AI on the span of control of those solvers who are not occupationally displaced.

6 Extension II: Nonautonomous AI

In our baseline setting, AI can perfectly mimic the behavior of humans with knowledge z_{AI} . This implies, in particular, that AI is “autonomous” in that it can engage in production without any assistance from humans. While our primary focus is to explore the impact of AI when it is indistinguishable from human intelligence, it is instructive to also analyze the case in which AI lacks human-like intelligence and cannot operate autonomously. This is what we do in this extension.

More concretely, we assume that knowledge has two dimensions. The first dimension corresponds to the knowledge considered in the baseline setting, while the second dimension corresponds to human knowledge that AI cannot mimic. Since solving a problem always requires some knowledge in both dimensions, AI always needs human assistance to produce in this case.

As shown below, AI’s lack of autonomy—together with the assumption that compute is abundant relative to time—implies that the equilibrium price of compute is zero and that all individuals specialize in assisting AI in the dimension in which it lacks human-like intelligence. Moreover, there is no technological unemployment, and all the benefits of AI accrue to labor. However, the introduc-

tion of AI may still generate losers when its knowledge is limited. Only when technology reaches a certain level of proficiency do wages across the entire knowledge distribution increase.

6.1 The Model

The model is exactly as described in Section 2, except that now knowledge and problems are multidimensional. In particular, an agent with knowledge $\mathbf{z} = (z_1, z_2) \in [0, 1]^2$ can solve a problem of difficulty $\mathbf{x} = (x_1, x_2) \in [0, 1]^2$ if her knowledge exceeds the problem's difficulty, i.e., if $\mathbf{z} \geq \mathbf{x}$.²² Each problem is ex-ante identical and its difficulty \mathbf{x} is distributed uniformly on $[0, 1]^2$ independently across problems.

We continue to assume that knowledge can be communicated at a cost. When a worker with knowledge $\mathbf{z} = (z_1, z_2)$ cannot solve a problem of difficulty \mathbf{x} and is matched with a solver of knowledge $\mathbf{s} = (s_1, s_2)$, she can ask the solver for help. If the knowledge of the solver combined with the knowledge of the worker is greater than \mathbf{x} , i.e., if $\max(z_1, s_1) \geq x_1$ and $\max(z_2, s_2) \geq x_2$, the worker produces one unit of output. Otherwise, no production takes place. In any case, this exchange consumes $h \in (0, 1)$ units of the solver's time, irrespective of the number of dimensions in which the worker needs help.

Note that, just as in the baseline setting, no two-layer organization has a solver who is less knowledgeable than its workers. In particular, letting \mathbf{z} and \mathbf{s} denote the knowledge of the workers and the solver of a two-layer organization, respectively, it is not possible that $\mathbf{z} \geq \mathbf{s}$. However, there can be two-layer organizations in which a solver may be more knowledgeable than her workers in one of the dimensions but less knowledgeable in the other dimension.

To stay as close as possible to the baseline model, we focus on the case in which humans are heterogeneous in terms of z_1 but homogenous in terms of z_2 . More precisely, while human knowledge z_1 is distributed according to an atomless probability distribution with full support on $[0, 1]$, all humans have $z_2 = 1$. This implies that humans can be identified by their knowledge z_1 and that the pre-AI equilibrium is identical to that of the baseline described in Proposition 1.

Our departure from the baseline setting is that AI is not fully comparable to human intelligence. To be precise, AI's knowledge is $\mathbf{z}_{\text{AI}} = (z_{\text{AI}}, 0)$, where $z_{\text{AI}} \in [0, 1]$.²³ Hence, AI is not "autonomous" in that it always needs human assistance to solve any given problem. This also implies that AI is never more knowledgeable than any human: Either a human is more knowledgeable than AI, or their knowledge is not comparable. We assume that compute is abundant relative to human time but not relative to production opportunities.

Wages, Prices, and Profits.—Let $w(z_1)$ be the wage of a human with knowledge z_1 and denote by r the

²²For any two vectors \mathbf{z} and \mathbf{z}' we write $\mathbf{z} \geq \mathbf{z}'$ to mean that both $z_1 \geq z'_1$ and $z_2 \geq z'_2$.

²³In contrast, our baseline setting corresponds to the case in which $\mathbf{z}_{\text{AI}} = (z_{\text{AI}}, 1)$.

rental rate of one unit of compute. The profit of a single-layer organization is as follows:

$$\Pi_1 = \begin{cases} z_1 - w(z_1) & \text{if the firm hires a human with knowledge } z_1 \\ -r & \text{if the firm uses AI} \end{cases}$$

That is, the profit of a single-layer nonautomated firm is the expected output z_1 of its independent producer net of her wage $w(z_1)$. The profit of a single-layer automated firm is $-r$ since an AI executing independent production work is unable to solve any problems.

Obtaining the profit of a two-layer organization is more involved. First, recall that $n(x) \equiv [h \times (1 - x)]^{-1}$ is the maximum number of workers that a given solver can assist when each worker asks a question with probability $1 - x$. This implies that tA and nA firms hire exactly $n(z_1)$ workers of knowledge z_1 (as human workers require assistance only when $x_1 > z_1$), while bA firms hire exactly $n(0)$ units of compute for production work (as the AI algorithm always requires assistance).

Second, given that no two-layer firm hires a solver who is less knowledgeable than its workers, then (i) nA firms necessarily hire solvers who are more knowledgeable than their workers in the first dimension (i.e., $s_1 > z_1$), and (ii) tA firms necessarily hire humans who are less knowledgeable than AI in the first dimension (i.e., $z_1 < z_{AI}$). However, all humans can potentially be solvers in bA firms because AI is never more knowledgeable than any human in both dimensions. This implies that both $s_1 \geq z_{AI}$ and $s_1 \leq z_{AI}$ are feasible when considering bA firms.

Consequently, the profit of a two-layer organization as a function of its configuration (tA , bA , or nA) is as follows:

$$\begin{aligned} \Pi_2^{tA}(z_1) &= n(z_1)[z_{AI} - w(z_1)] - r && \text{(where } z_1 \leq z_{AI}) \\ \Pi_2^{bA}(s_1) &= n(0)[\max\{s_1, z_{AI}\} - r] - w(s_1) \\ \Pi_2^{nA}(s_1, z_1) &= n(z_1)[s_1 - w(z_1)] - w(s_1) && \text{(where } z_1 \leq s_1) \end{aligned}$$

where z_1 and s_1 denote the knowledge of a human worker and a human solver, respectively, in the first dimension.

Competitive Equilibrium.— Define the sets (W_p, W_a, I, S_p, S_a) , the masses μ_i, μ_w , and μ_s , and the matching function $m : W_p \rightarrow S_p$ as in the baseline model (although, in this case, the sets refer to the first dimension of knowledge).

Definition (Competitive Equilibrium). An equilibrium consists of nonnegative masses (μ_i, μ_w, μ_s) , sets (W_p, W_a, I, S_p, S_a) , a matching function $m : W_p \rightarrow S_p$, a wage schedule $w : [0, 1] \rightarrow \mathbb{R}_{\geq 0}$, and a rental rate of compute $r \in \mathbb{R}_{\geq 0}$ such that:

1. Firms optimally choose their structure (while earning zero profits).
2. tA firms hiring $n(z_1)$ workers with knowledge $z_1 \in W_a$ rent one unit of compute.
3. bA firms hiring a solver with knowledge $s_1 \in S_a$ rent $n(0)$ units of compute.
4. nA firms hiring $n(z_1)$ workers with knowledge $z_1 \in W_p$ hire a solver with knowledge $m(z_1) \in S_p$.

5. Markets clear: (i) $\mu_i + \mu_w + \mu_s = \mu$, and (ii) the union of the sets (W_p, W_a, I, S_p, S_a) is $[0, 1]$ and the intersection of any two of these sets has measure zero.

6.2 The Effects of a Nonautonomous AI

The pre-AI equilibrium is the same as that in the baseline. The post-AI equilibrium, in turn, is characterized by the next proposition. We index the post-AI equilibrium with the superscript “**” to distinguish it from those of Sections 3 and 5.

Proposition 8. *In the presence of AI, there is a unique equilibrium. The equilibrium allocations are:*

$$S_a^{**} = [0, 1], W_a^{**} = W_p^{**} = I^{**} = S_p^{**} = \emptyset$$

$$\mu_w^{**} = n(0), \mu_s^{**} = 0, \mu_i^{**} = \mu - \mu_w^{**}$$

*The equilibrium prices are $r^{**} = 0$ and $w^{**}(z_1) = n(0) \max\{z_1, z_{AI}\}$ for any $z_1 \in S_a^{**}$.*

Proof. See Section 6 of the Online Appendix. □

According to Proposition 8, which is illustrated in Figure 9, the equilibrium when AI lacks autonomy is closer to that of Proposition 7 (where compute is more abundant than production opportunities) than to that of the baseline setting. In particular, in both Propositions 7 and 8, the equilibrium price of compute is zero, and humans specialize in solving the problems that AI cannot solve. However, there is a crucial difference between the two outcomes: In Proposition 7, AI results in unemployment, whereas in Proposition 8, everyone remains employed following AI’s introduction.

Intuitively, in both cases, the equilibrium price of compute drops to zero due to an abundance of compute relative to its applications. However, in Proposition 7, this situation arises from compute being more abundant than production opportunities, whereas in Proposition 8, this situation is due to AI lacking autonomy. Consequently, in the latter case, there remains demand for labor post-AI, as humans are crucial for assisting AI where it lacks human-like intelligence. This result implies that all income flows to labor, as compute is free and production opportunities are plentiful.

Moreover, AI’s lack of autonomy also implies that the least knowledgeable humans leverage AI to not only perform production work but also solve problems of greater difficulty than those they can handle on their own. Indeed, as shown in the statement of Proposition 8, those humans with $z_1 \leq z_{AI}$ earn a wage of $z_{AI}n(0)$, as they rent $n(0)$ units of compute and benefit from AI solving problems of z_{AI} in the first dimension while providing AI assistance in the second dimension. In contrast, humans with $z_1 > z_{AI}$ earn $n(0)z_1$, as they assist AI across both dimensions.

Nevertheless, AI does not always result in benefits across the entire wage distribution. Specifically, individuals with a low z_1 may still fare worse than in the pre-AI equilibrium when z_{AI} is sufficiently low. This is because when AI’s knowledge is limited, it slightly increases the difficulty of problems that individuals with limited knowledge can handle while entirely erasing higher-paying worker positions (as the most knowledgeable individuals switch to using AI for production at no cost).

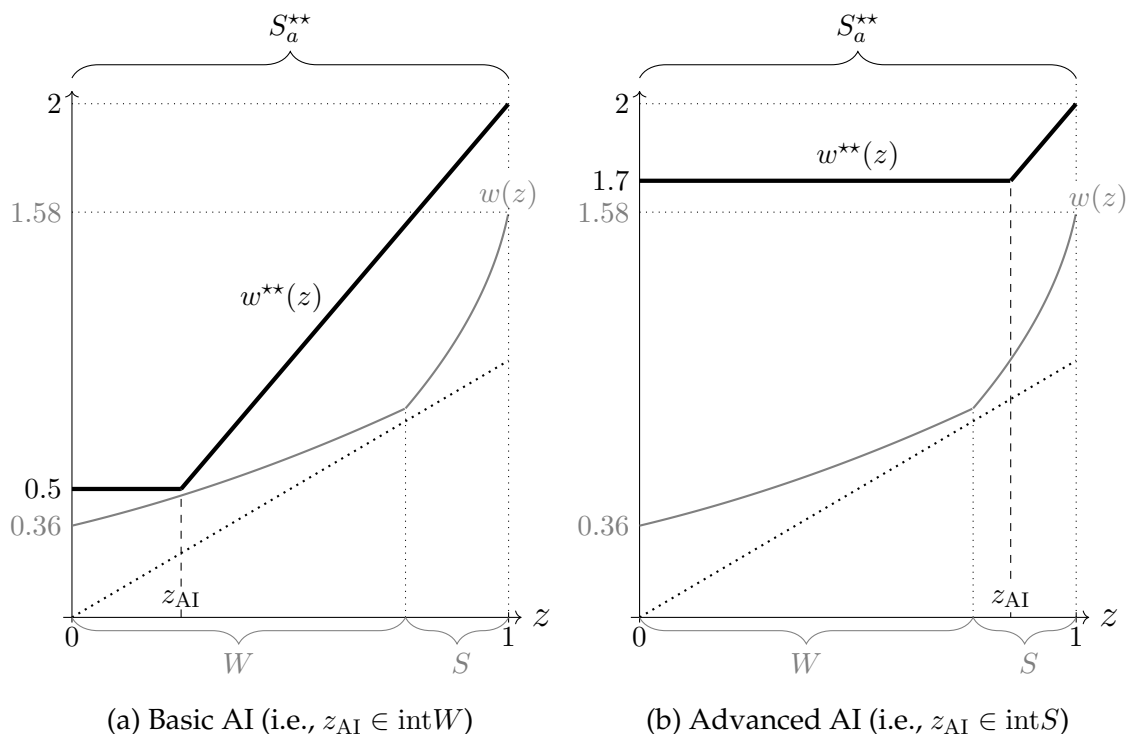


Figure 9: The Effects of a Nonautonomous AI

Notes. Distribution of knowledge: $G(z) = z$. Parameter values: For both panels, $h = 1/2$. For panel (a), $z_{AI} = 0.25$, while for panel (b), $z_{AI} = 0.85$. The thick gray line depicts the pre-AI equilibrium wage function w . The thick black line depicts the post-AI equilibrium wage function w^{**} .

7 Final Remarks

The potential impact of AI is undeniable, yet its precise implications for the future of work are controversial. This paper introduces a new framework for examining AI's equilibrium impact on organizational and labor outcomes. The novelty of this framework is that it explicitly incorporates the peculiarities of AI and knowledge work, enabling us to provide predictions about the potential effects of AI as a function of its capabilities and other parameters of the economy. Our goal is to enhance understanding of AI's economic impact, which is essential for designing policies to prepare for these changes.

Four broad lessons emerge from our analysis. First, AI's comparative advantage is shaped by its knowledge and autonomy as well as the availability of compute. In particular, even if AI is comparable to human intelligence and has an absolute advantage over a fraction of the population, it does not necessarily lead to unemployment. The reason for this is that it may not be cost-effective to deploy AI in all tasks in which AI is superior to humans.

Second, by changing the endogenous sorting of humans into firms, AI has major effects on labor and organizational outcomes. This has implications for (i) the knowledge content of human work, (ii) the distribution of firm size, productivity, and decentralization, and (iii) the productivity and the

resources supervised by humans at different knowledge levels. For example, individuals who do not directly interact with AI may experience a decrease in productivity if its introduction forces them to relocate to less productive firms.

Third, AI always benefits the most knowledgeable individuals because it provides them with less expensive technology with which to leverage their knowledge. In contrast, AI benefits the least knowledgeable individuals only when the technology is sufficiently good and compute scarce relative to the needs of society. Only in such a case do the gains these individuals achieve from AI—in terms of being able to solve more difficult problems inexpensively—outweigh the potentially negative effects caused by its introduction.

Finally, the overall increase in income generated by AI is shared by the owners of labor, compute, and production opportunities. The distribution of these gains is determined by the relative scarcities of these three different factors of production and AI's capabilities.

To conclude, we believe that this paper opens several interesting areas for future research. These areas include investigating (i) firms' incentives to develop AI, (ii) the effectiveness of reskilling programs in response to AI, and (iii) the effects of AI on the organization of international trade, offshoring, and, more generally, economic development.

APPENDIX

A The Pre-AI Equilibrium: Complete Characterization

In this Appendix, we provide the full characterization of the equilibrium without AI. As noted in the main text, this equilibrium was first described in general by [Fuchs et al. \(2015\)](#). Proposition 1 follows directly from Lemmas A.1 and A.2 and Corollary A.1.

To start, note that the First Welfare Theorem holds in this setting. Thus, the competitive equilibrium is efficient. The following lemma shows that in the pre-AI world, there is a unique efficient allocation:

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

- When $h \geq h_0 \in (0, 1)$, then $W = [0, \underline{z}]$, $I = (\underline{z}, \bar{z})$, and $S = [\bar{z}, 1]$. Moreover, workers and solvers are matched according to the strictly increasing function $m(z; \bar{z})$ given by $\int_{\bar{z}}^{m(z; \bar{z})} dG(u) = \int_0^z h(1-u)dG(u)$. Finally, the cutoffs $0 < \underline{z} < \bar{z} < 1$ satisfy:

$$(2) \quad \frac{1}{h} - \bar{z} = \int_{\bar{z}}^1 n(e(u; \bar{z})) du \quad \text{and} \quad m(\underline{z}; \bar{z}) = 1 \quad \text{where} \quad e(z; \bar{z}) = m^{-1}(z; \bar{z})$$

- When $h < h_0$, then $W = [0, \hat{z}]$, $I = \emptyset$, and $S = [\hat{z}, 1]$. Moreover, workers and solvers are matched according to the strictly increasing function $m(z; \hat{z})$ given by $\int_{\hat{z}}^{m(z; \hat{z})} dG(u) = \int_0^z h(1-u)dG(u)$. Finally, the cutoff $\hat{z} \in (0, 1)$ satisfies $m(0; \hat{z}) = \hat{z}$, $m(\hat{z}; \hat{z}) = 1$, and:

$$(3) \quad \frac{1}{h} - \hat{z} > \int_{\hat{z}}^1 n(e(z; \hat{z})) dz \quad \text{where} \quad e(z; \hat{z}) = m^{-1}(z; \hat{z})$$

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

Note that the efficient (and, therefore, equilibrium) allocation satisfies occupational stratification and strict positive assortative matching. Moreover, $W \neq \emptyset$ and $S \neq \emptyset$ but $I \neq \emptyset$ if and only if $h > h_0 \in (0, 1)$. The next step is characterizing the wages that support the allocation of Lemma A.1 as a competitive equilibrium:

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

- When $h \geq h_0$:

$$w(z) = \begin{cases} m(z; \bar{z}) - w(m(z; \bar{z}))/n(z) & \text{if } z \in W \\ z & \text{if } z \in I \\ \bar{z} + \int_{\bar{z}}^z n(e(u; \bar{z})) du & \text{if } z \in S \end{cases}$$

- When $h < h_0$:

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

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

We end by showing that w is continuous, strictly increasing, and convex (strictly so when $z \in W \cup S$) and that $w(z) > z$ for all $z \notin \text{cl}I$.

Corollary A.1. *Irrespective of whether $h \geq h_0$, 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 $z \in S$ (except possibly at $z = \inf S$), and $w'(z) > 1$, and $w''(z) > 0$ for $z \in \text{int}S$.

Proof. Consider first $h \geq h_0$. To show continuity, it is sufficient to verify that $w(z)$ is continuous at $z = \underline{z}$ and $z = \bar{z}$. Continuity at $z = \underline{z}$ follows from the fact that $\lim_{z \uparrow \underline{z}} w(z) = 1 - w(1)h(1 - \underline{z}) = \underline{z} = \lim_{z \downarrow \underline{z}} w(z)$, where we are using the fact that $m(\underline{z}; \bar{z}) = 1$ and that $w(1) = \int_{\bar{z}}^1 n(e(u; \bar{z}))du + \bar{z} = \frac{1}{h}$ (due to condition (2)). Continuity at $z = \bar{z}$ follows from Lemma A.2 because it implies that $\lim_{z \downarrow \bar{z}} w(z) = w(\bar{z}) = \bar{z}$.

We now show the remaining properties of $w(z)$. Note that if $z \in \text{int}S$, then $w'(z) = n(e(z; \bar{z})) > 1$, which also implies that $w''(z) > 0$ because 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 $z \in (\bar{z}, 1]$. Similarly, if $z \in \text{int}W$, then $w'(z) = hw(m(z; \bar{z})) > 0$, which immediately implies that $w''(z) > 0$, since both $w(z)$ and $m(z; \bar{z})$ are strictly increasing in their arguments. Given that $\lim_{z \uparrow \underline{z}} w(z) = \underline{z}$, the previous results then imply that $w(z) > z$ for all $z \in [0, \underline{z})$. Finally, $w'(z) \in (0, 1)$ follows from the fact that:

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

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

Now consider $h \leq h_0$. To show continuity, it is sufficient to verify that $w(z)$ is continuous at $z = \hat{z}$. Given that $m(\hat{z}; \hat{z}) = 1$, from Lemma A.2, we have that:

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

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

B The AI Equilibrium: Complete Characterization

In this Appendix, we provide a complete characterization of the AI equilibrium. As noted in the main text, we focus on $h < h_0$. The proof of Proposition 2 is a direct implication of Lemmas B.1 and B.2 below.

We begin the characterization with the following set of results:

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

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

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

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

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

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

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

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

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

- **Type 1 configuration:**

$$W_a^* = \emptyset, W_p^* = [0, z_{\text{AI}}], I^* = (z_{\text{AI}}, \underline{z}_1^*), S_p^* = [\underline{z}_1^*, \bar{z}_1^*], S_a^* = [\bar{z}_1^*, 1], \text{ where } z_{\text{AI}} < \underline{z}_1^* \leq \bar{z}_1^* \leq 1$$

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

- **Type 2 configuration:**

$$W_a^* = [0, \underline{z}_2^*], W_p^* = [\underline{z}_2^*, z_{AI}], I^* \subseteq \{z_{AI}\}, S_p^* = [z_{AI}, \bar{z}_2^*], S_a^* = [\bar{z}_2^*, 1], \text{ where } 0 \leq \underline{z}_2^* \leq z_{AI} \leq \bar{z}_2^* \leq 1$$

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

- **Type 3 configuration:**

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

$$\text{So } \mu_w^* = 0, \mu_s^* = \int_0^{\bar{z}_3^*} n(z)^{-1} dG(z), \mu_i^* = \mu - \mu_s^*$$

- **Type 4 configuration:**

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

$$\text{So } \mu_w^* = 0, \mu_s^* = 0, \mu_i^* = \mu$$

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

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

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

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

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

$$\Gamma_s(x) \equiv \frac{1}{h} - x - \int_x^1 n(e_s(u; x)) du$$

Consider the following partition of the knowledge space (note that $W \cup S = [0, 1]$ when $h < h_0$): $\mathcal{R}_1 \equiv \{z \in W : \Gamma_w(z) \leq 0\}$, $\mathcal{R}_2 \equiv \{z \in W : \Gamma_w(z) > 0\} \cup \{z \in S : \Gamma_s(z) > 0\}$, and $\mathcal{R}_3 \equiv \{z \in S : \Gamma_s(z) \leq 0\}$. The following lemma—which is the main result of this appendix—characterizes in detail the post-AI equilibrium:

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

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

$$\bar{z}_1^* = m_1^*(z_{AI}; \underline{z}_1^*) \quad \text{and} \quad n(z_{AI})(m_1^*(z_{AI}; \underline{z}_1^*) - z_{AI}) = \underline{z}_1^* + \int_{\underline{z}_1^*}^{m_1^*(z_{AI}; \underline{z}_1^*)} n(e_1^*(z; \underline{z}_1^*)) dz$$

where $m_1^* : [0, z_{AI}] \rightarrow [\underline{z}_1^*, \bar{z}_1^*]$ is given by $\int_{\underline{z}_1^*}^{m_1^*(z; \underline{z}_1^*)} dG(u) = \int_0^z h(1-u)dG(u)$ and $e_1^*(z; \cdot) = (m_1^*)^{-1}(z; \cdot)$

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

$$\bar{z}_2^* = m_2^*(z_{AI}; \underline{z}_2^*) \quad \text{and} \quad n(z_{AI})(m_2^*(z_{AI}; \underline{z}_2^*) - z_{AI}) = z_{AI} + \int_{z_{AI}}^{m_2^*(z_{AI}; \underline{z}_2^*)} n(e_2^*(z; \underline{z}_2^*)) dz$$

where $m_2^* : [\underline{z}_2^*, z_{AI}] \rightarrow [z_{AI}, \bar{z}_2^*]$ given by $\int_{z_{AI}}^{m_2^*(z; \underline{z}_2^*)} dG(u) = \int_{\underline{z}_2^*}^z h(1-u)dG(u)$ and $e_2^*(z; \cdot) = (m_2^*)^{-1}(z; \cdot)$

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

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

where $m_3^* : [\underline{z}_3^*, \bar{z}_3^*] \rightarrow [z_{AI}, 1]$ given by $\int_{z_{AI}}^{m_3^*(z; \underline{z}_3^*)} dG(u) = \int_{\underline{z}_3^*}^z h(1-u)dG(u)$ and $e_3^*(z; \cdot) = (m_3^*)^{-1}(z; \cdot)$

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

- $w^*(z) = z_{AI}(1 - 1/n(z)) > z$ for all $z \in W_a^*$.
- $w^*(z) = m^*(z) - w^*(m^*(z))/n(z)$ for all $z \in W_p^*$.
- $w^*(z) = z$ for all $z \in I^*$.
- $w^*(z) = \inf S_p^* + \int_{\inf S_p^*}^z n(e^*(u))du$ for all $z \in S_p^*$.
- $w^*(z) = n(z_{AI})(z - z_{AI}) > z$ for all $z \in S_a^*$.

For the interested reader, in the Online Appendix, we also provide a version of Lemma B.2 for the case in which human knowledge is uniformly distributed. In such a case, the equilibrium expressions (i.e., the wage function w^* , the equilibrium cutoffs, and the partition \mathcal{R}_j for $j = 1, 2, 3$) can be obtained in (almost) closed form.

According to Lemma B.1, AI is always used for independent production. Lemma B.2 adds to this point by stating that if $z_{AI} \in W$, then AI is also used as a worker and possibly as a solver, while if $z_{AI} \in S$, then AI is also used as a solver and possibly as a worker. Indeed, if $z_{AI} \in W$, then the post-AI equilibrium is either Type 1 (in which AI is a worker but not a solver) or Type 2 (in which AI is both a worker and solver). Similarly, if $z_{AI} \in S$, then the post-AI equilibrium is either Type 3 (in which AI is a solver but not a worker) or Type 2.

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

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

$$W_a^* = [0, \underline{z}_2^*], W_p^* = [\underline{z}_2^*, z_{AI}], I^* = \emptyset, S_p^* = [z_{AI}, \bar{z}_2^*], S_a^* = [\bar{z}_2^*, 1], \text{ where } 0 \leq \underline{z}_2^* \leq z_{AI} \leq \bar{z}_2^* \leq 1$$

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

Now consider those firms that do not use AI, i.e., the nA firms. First, let $m_2^*(z)$ be the equilibrium matching function in this case. This function must be strictly increasing and satisfy the following resource constraint $\int_{z_{AI}}^{m_2^*(z)} dG(u) = \int_{z_2^*}^z h(1-u)dG(u)$ for all $z \in [z_2^*, z_{AI}]$, which states that the total time required to consult on the problems left unsolved by workers in the interval $[z_2^*, z]$ must equal the total time available of solvers in the interval $[z_{AI}, m_2^*(z)]$. Moreover, given that $\sup W_p^* = z_{AI}$ and $\sup S_p^* = \bar{z}_2^*$, it must also be that $m_2^*(z_{AI}) = \bar{z}_2^*$.

Note then that for any given $z \in W_p$, there exists a unique increasing function $m_2^*(z)$ that satisfies both constraints. It is given by the solution to the differential equation $m_2^{*'}(z) = h(1-z)g(z)/g(m_2^*(z))$ with border condition $m_2^*(z_{AI}) = \bar{z}_2^*$ (which comes from differentiating both sides of the resource constraint). We denote such a unique function by $m_2^*(z; \bar{z}_2^*)$ (as it depends on \bar{z}_2^* through its domain).

Consider then the problem of an nA firm that recruited $n(z)$ workers of type $z \in W_p^*$ and is deciding which solver $s \in S_p^*$ to hire: $\max_{s \in S_p^*} \Pi_2^{nA}(s, z) = n(z)[s - w(z)] - w(s)$. The corresponding first-order condition evaluated at $s = m_2^*(z; \bar{z}_2^*)$ implies that $w'(z) = n(e_2^*(z; \bar{z}_2^*))$ for any $z \in S_p^*$. Thus, $w^*(z) = C^* + \int_{z_{AI}}^z n(e_2^*(u; \bar{z}_2^*))du$ for any $z \in S_p^*$. The wages of the workers hired by such firms are then determined by the zero-profit condition of nA firms: $w^*(z) = m_2^*(z; \bar{z}_2^*) - w^*(m_2^*(z; \bar{z}_2^*)/n(z)$ for any $z \in W_p^*$.

The final step is determining the constant C^* , the cutoff \bar{z}_2^* , and arguing that no firms have incentives to deviate. To do this, note that the least knowledgeable human solver has the same knowledge as AI, i.e., $\inf S_p^* = z_{AI}$. Hence, her wage must be equal to the price of one unit of compute, so $C^* = z_{AI}$. Moreover, the most knowledgeable solver hired by an nA firm has the same knowledge as the least knowledgeable solver of a bA firm. As a result, these two individuals must receive the same wage, i.e., $\lim_{z \uparrow \bar{z}_2^*} w^*(z) = \lim_{z \downarrow \bar{z}_2^*} w^*(z)$. Since $\bar{z}_2^* = m_2^*(z_{AI}; \bar{z}_2^*)$, we obtain the following:

$$(4) \quad n(z_{AI})(m_2^*(z_{AI}; \bar{z}_2^*) - z_{AI}) = z_{AI} + \int_{z_{AI}}^{m_2^*(z_{AI}; \bar{z}_2^*)} n(e_2^*(z; \bar{z}_2^*))dz$$

which is the equilibrium condition in the statement of Lemma B.2. It is then possible to prove that there is a unique cutoff \bar{z}_2^* that satisfies this condition and that such a cutoff is contained in $(0, z_{AI}]$ if and only if $z_{AI} \in \mathcal{R}_2$. This explains why this equilibrium can only arise in such a region of the parameter space. The fact that $C^* = z_{AI}$ and that \bar{z}_2^* satisfies (4) then implies that $w^*(z)$ is also continuous at the juncture between W_p^* and S_p^* :

$$\begin{aligned} \lim_{z \downarrow \bar{z}_2^*} w^*(z) &= m_2^*(\bar{z}_2^*; \bar{z}_2^*) - \frac{w^*(m_2^*(\bar{z}_2^*; \bar{z}_2^*))}{n(\bar{z}_2^*)} = z_{AI} \left(1 - \frac{1}{n(\bar{z}_2^*)}\right) = \lim_{z \uparrow \bar{z}_2^*} w^*(z) \\ \lim_{z \uparrow z_{AI}} w^*(z) &= m_2^*(z_{AI}; \bar{z}_2^*) - \frac{w^*(m_2^*(z_{AI}; \bar{z}_2^*))}{n(z_{AI})} = z_{AI} = \lim_{z \downarrow z_{AI}} w^*(z) \end{aligned}$$

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

From here, arguing that no firm has incentives to deviate is straightforward. Indeed, note that the wage function is continuous, strictly increasing, and weakly convex. This implies that if a firm does not have incentives to deviate “locally,” then it does not have incentives to deviate globally either. That bA firms do not want to deviate locally, i.e., hire a different human solver in S_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. \square

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

Proof of Step 1. We show this part only for $z_{AI} \in \mathcal{R}_1$, as the other two cases are analogous. We begin by verifying market clearing in the market for compute. By Corollary B.1, it is immediate that $\mu_i^* + \mu_w^* + \mu_s^* = \mu$. Moreover, the total time required to consult on the problems left unsolved by AI is equal to the total time available of solvers in S_a^* , i.e., $h(1 - z_{AI})\mu_w^* = \int_{z_1^*}^1 dG(z)$.

We now turn to verifying labor market clearing. First, it must be that the total time required to consult on the problems left unsolved by the human workers in the interval $[0, z] \subseteq W_p^*$ is equal to the total time available of human solvers in the interval $[z_1^*, m_1^*(z; z_1^*)] \subseteq S_p^*$. This resource constraint is satisfied as $m_1^*(z; z_1^*)$ is given by $\int_{z_1^*}^{m_1^*(z; z_1^*)} dG(u) = \int_0^z h(1 - u)dG(u)$ and $\bar{z}_1^* = m_1^*(z_{AI}; z_1^*)$.

Second, it must be that the union of the sets $(W_p^*, W_a^*, I^*, S_p^*, S_a^*)$ is $[0, 1]$, and the intersection of any two of these sets has measure zero. By Corollary B.1, this occurs if and only if $z_{AI} < z_1^* \leq \bar{z}_1^* \leq 1$. Verifying that that $z_1^* \leq \bar{z}_1^*$ is straightforward: It follows because $m_1^*(z; z_1^*)$ is strictly increasing in z plus the fact that $z_1^* = m_1^*(0; z_1^*)$ and $\bar{z}_1^* = m_1^*(z_{AI}; z_1^*)$.

Showing that $z_{AI} < z_1^*$ requires more work. Note that z_1^* is given by the solution $\Gamma_1(z_1^*; z_{AI}) = 0$, where $\Gamma_1(x; z_{AI}) \equiv n(z_{AI})(m_1^*(z_{AI}; x, z_{AI}) - z_{AI}) - x - \int_x^{m_1^*(z_{AI}; x, z_{AI})} n(e_1^*(z; x, z_{AI}))dz$ (here we are making explicit that $m_1^*(\cdot)$ and $e_1^*(\cdot)$ also depend indirectly on z_{AI} through the boundary of the set $W_p^* = [0, z_{AI}]$ to avoid any type of confusion²⁴). It is then not difficult to prove that $\Gamma_1(x; z_{AI})$ is strictly increasing in x , so $z_{AI} < z_1^*$ if and only if $\Gamma_1(0; z_{AI}) < 0$. Furthermore, note that $m_1^*(z; 0, z_{AI})$ satisfies $\int_0^{m_1^*(z; 0, z_{AI})} dG(u) = \int_0^z h(1 - u)dG(u)$ for $z \in [0, z_{AI}]$, which implies that $m_1^*(z; 0, z_{AI}) = m_w(z; z_{AI})$ (and, therefore, that $e_1^*(z; 0, z_{AI}) = e_w(z; z_{AI})$). Hence, $\Gamma_1(0; z_{AI}) = \Gamma_w(z_{AI}) < 0$ as $z_{AI} \in \mathcal{R}_1$.

Finally, we show that $\bar{z}_1^* \leq 1$. To prove it, we show that $z_1^* < \hat{z}$ and then use this result to conclude that $\bar{z}_1^* \leq 1$. As a first step, note that $m_1^*(z; \hat{z}, \hat{z})$ satisfies $\int_{\hat{z}}^{m_1^*(z; \hat{z}, \hat{z})} dG(u) = \int_0^z h(1 - u)dG(u)$ for $z \in [0, \hat{z}]$, so $m_1^*(z; \hat{z}, \hat{z}) = m(z; \hat{z})$ where $m(z; \hat{z})$ is the pre-AI matching function. Recall then that z_1^* is the unique solution $\Gamma_1(z_1^*; z_{AI}) = 0$, where $\Gamma_1(x; z_{AI})$ is strictly increasing in x . It is not difficult to prove that $\Gamma_1(x; z_{AI})$ is also strictly decreasing in z_{AI} for any given x . We then claim that $\Gamma_1(\hat{z}; z_{AI}) > 0$, which immediately implies that $z_1^* < \hat{z}$. Indeed, note that $\Gamma_1(\hat{z}; z_{AI}) \geq \Gamma_1(\hat{z}; \hat{z}) = \frac{1}{h} - \hat{z} - \int_{\hat{z}}^1 n(e(z; \hat{z}))dz$, where the first inequality follows because $\Gamma_1(x; z_{AI})$ is strictly decreasing in z_{AI} and $z_{AI} \leq \hat{z}$ (as $z_{AI} \in W$) and the last equality follows because $e_1^*(z; \hat{z}, \hat{z}) = e(z; \hat{z})$ for all $z \in [\hat{z}, 1]$. However, by Lemma A.1, we know that $\frac{1}{h} - \hat{z} - \int_{\hat{z}}^1 n(e(z; \hat{z}))dz > 0$ when $h < h_0$, so $\Gamma_1(\hat{z}; z_{AI}) > 0$.

Having proven that $z_1^* < \hat{z}$, we now show that $\bar{z}_1^* \leq 1$. By construction, \bar{z}_1^* satisfies $\int_{z_1^*}^{\bar{z}_1^*} dG(u) =$

²⁴In the statement of Lemma B.2, we simply wrote $m_1^*(z; z_1^*)$ instead of $m_1^*(z; z_1^*, z_{AI})$ to avoid cluttering notation.

$\int_0^{z_{AI}} h(1-u)dG(u)$. Note, moreover, that $z \in W$ implies that $\int_0^{z_{AI}} h(1-u)dG(u) \leq \int_0^{\hat{z}} h(1-u)dG(u) = \int_{\hat{z}}^1 dG(u)$ (as $z_{AI} \leq \hat{z}$ and $\int_0^{\hat{z}} h(1-u)dG(u) = \int_{\hat{z}}^1 dG(u)$). Hence, $\int_{\underline{z}_1^*}^{\bar{z}_1^*} dG(u) \leq \int_{\hat{z}}^1 dG(u)$. Given that $\underline{z}_1^* < \hat{z}$, it must be that $\bar{z}_1^* < 1$.

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

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

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

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

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

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

By Lemma B.3, we have that if $z_{AI} \in \mathcal{R}_1$, then $z_{AI} < \hat{z}$. This implies the following:

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

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

Now suppose for contradiction that there is a Type 2 equilibrium. As explained above (see “Informal Construction of the Equilibrium”), for this to be an equilibrium, there must exist a cutoff $\underline{z}_2^* \in [0, z_{AI}]$ with the property that $\Gamma_2(\underline{z}_2^*; z_{AI}) = 0$, where $\Gamma_2(x; z_{AI}) \equiv n(z_{AI})(m_2^*(z_{AI}; x) - z_{AI}) - z_{AI} - \int_{z_{AI}}^{m_2^*(z_{AI}; x)} n(e_2^*(z; x)) dz$. It is then not difficult to prove that $\Gamma_2(x; z_{AI})$ is strictly decreasing in x . Hence, there exists at most one \underline{z}_2^* that satisfies $\Gamma_2(\underline{z}_2^*; z_{AI}) = 0$, and for $\underline{z}_2^* \geq 0$, it must be that

$\Gamma_2(0; z_{AI}) \geq 0$, where $m_2^*(z; 0)$ is given by $\int_{z_{AI}}^{m_2^*(z; 0)} dG(u) = \int_0^z h(1-u)dG(u)$ for $z \in [0, z_{AI}]$. However, from this last condition we have that $m_2^*(z; 0) = m_w(z; z_{AI})$ (as $m_2^*(z; 0)$ and $m_w(z; z_{AI})$ satisfy the same condition), so $\Gamma_2(0; z_{AI}) = \Gamma_w(z_{AI})$. Consequently, for $\underline{z}_2^* \geq 0$, we need that $\Gamma_w(z_{AI}) \geq 0$, which contradicts the fact that $z_{AI} \in \mathcal{R}_1$.

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

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

Moreover, following similar reasoning as that developed above for a Type 2 equilibrium, for this to be an equilibrium, (i) it must be that $\bar{z}_4^* = m_4^*(0; \underline{z}_4^*)$, where $m_4^*(0; \underline{z}_4^*)$ satisfies $m_4^*(\underline{z}_4^*; \underline{z}_4^*) = 1$ and $\int_{m_4^*(0; \underline{z}_4^*)}^{m_4^*(z; \underline{z}_4^*)} dG(u) = \int_{\underline{z}_4^*}^z h(1-u)dG(u)$ for $z \in [0, \underline{z}_4^*]$, and (ii) there must exist a cutoff $\underline{z}_4^* < m_4^*(0; \underline{z}_4^*)$ such that $1/h - m_4^*(0; \underline{z}_4^*) - \int_{m_4^*(0; \underline{z}_4^*)}^1 n(e_4^*(z; \underline{z}_4^*))dz = 0$. However, by Lemma A.1, we know that the unique solution to these equilibrium conditions is $\underline{z}_4^* = \underline{z}$ and $\bar{z}_4^* = m_4^*(0; \underline{z}_4^*) = \bar{z}$, 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 $z_{AI} \in I^* = (\underline{z}, \bar{z})$, which contradicts the assumption that $z_{AI} \in \mathcal{R}_1$. \square

We end this appendix with some properties of the partition $(\mathcal{R}_1, \mathcal{R}_2, \mathcal{R}_3)$ that will be useful later. In particular, the next lemma states that when z_{AI} is sufficiently low, the equilibrium is Type 1, so AI is used as a worker and as an independent producer (but not as a solver) in that case. Similarly, when $z_{AI} = \hat{z}$ or $z_{AI} \rightarrow 1$ (where \hat{z} is the knowledge cutoff to become a solver in the pre-AI equilibrium), the equilibrium is Type 2, so AI is used in all three possible roles.

Lemma B.3. (i) $z_{AI} \in \mathcal{R}_1$ if $z_{AI} \in [0, \epsilon)$ with $\epsilon \downarrow 0$, (ii) $z_{AI} = \hat{z} \in \mathcal{R}_2$ (where $\{\hat{z}\} = W \cap S$), and (iii) $z_{AI} \in \mathcal{R}_2$ if $z_{AI} \in [1 - \epsilon, 1)$ with $\epsilon \downarrow 0$.

Proof. First we show that $z_{AI} \in \mathcal{R}_1$ if $z_{AI} \in [0, \epsilon)$ with $\epsilon \downarrow 0$. Note that $\Gamma_w(0) = 0$ (as $m_w(z; 0) = 0$), and that $\Gamma'_w(0) = -1$, as:

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

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

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

Finally, we show that $z_{AI} \in \mathcal{R}_2$ if $z_{AI} \in [1 - \epsilon, 1)$ with $\epsilon \downarrow 0$. First, note that if $z_{AI} \in [1 - \epsilon, 1)$ with $\epsilon \downarrow 0$, then $z_{AI} \in S$. Second, is not difficult to prove that $\Gamma_s(1) = 1/h - 1 > 0$, so by continuity, $\Gamma_s(z_{AI}) > 0$ for all $z_{AI} \in [1 - \epsilon, 1)$ with $\epsilon \downarrow 0$. Hence, if $z_{AI} \in [1 - \epsilon, 1)$ with $\epsilon \downarrow 0$, then $z_{AI} \in S$ and $\Gamma_s(z_{AI}) > 0$, so $z_{AI} \in \mathcal{R}_2$. \square

C Proofs Omitted from Section 4

C.1 Proof of Proposition 3

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

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

Now, suppose the equilibrium is Type 3. Then, $\inf S^* = z_{AI} > \hat{z}$, immediately implying that $S^* \subset S$. To prove that $W^* \supset W$, we show that $\sup W^* = \bar{z}_3^* > \hat{z}$. Indeed, recall that \bar{z}_3^* is given by:²⁵

$$\bar{z}_3^* = e_3^*(1; \underline{z}_3^*, z_{AI}) \quad \text{and} \quad \frac{1}{h} = z_{AI} + \int_{z_{AI}}^1 n(e_3^*(z; \underline{z}_3^*, z_{AI})) dz$$

where $\int_{z_{AI}}^{m_3^*(z; \underline{z}_3^*, z_{AI})} dG(u) = \int_{\underline{z}_3^*}^z h(1-u)dG(u)$ for $z \in [\underline{z}_3^*, \bar{z}_3^*]$

Using the fact that $m_3^*(\bar{z}_3^*; \underline{z}_3^*, z_{AI}) = 1$, the equilibrium conditions that determine \underline{z}_3^* and \bar{z}_3^* can be written as follows:²⁶

$$\underline{z}_3^* = \bar{e}_3^*(z_{AI}; \bar{z}_3^*, z_{AI}) \quad \text{and} \quad \frac{1}{h} = z_{AI} + \int_{z_{AI}}^1 n(\bar{e}_3^*(z; \bar{z}_3^*, z_{AI})) dz$$

where $G(z) = 1 - \int_{\bar{e}_3^*(z; \bar{z}_3^*, z_{AI})}^{\bar{z}_3^*} h(1-u)dG(u)$ for $z \in [z_{AI}, 1]$

Define $\Gamma_3(x; z_{AI}) \equiv \frac{1}{h} - z_{AI} - \int_{z_{AI}}^1 n(\bar{e}_3^*(z; x, z_{AI})) dz$. It is not difficult to prove that $\Gamma_3(x; z_{AI})$ is strictly decreasing in x and strictly increasing in z_{AI} . Moreover, \bar{z}_3^* is given by the unique solution to $\Gamma_3(\bar{z}_3^*; z_{AI}) = 0$. To prove that $\bar{z}_3^* > \hat{z}$, it suffices to show that $\Gamma_3(\hat{z}; z_{AI}) > 0$. Since $z_{AI} > \hat{z}$, we have that $\Gamma_3(\hat{z}; z_{AI}) > \Gamma_3(\hat{z}; \hat{z}) = \frac{1}{h} - \hat{z} - \int_{z_{AI}}^1 n(e(z; \hat{z})) dz > 0$, where the second-to-last inequality follows because $\bar{e}_3^*(z; \hat{z}, \hat{z}) = e(z; \hat{z})$ for all $z \in [\hat{z}, 1]$, and the last inequality comes from the pre-AI equilibrium characterized in Lemma A.1. \square

C.2 Proof of Corollary 1

Since all two-layer organizations hire a single solver, to prove this corollary, it is sufficient to show that AI increases the overall number of solvers in the economy. When $z_{AI} \in \text{int}W$, this follows because $S^* \supset S$. When $z_{AI} \in \text{int}S$, more humans become workers after AI's introduction (i.e., $W^* \supset W$). Hence, the overall number of solvers—human plus AI—must increase as each worker requires the same amount of help post-AI as that required pre-AI. \square

²⁵To avoid any type of confusion, we are making explicit that $m_3^*(z; \underline{z}_3^*, z_{AI})$ and $e_3^*(1; \underline{z}_3^*, z_{AI})$ depend on both \underline{z}_3^* and z_{AI} (in the statement of Lemma B.2, we simply wrote $m_3^*(z; \underline{z}_3^*)$ instead of $m_3^*(z; \underline{z}_3^*, z_{AI})$ to avoid cluttering notation).

²⁶Note that $\bar{e}_3^*(z; \bar{z}_3^*, z_{AI})$ is the equilibrium employee function indexed by \bar{z}_3^* instead of \underline{z}_3^* .

C.3 Proof of Proposition 4

Consider first span of control. The most decentralized firm of the pre-AI economy has a span of control $n(\sup W)$, while the most decentralized firm of the post-AI world has a span of control of $n(\sup W^*)$. That AI decreases the maximum span of control when $z_{AI} \in \text{int}W$, and increases it when $z_{AI} \in \text{int}S$, follows from occupational displacement, i.e. Proposition 3, and the fact that $n(z) = [h \times (1 - z)]^{-1}$ is strictly increasing in z .

Consider next productivity. The least productive firm of the pre-AI economy has productivity $\inf S$, while the least productive firm of the post-AI world has productivity $\inf S^*$. That AI decreases the minimum firm productivity when $z_{AI} \in \text{int}W$, and increases it when $z_{AI} \in \text{int}S$, follows directly from occupational displacement (i.e. Proposition 3).

Finally, consider size. As noted in the main text, the size of the smallest firm is $n(0)$ times the minimum firm productivity, while the size of the largest firm is 1 times the maximum span of control. Consequently, when $z_{AI} \in \text{int}W$, AI reduces the minimum and maximum firm size, as it decreases the maximum span of control and the minimum firm productivity. Likewise, when $z_{AI} \in \text{int}S$, AI increases the minimum and maximum firm size, as it increases the maximum span of control and the minimum firm productivity. \square

C.4 Proof of Corollary 2

First, note that when $z_{AI} \in \text{int}S$, the largest post-AI firms are larger than the largest pre-AI firms.

Second, we know by Lemma B.3 that $z_{AI} \in \mathcal{R}_2$ if $z_{AI} \in [1 - \epsilon, 1) \subset S$ with $\epsilon \downarrow 0$. Thus, by Lemma B.2, firms use AI in all three possible roles when z_{AI} is sufficiently close to 1. Since no worker is better than AI and the equilibrium exhibits positive assortative matching (Lemma B.1), we then have that AI is the best worker in the economy and is therefore supervised by the most knowledgeable humans. Hence, when z_{AI} is sufficiently close to 1, the largest firms post-AI are bottom-automated firms.

Combining these two results, we have that when z_{AI} is sufficiently close to 1, AI leads to the creation of superstar firms with scale but no mass. \square

C.5 Proof of Proposition 5

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

Moreover, for the proof that follows, it is important that we make explicit that the pre-AI matching and employee functions depend on \hat{z} (i.e., the knowledge threshold that separates workers and solvers in the pre-AI equilibrium). For this reason, we write $m(z; \hat{z})$ and $e(z; \hat{z})$ instead of $m(z)$ and $e(z)$, respectively.

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

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

$$\begin{aligned} \int_{\underline{z}_1^*}^{m_1^*(z; \underline{z}_1^*)} dG(u) &= \int_0^z h(1-u) dG(u) \text{ for } z \in W_p^* = [0, z_{AI}] \\ \int_{\hat{z}}^{m(z; \hat{z})} dG(u) &= \int_0^z h(1-u) dG(u) \text{ for } z \in W = [0, \hat{z}] \end{aligned}$$

Thus, if $z \in W_p^* \cap W = W_p^*$, then $\int_{\underline{z}_1^*}^{m_1^*(z; \underline{z}_1^*)} dG(u) = \int_{\hat{z}}^{m(z; \hat{z})} dG(u)$, so $m_1^*(z; \underline{z}_1^*) < m(z; \hat{z})$ 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:

$$\begin{aligned} \int_{z_{AI}}^{m_2^*(z; \underline{z}_2^*)} dG(u) &= \int_{\underline{z}_2^*}^z h(1-u) dG(u) \text{ for } z \in W_p^* = [\underline{z}_2^*, z_{AI}] \\ \int_{\hat{z}}^{m(z; \hat{z})} dG(u) &= \int_0^z h(1-u) dG(u) \text{ for } z \in W = [0, \hat{z}] \end{aligned}$$

Consequently, if $z \in W_p^* \cap W = W_p^*$, then $\int_{\hat{z}}^{m(z; \hat{z})} dG(u) - \int_{z_{AI}}^{m_2^*(z; \underline{z}_2^*)} dG(u) = \int_0^{\underline{z}_2^*} h(1-u) dG(u) > 0$, which implies that $m_2^*(z; \underline{z}_2^*) < m(z; \hat{z})$ as $z_{AI} < \hat{z}$.

We now turn to solvers, i.e., those $z \in S \subset S^*$. We first claim that if $e(z; \hat{z}) = z_{AI}$, then $z \in S_a^* \cap S$, which immediately implies that if $e(z; \hat{z}) = z_{AI}$, then z has the same span of control pre- and post-AI. The proof is via the contrapositive. Suppose that $z \notin S_a^* \cap S$ (but that z is a solver). Then $z \in S_p^* \cap S$. However, if this is the case, then $z_{AI} \geq e_j^*(z; \underline{z}_j^*) > e(z; \hat{z})$, where the first inequality follows because AI is the best worker, and the second inequality follows because $e_j^*(z'; \underline{z}_j^*) > e(z'; \hat{z})$ for all $z' \in S_p^* \cap S$ if $m_j^*(z''; \underline{z}_j^*) < m(z''; \hat{z})$ for all $z'' \in W_p^* \cap W = W_p^*$ (which we already showed is true for $j = 1, 2$). Hence, $e(z; \hat{z}) \neq z_{AI}$.

The previous claim implies that if $e(z; \hat{z}) < z_{AI}$, then z 's span of control increases with AI, while if $e(z; \hat{z}) > z_{AI}$, then z 's span of control decreases with AI. Indeed, if $e(z; \hat{z}) < z_{AI}$, then either $z \in S_p^* \cap S$ or $z \in S_a^* \cap S$. In either case, z is assisting more knowledgeable workers post-AI than pre-AI: If $z \in S_p^*$, then we already know that $e_j^*(z; \underline{z}_j^*) > e(z; \hat{z})$, while if $z \in S_a^*$, then, post-AI, she is assisting AI, while pre-AI, she was humans with knowledge $e(z; \hat{z}) < z_{AI}$. Similarly, if $e(z; \hat{z}) > z_{AI}$, then $z \in S_a^* \cap S$. The latter implies that z is assisting less knowledgeable workers post-AI than pre-AI, as post-AI, she is assisting the work of AI, while pre-AI, she was assisting humans with knowledge $e(z; \hat{z}) > z_{AI}$. \square

• $z_{AI} \in \text{int}S$.— First, we show that every $z \in S^* \subset S$ has a larger span of control post-AI than pre-AI. Note that if $z \in S_a^*$, then such a solver is matched with AI in the post-AI equilibrium. However, if so, then $e(z; \hat{z}) \leq \hat{z} < z_{AI}$, as $z_{AI} \in \text{int}S$.

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

given by:

$$\int_z^{\bar{z}_2^*} dG(u) = \int_{e_2^*(z; \underline{z}_2^*)}^{z_{AI}} h(1-u) dG(u) \text{ for } z \in S_p^* = [z_{AI}, \bar{z}_2^*]$$

$$\int_z^1 dG(u) = \int_{e(z; \hat{z})}^{\hat{z}} h(1-u) dG(u) \text{ for } z \in S = [\hat{z}, 1]$$

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

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

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

$$\int_z^1 dG(u) = \int_{e(z; \hat{z})}^{\hat{z}} h(1-u) dG(u) \text{ for } z \in S = [\hat{z}, 1]$$

Consequently, for $z \in S_p^* \cap S = S_p^*$, 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 this is the case, 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 $z = e(z_{AI}; \hat{z})$, then $z \in W_a^* \cap W$, which immediately implies that if $z = e(z_{AI}; \hat{z})$, 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 $z_{AI} \leq m_j^*(z; \underline{z}_j^*) < m(z; \hat{z})$, where the first inequality follows because AI is the worst solver, and the second inequality follows because $m_j^*(z'; \underline{z}_j^*) < m(z'; \hat{z})$ for all $z' \in W_p^* \cap W$ if $e_j^*(z''; \underline{z}_j^*) > e(z''; \hat{z})$ for all $z'' \in S_p^* \cap S = S_p^*$ (which we already showed is true for $j = 2, 3$). Hence, $z \neq e(z_{AI}; \hat{z})$.

The previous claim implies that if $z < e(z_{AI}; \hat{z})$, then z is strictly more productive post-AI than pre-AI, while if $z > e(z_{AI}; \hat{z})$, then z is strictly less productive post-AI than pre-AI. Indeed, if $z < e(z_{AI}; \hat{z})$, then $z \in W_a^* \cap W$. Hence, z is matched with a better solver post-AI than pre-AI because post-AI, she is assisted by AI, while pre-AI, she was assisted by a human with knowledge $m(z; \hat{z}) < z_{AI}$. Similarly, if $z > e(z_{AI}; \hat{z})$, then $z \in W_a^* \cap W$ or $z \in W_p^* \cap W$. If $z \in W_a^* \cap W$, then z is matched with a worse solver post-AI than pre-AI because post-AI, she is assisted by AI, while pre-AI, she was assisted by a human with knowledge $m(z; \hat{z}) > z_{AI}$. Similarly, if $z \in W_p^* \cap W$, then z is also matched with a worse solver post-AI than pre-AI because we already established that $m_j^*(z; \underline{z}_j^*) < m(z; \hat{z})$ for $j = 2, 3$. \square

C.6 Proof of Lemma 2

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

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

Proof. That $\Delta(z_{AI}) = z_{AI} - w(z_{AI}) < 0$ follows directly from the fact that $w(z) > z$ for all $z \in [0, 1]$ when $h < h_0$ (see Proposition 1). Consider next $\Delta(1) = w^*(1) - w(1)$. By Lemma B.2, we have that $w^*(1) = 1/h$, while by Lemma A.1 and Corollary A.1, we have that $w(1) < 1/h$. Hence, $\Delta(1) > 0$.

Finally, consider $\Delta(0) = w^*(0) - w(0)$ and suppose that $z_{AI} \in S$. From Lemma B.2, we have that $w^*(0) = z_{AI}(1 - h)$ as the human with zero knowledge is assisted by AI irrespective of whether $z_{AI} \in \mathcal{R}_2$ or $z_{AI} \in \mathcal{R}_3$. Moreover, from Lemma A.1 and Corollary A.1 we have that $w(0) = \hat{z} - hw(\hat{z})$. Thus, $\Delta(0) = z_{AI}(1 - h) - (\hat{z} - hw(\hat{z})) > (1 - h)(z_{AI} - \hat{z}) \geq 0$, where the first inequality follows because $w(\hat{z}) > \hat{z}$, and the second inequality follows because $z_{AI} \geq \hat{z}$. \square

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

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

• $z_{AI} \in \text{int}W$.— Lemma B.2 implies that $\sup W^* = z_{AI}$, while Proposition 3 that $W^* \subset W$ and $S^* \supset S$. Hence, if $z > z_{AI}$, then z can only belong to either $I^* \cap W$, $S^* \cap W$, or $S^* \cap S$.

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

Consider then $z \in S^* \cap S$. Since $S^* = S_p^* \cup S_a^*$, here we have two cases to consider: $z \in S_p^* \cap S$ and $z \in S_a^* \cap S$. If $z \in S_p^* \cap S$, then $\Delta'(z) = n(e_j^*(z; \underline{z}_j^*)) - n(e(z; \hat{z}))$, where $e_j^*(z; \underline{z}_j^*)$ is the employee function in a Type $j = 1, 2$ equilibrium and $e(z; \hat{z})$ employee function in the pre-AI equilibrium. However, by Proposition 5, we know that $e_j^*(z; \underline{z}_j^*) > e(z; \hat{z})$ as every $z \in S_p^* \cap S$ supervises better workers post-AI than pre-AI. Hence, in this case, $\Delta'(z) > 0$ also. Consequently, if $\Delta(z)$ crosses zero when $z \in S_p^* \cap S$, then it necessarily crosses from below.

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

• $z_{AI} \in \text{int}S$.— From Lemma B.2, we know that $\inf S^* = z_{AI}$. Moreover, from Proposition 3, we have that $W^* \supseteq W$ and $S^* \subseteq S$. Consequently, if $z > z_{AI}$, then $z \in S^* \cap S$ necessarily. Since $S^* = S_p^* \cup S_a^*$, here we have two cases to consider: $z \in S_p^* \cap S$ and $z \in S_a^* \cap S$.

If $z \in S_p^* \cap S$, then $\Delta'(z) = n(e_j^*(z; \underline{z}_j^*)) - n(e(z; \hat{z}))$, while if $z \in S_a^* \cap S$, then $\Delta'(z) = n(z_{AI}) - n(e(z; \hat{z}))$, where $e_j^*(z; \underline{z}_j^*)$ is the employee function in a Type $j = 2, 3$ equilibrium and $e(z; \hat{z})$ the employee function in the pre-AI equilibrium. In either case, $\Delta'(z) > 0$ since Proposition 5 states that $e_j^*(z; \underline{z}_j^*) > e(z; \hat{z})$ and $z_{AI} > e(z; \hat{z})$ when AI has the knowledge of a pre-AI solver. Consequently, $\Delta'(z) > 0$ for all $z \geq z_{AI}$, so if $\Delta(z)$ crosses zero at some $z > z_{AI}$, then it always crosses it from below. \square

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

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

• $z_{AI} \in \text{int}W$.— Lemma B.2 implies that $\sup W^* = z_{AI}$, while Proposition 3 that $W^* \subseteq W$ and $S^* \supseteq S$. Consequently, if $z < z_{AI}$, then $z \in W^* \cap W$, where $W^* = W_a^* \cup W_p^*$.

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

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

Consequently, if $\Delta(z) = 0$ at some z in this interval, say at $z = \zeta$, then $\Delta'(\zeta) = m_j^*(\zeta; \underline{z}_j^*) - m(\zeta; \hat{z}) \leq 0$, where the last inequality follows because every worker is assisted by a worse solver post-AI than pre-AI when $z_{AI} \in \text{int}W$ (as shown in Proposition 5).

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

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

Consequently, if $\Delta(z) = 0$ at some z in this interval, say at $z = \zeta$, then $\Delta'(\zeta) = z_{AI} - m(\zeta; \hat{z}) \leq 0$, where the inequality follows, again, from Proposition 5.

• $z_{AI} \in \text{int}S$.— In this case, Lemma B.2 implies that $\inf S^* = z_{AI}$, while Proposition 3 that $W^* \supset W$ and $S^* \subset S$. Consequently, if $z < z_{AI}$, then z can only belong to either $I^* \cap S$, $W^* \cap S$, or $W^* \cap W$.

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

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

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

where $m_j^*(z; \underline{z}_j^*)$ is the matching function in a Type $j = 2, 3$ equilibrium, and $m(z; \hat{z})$ the matching function of the pre-AI equilibrium. Consequently, if $\Delta(z) = 0$ at some z in this interval, say at $z = \zeta$, then $\Delta'(\zeta) = m_j^*(\zeta; \underline{z}_j^*) - m(\zeta; \hat{z}) \leq 0$, where the last inequality follows because every $z \in W^* \cap W$ is assisted by a worse solver post-AI than pre-AI when $z_{AI} \in \text{int}S$.

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

C.7 Proof of Proposition 6

Part (i) (“there always exists a set $[z, 1] \subseteq (z_{AI}, 1]$ of winners”) follows directly from Lemma 2 and the fact that $\Delta(1) > 0$ (see Claim C.1). Hence, part (ii) (“there exists a set $[0, z] \subseteq [0, z_{AI}]$ of winners if and only if $z_{AI} > \bar{z}_{AI}$, where $\bar{z}_{AI} \in \text{int}W$ ”) is the only part that remains to be proven. To do so, we begin by constructing \bar{z}_{AI} and then show that the statement is true.

Let $\Delta(0; z_{AI}) \equiv w^*(0; z_{AI}) - w(0)$, and define \bar{z}_{AI} as the solution to $\Delta(0; \bar{z}_{AI}) = 0$. We first show that \bar{z}_{AI} exists, is unique, and that $\bar{z}_{AI} \in (0, \hat{z})$. We do this by showing that $\Delta(0; z_{AI})$ crosses zero once as we move from $z_{AI} = 0$ to $z_{AI} = 1$, and that this crossing point is at a $z_{AI} < \hat{z}$. Indeed, if $z_{AI} \geq \hat{z}$, then Claim C.1 states that $\Delta(0; z_{AI}) > 0$ (as $z_{AI} \in S$ in this case). Moreover, as shown in Lemma B.3, $0 \in \mathcal{R}_1$, so when $z_{AI} = 0$, the equilibrium is always Type 1. The latter implies that $\Delta(0; 0) = z_1^*(0)(1 - h) - w(0) = -w(0) < 0$,²⁷ where the second-to-last equality follows because $z_1^*(0) = 0$, as can be easily be proven from the condition that determines $z_1^*(z_{AI})$ (see the statement of Lemma B.2).

Now, when $z \in [0, \hat{z}] = W$, the equilibrium is either Type 1, in which case $w^*(0; z_{AI}) = z_1^*(z_{AI})(1 - h)$, or Type 2, in which case $w^*(0; z_{AI}) = z_{AI}(1 - h)$. Using the equilibrium condition that determines $z_1^*(z_{AI})$, it is not difficult to prove that (i) $z_1^*(z_{AI})$ is strictly increasing in z_{AI} , and that (ii) $z_1^*(z_{AI}) = z_{AI}$ whenever we switch from a Type 1 to a Type 2 equilibrium (and vice versa). Consequently, irrespective of the equilibrium type in this region, $w^*(0; z_{AI})$ is continuous and strictly increasing in z_{AI} , which implies that $\Delta(0; z_{AI})$ is also continuous and strictly increasing in z_{AI} . This result, combined with the fact that $\Delta(0; 0) < 0$ and $\Delta(0; \hat{z}) > 0$, immediately yields the desired result.

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

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

References

ACEMOGLU, D. (2024): “The Simple Macroeconomics of AI,” Working Paper.

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

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

²⁷To avoid any confusion, here we make explicit that $z_1^*(z_{AI})$ depends on z_{AI} as seen from the condition that determines $z_1^*(z_{AI})$ in the statement of Lemma B.2.

- ACEMOGLU, D. AND P. RESTREPO (2018): “The Race Between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment,” *American Economic Review*, 108, 1488–1542.
- (2019): “8. Artificial Intelligence, Automation, and Work,” in *The Economics of Artificial Intelligence: An Agenda*, ed. by A. Agrawal, J. Gans, and A. Goldfarb, University of Chicago Press, 197–236.
- (2022): “Tasks, Automation, and the Rise in U.S. Wage Inequality,” *Econometrica*, 90, 1973–2016.
- AGHION, P., B. F. JONES, AND C. I. JONES (2017): “Artificial Intelligence and Economic Growth,” NBER Working Paper No. 23928.
- AGRAWAL, A., J. GANS, AND A. GOLDFARB (2019): “Introduction,” in *The Economics of Artificial Intelligence: An Agenda*, ed. by A. Agrawal, J. Gans, and A. Goldfarb, University of Chicago Press, 1–19.
- ALONSO, R., W. DESSEIN, AND N. MATOUSCHEK (2008): “When Does Coordination Require Centralization?” *American Economic Review*, 98, 145–179.
- ANTRÀS, P., L. GARICANO, AND E. ROSSI-HANSBERG (2006): “Offshoring in a Knowledge Economy,” *The Quarterly Journal of Economics*, 121, 31–77.
- AUTOR, D. (2024): “Applying AI to Rebuild Middle Class Jobs,” NBER Working Paper No. 32140.
- AUTOR, D., D. DORN, L. F. KATZ, C. PATTERSON, AND J. VAN REENEN (2020): “The fall of the labor share and the rise of superstar firms,” *The Quarterly Journal of Economics*, 135, 645–709.
- AUTOR, D. H., F. LEVY, AND R. J. MURNANE (2003): “The Skill Content of Recent Technological Change: An Empirical Exploration*,” *The Quarterly Journal of Economics*, 118, 1279–1333.
- AZAR, J., M. CHUGUNOVA, K. KELLER, AND S. SAMILA (2023): “Monopsony and Automation,” Max Planck Institute for Innovation & Competition Research Paper No. 23-21.
- BLOOM, N., L. GARICANO, R. SADUN, AND J. VAN REENEN (2014): “The Distinct Effects of Information Technology and Communication Technology on Firm Organization,” *Management Science*, 60, 2859–2885.
- BRYNJOLFSSON, E. (2022): “The Turing Trap: The Promise & Peril of Human-Like Artificial Intelligence,” *Daedalus*, 151, 272–287.
- BRYNJOLFSSON, E., D. LI, AND L. R. RAYMOND (2023): “Generative AI at Work,” NBER Working Paper No. 31161.
- BRYNJOLFSSON, E. AND A. MCAFEE (2016): *The Second Machine Age*, WW Norton.
- BRYNJOLFSSON, E., A. MCAFEE, M. SORELL, AND F. ZHU (2008): “Scale without mass: business process replication and industry dynamics,” *Harvard Business School Technology & Operations Mgt. Unit Research Paper*.
- 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*.
- CREMER, J., L. GARICANO, AND A. PRAT (2007): "Language and the Theory of the Firm," *The Quarterly Journal of Economics*, 122, 373–407.
- DESSEIN, W., A. GALEOTTI, AND T. SANTOS (2016): "Rational Inattention and Organizational Focus," *American Economic Review*, 106, 1522–1536.
- DESSEIN, W. AND T. SANTOS (2006): "Adaptive Organizations," *Journal of Political Economy*, 114, 956–995.
- ECKHOUT, J. AND P. KIRCHER (2018): "Assortative Matching with Large Firms," *Econometrica*, 86, 85–132.
- FUCHS, W., L. GARICANO, AND L. RAYO (2015): "Optimal Contracting and the Organization of Knowledge," *The Review of Economic Studies*, 82, 632–658.
- GARICANO, L. (2000): "Hierarchies and the Organization of Knowledge in Production," *Journal of Political Economy*, 108, 874–904.
- GARICANO, L. AND T. HUBBARD (2007): "Managerial Leverage Is Limited by the Extent of the Market: Hierarchies, Specialization, and the Utilization of Lawyers' Human Capital," *The Journal of Law & Economics*, 50, 1–43.
- GARICANO, L. AND T. N. HUBBARD (2012): "Learning About the Nature of Production from Equilibrium Assignment Patterns," *Journal of Economic Behavior & Organization*, 84, 136–153.
- GARICANO, L. AND E. ROSSI-HANSBERG (2004): "Inequality and the Organization of Knowledge," *American Economic Review Papers and Proceedings*, 94, 197–202.
- (2006): "Organization and Inequality in a Knowledge Economy," *The Quarterly Journal of Economics*, 121, 1383–1435.
- (2012): "Organizing growth," *Journal of Economic Theory*, 147, 623–656.
- (2015): "Knowledge-Based Hierarchies: Using Organizations to Understand the Economy," *Annual Review of Economics*, 7, 1–30.
- GIBBONS, R. AND J. ROBERTS (2013): *The Handbook of Organizational Economics*, Princeton University Press.
- GOLDFARB, A. AND C. TUCKER (2019): "Digital Economics," *Journal of Economic Literature*, 57, 3–43.
- GUMPERT, A., H. STEIMER, AND M. ANTONI (2022): "Firm Organization with Multiple Establishments," *The Quarterly Journal of Economics*, 137, 1091–1138.

- IDE, E. AND E. TALAMÀS (2024): "Artificial Intelligence in the Knowledge Economy," ArXiv preprint arXiv:2312.05481 (version 6).
- JOHNSON, S. AND D. ACEMOGLU (2023): *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity*, Hachette UK.
- JONES, C. I. (2024): "The A.I. Dilemma: Growth versus Existential Risk," NBER Working Paper No. 31837.
- KAHNEMAN, D. (2019): "24. Artificial Intelligence and Behavioral Economics: A Comment," in *The Economics of Artificial Intelligence: An Agenda*, ed. by A. Agrawal, J. Gans, and A. Goldfarb, University of Chicago Press, 608–610.
- KORINEK, A. AND D. SUH (2024): "Scenarios for the Transition to AGI," NBER Working Paper No. 32255.
- MCELHERAN, K., J. F. LI, E. BRYNJOLFSSON, Z. KROFF, E. DINLERSOZ, L. S. FOSTER, AND N. ZOLAS (2023): "AI Adoption in America: Who, What, and Where," NBER Working Paper No. 31788.
- MESEROLE, C. O. (2018): "What is Machine Learning?" Tech. rep., Brookings Institution.
- MOLL, B., L. RACHEL, AND P. RESTREPO (2022): "Uneven Growth: Automation's Impact on Income and Wealth Inequality," *Econometrica*, 90, 2645–2683.
- MURO, M., J. WHITON, AND R. MAXIM (2019): "What Jobs are Affected by AI?" Tech. rep., Brookings Institution.
- NORDHAUS, W. D. (2007): "Two Centuries of Productivity Growth in Computing," *The Journal of Economic History*, 67, 128–159.
- NOY, S. AND W. ZHANG (2023): "Experimental evidence on the productivity effects of generative artificial intelligence," *Science*, 381, 187–192.
- OTIS, N., R. CLARKE, S. DELECOURT, D. HOLTZ, AND R. KONING (2023): "The Uneven Impact of Generative AI on Entrepreneurial Performance," Available at SSRN 4671369.
- PENG, S., E. KALLIAMVAKOU, P. CIHON, AND M. DEMIRER (2023): "The Impact of AI on Developer Productivity: Evidence from GitHub Copilot," ArXiv preprint arXiv:2302.06590.
- SVANBERG, M., W. LI, M. FLEMING, B. GOEHRING, AND N. THOMPSON (2024): "Beyond AI Exposure: Which Tasks are Cost-Effective to Automate with Computer Vision?" Available at SSRN 4700751.
- ZEIRA, J. (1998): "Workers, Machines, and Economic Growth," *The Quarterly Journal of Economics*, 113, 1091–1117.