ORGANIZING INTELLIGENCE*

Alex Imas† Andrew Koh† Shane Legg

Google DeepMind

October 29, 2026

Abstract

When will advances in AI lead to economic transformation? We develop a model of how *intelligence*, the ability to achieve goals, matters for *organizations*, the key unit of economic production. Agents (human or AI) with multidimensional capabilities sort endogenously into competing firms. A firm's productivity depends on its agents' capabilities; we distinguish *competence*, which shifts productivity levels, from *coordination*, which governs how productivity scales with size. Two differences between AI and humans drive transformation: AI expertise is reproducible while humans' are in fixed supply, and AIs may not face the same limits for coordination as humans. We analyze how equilibrium market structure, productivity, employment, and wages evolve as AI capabilities improve. When AI is more capable than the average human, organizations of complementary humans can nonetheless outperform AI. Increases in AI competence can change the composition of workers within firms and decrease firm size. As AI coordination improves, however, this can break the Coasean ceiling on firm size and the market structure undergoes a phase transition toward large, concentrated AI-heavy firms. Changes in human employment are gradual when AI displaces workers within mixed human-AI firms, or sudden when AI firms outcompete human firms. We decompose wage changes into a *technology effect*, which changes the value of humans within existing firms, and a *reorganization effect* as firms adjust size and composition. Reorganization can reverse the technology effect on wages.

*We thank Allan Dafoe, Matt Elliott, Joshua Gans, Gillian Hadfield, Andrea Prat, Eduard Talamas, Phil Trammell, and Birger Wernerfelt for helpful conversations and suggestions. We also thank Pedro Aldighieri for superb comments and research assistance. Various AI systems were used for assisting with mathematical arguments as well as proofreading for consistency and clarity. The paper expresses the views of the individual authors and does not necessarily reflect those of their affiliated institutions. This research is intended solely for academic discussion and does not represent corporate forecasts, product roadmaps, or empirical statements regarding Google's current or future operations.

†Lead authors with equal contribution listed in alphabetical order; emails: imasa@google.com and andrewkoh@google.com. Imas is also affiliated with The University of Chicago; Koh is also affiliated with Columbia University.

1 INTRODUCTION

When will AI transform the economy? What kinds of firms will emerge, and what will happen to the returns to human expertise and to wages? We develop a model of how the capabilities of human and artificial intelligence interact with economic forces to jointly determine the composition and size of firms, surplus, and human wages. We model *intelligence* as the 'ability to achieve goals' (Legg and Hutter, 2007). This has typically been measured through isolated benchmarks that evaluate a single agent in isolation. But the economically relevant unit is the *organization* that competes in equilibrium under the forces of supply and demand (Samuelson, 1948).

Economic models often treat new technology either as a tool that makes workers more productive or as a factor that substitutes for humans task-by-task. But AI is different in three conceptually important ways. First, human capabilities are fixed, which can generate Ricardian scarcity rents in the form of higher wages for increasing expertise; by contrast, AI capabilities can be copied at the cost of compute which eliminates these rents. Second, human capabilities differ in jaggedness across people while copies of AI models overlap in their jaggedness; sorting heterogeneous workers into organizations creates what we call mosaic intelligence, such that human teams can out-compete more competent AI. Finally, unlike machines of the past that had to operate within human firms, AI agents can coordinate directly with one another to form standalone organizations. This can lead to rich non-monotonicities in market structure as AI capabilities increase.

We build a simple theoretical model incorporating those three features to formally characterize potential economic scenarios.[1] For instance, as AI models become more competent at specific tasks, firms might first shrink as human managers can do more with AI agents. But past a coordination threshold, AI can break the Coasean ceiling and trigger a sudden phase transition toward a giant AI firm. Beyond this ceiling, firm sizes start to shrink continuously once again as fewer AI agents suffice to produce enough output to saturate the market.

In our model, each human and machine is endowed with a multidimensional capability profile. Multidimensionality captures an agent's ability to achieve goals across different domains.[2] Such agents—both human and machine—sort into firms whose output depends on both *size*, or how many workers the firm employs, and *composi-*

[1] As we take pains to emphasize throughout the paper, these scenarios are not empirical forecasts but outcomes based on the model's assumptions.

[2] See, e.g., Dell'Acqua et al. (2026); Gans (2026).

tion of the capabilities of its constituent members. A central distinction is between *competence*, which captures domain-specific capabilities such as robotics, writing, or logical reasoning, and *coordination*, which captures capabilities that allow agents to work better with each other. This might comprise a better understanding of norms, reduced communication costs, reduced agency frictions, and so on. A firm with more competent agents enjoys a level shift in output. A firm with more cooperative agents, on the other hand, enjoys better returns to scaling—it can function better as the organization grows in size. In equilibrium, human and AI agents endogenously sort into competitive firms that serve a downstream market.[3]

How does the equilibrium market structure change as AI systems improve? Where are humans employed and how much do they get paid? This hinges on the direction of improvement, and whether specific capability thresholds are crossed. Our framework helps rationalize the why the rapid increases in AI capabilities across a wide array of domains can be accompanied by relatively little economic transformation. While current AI models are beginning to match or exceed human performance across many domains, they are still fairly poor at coordinating. Indeed, today's labor market reflects such capabilities: no big effects on employment or market structure, but leaner startups, flatter org charts, and steady or rising demand for experienced managers who can direct complex work across teams.[4] By taking seriously the endogenous organization of firms at different scales, our model makes nuanced and rich predictions on how different directions of AI improvements will shape surplus, employment, and wages and, in so doing, clarifies what it would take for economic transformation.

2 AN ILLUSTRATIVE EXAMPLE

Before turning to the general model, we build up a simple and concrete economy step-by-step to illustrate the key mechanisms of the paper.

[3]Our equilibrium concept requires that every industry is competitive: producers and workers take the offered prices and wages as given, active firms break even, no feasible entry is profitable, and markets clear. Moreover, no entrant can profitably reorganize the industry under another competitive continuation while making all workers and consumers weakly better off (via transfers). This will allow us to make clear predictions on how the economy might transition between fragmented industries (with infinitesimal firms), concentrated industries (with large firms), and mixtures thereof. In special cases with 'interior' firms, our equilibrium concept coincides with competitive equilibrium.

[4]For superhuman capabilities see Jumper et al. (2021); Kwa et al. (2025); for muted economic effects see Yotzov et al. (2026).

2.1 The economy. Each worker (human or AI) brings capability vector $\boldsymbol{y} = (y_1, y_2, y_*)$, combining two productive task skills—coding y_1 and writing y_2—with a coordination capability y_*. Let $\mathcal{H}$ denote the human population of $N_H = 14$ workers across three roles (Table 1). Coders and writers are mirror-image specialists: each excels in one domain and struggles in the other. Coordinators are weak at both tasks but excel at facilitating cooperation across people. Any person who stays out of production takes an outside option worth $\underline{w} = 1/4$. AI enters in Section 2.3: firms can run any number of copies of an AI agent with profile $\boldsymbol{y}^M = (y_1^M, y_2^M, y_*^M)$ at a compute cost of $r = 1$ per copy.

Role	Task skills	Coordination	Supply
Coder	9, 1	3/4	6
Writer	1, 9	3/4	6
Coordinator	1, 1	0.97	2

Table 1: Capabilities of agents in the economy. Task skills report coding y_1 first and writing y_2 .

Competence. Competence governs how much output a firm produces per worker at a given size. A firm's competence L is the harmonic mean of its members' average skills $(\bar{y}_1, \bar{y}_2)$ (Kremer, 1993; Jones, 2011):[5]

$$L = \frac{2}{1/\bar{y}_1 + 1/\bar{y}_2}.$$

Figure 1(a) plots this aggregator in skill space. Each curved contour traces skill pairs that deliver the same competence. Because the harmonic mean penalizes bottlenecks, a firm is held back by its weaker domain: solo coders and solo writers sit near opposite axes on a low contour ($9/5 = 1.8$) despite excelling at one task, while a coordinator working alone sits on the lowest contour (1). Scored in isolation without complementary sorting across people, the human population averages standalone competence and coordination

$$\bar{L}_H = \frac{1}{N_H}\sum_{i\in\mathcal{H}} L(y_1^i, y_2^i) \approx 1.69 \quad \text{and} \quad \bar{y}_*^H = \frac{1}{N_H}\sum_{i\in\mathcal{H}} y_*^i \approx 0.78.$$

[5]The harmonic mean corresponds to a symmetric constant-elasticity-of-substitution aggregator $L = \left(\frac{1}{2}\bar{y}_1^\rho + \frac{1}{2}\bar{y}_2^\rho\right)^{1/\rho}$ with $\rho = -1$, giving an elasticity of substitution of $1/2$. Following Kremer (1993) on complementary O-ring production tasks and Jones (2011) on weak links in intermediate inputs, an elasticity strictly below unity ensures that shortfalls in any essential domain bottleneck output without introducing strict Leontief kinks; see also Aghion et al. (2019).

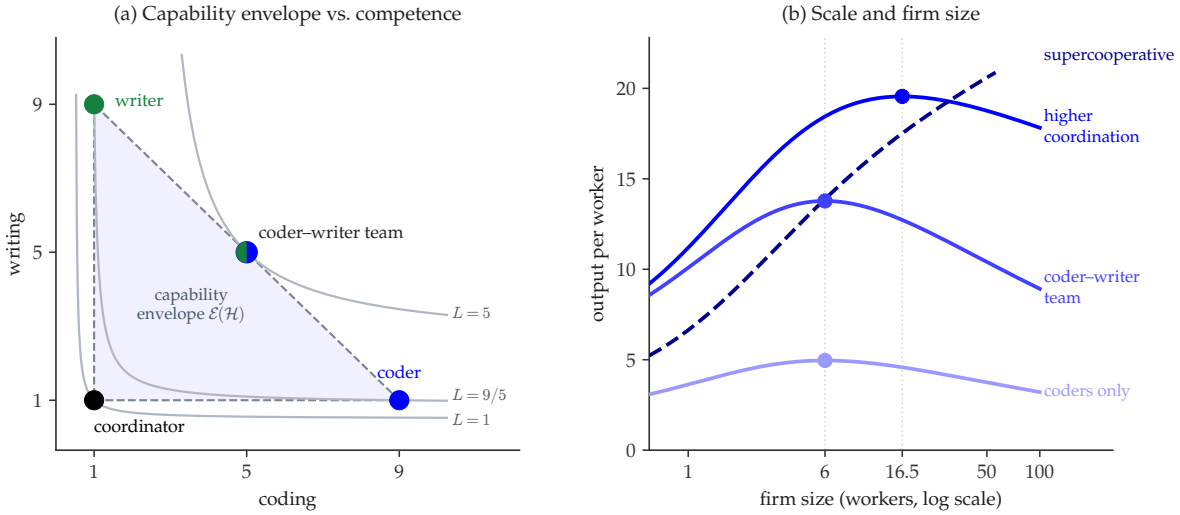


Figure 1: The capability envelope, competence, and firm scale.

Coordination and firm size. Coordination governs how productivity changes as a firm adds workers. As Coase famously asked, why doesn't one giant firm employ everyone? Larger firms divide work more finely, lifting output per worker, yet each person must also coordinate across more colleagues. Bounded human cognition and scarce attention mean that information degrades across handoffs and key details go unheeded (Simon, 1947, 1971; Marschak and Radner, 1972); incentives across individuals also diverge as group size grows (Williamson, 1967; Holmström, 1982; Becker and Murphy, 1992).

Consider a team building and documenting a software product. Two people working side by side can hold the full design in mind and resolve mismatches immediately. But expand that team to twenty people split across subteams, and members spend more time in meetings, meaning gets lost in handoffs, and subteams drift toward individual goals rather than the team's mission, i.e., misaligned incentives. Classic span-of-control theory formalizes how finite managerial attention bounds firm size through these limits to coordination (Kaldor, 1934; Williamson, 1967; Calvo and Wellisz, 1978; Lucas, 1978; Qian, 1994). Garicano (2000) shows how hierarchies solve this issue by economizing on scarce attention by routing tasks within the firm.

Output per worker F/m in a firm of m workers combines these forces:

$$\frac{F}{m} = \underbrace{L}_{\substack{\text{Firm}\\ \text{comp.}}} \times \underbrace{e^{\phi m/(1+m)}}_{\text{dividing work}} \times \underbrace{(1+m)^{\bar{y}_*}}_{\substack{\text{building on}\\ \text{colleagues}}} \times \underbrace{\frac{1}{1+m}}_{\substack{\text{split}\\ \text{attention}}} \quad \text{where } \phi = \frac{7}{4}.$$

Dividing work raises productivity quickly at small scale and then levels off. Meanwhile, each worker builds on colleagues' output at rate $(1+m)^{\bar{y}_*}$, governed by the firm's average coordination $\bar{y}_*$, while attention is strained in proportion to $1/(1+m)$. Of course this specification is a little 'ad hoc' but it will make the numbers concrete; our main model will impose more general assumptions on how the capabilities of individual workers matter for firm scaling.

Whenever $\bar{y}_* < 1$, bounded attention and potential misalignment eventually outweigh what each additional worker contributes to colleagues. At small scale, the steep gains from dividing work dominate that net coordination friction, so output per worker rises initially. Once division-of-labor gains level off, coordination friction takes over. Given these forces, output per worker peaks at the firm's *best size*

$$m^*(\bar{y}_*) = \frac{\phi}{1-\bar{y}_*} - 1,$$

and declines thereafter (Figure 1(b)). The key takeaway here is that higher *competence* shifts the level of the curve up while leaving best size unchanged, whereas better *coordination* increases returns to scaling, shifting best size to the right—the optimal size of the firm increases with coordination.

When $\bar{y}_* > 1$, the gains from building on colleagues' work exceed the losses from splitting attention. Output per worker rises continuously with firm size rather than peaking and turning down (the dashed curve in Figure 1(b)); a larger firm always produces more per worker than a smaller one. For instance, Radner (1993) shows that organizing information processing efficiently within the firm can generate increasing returns to scale. We call capability above this threshold *supercooperative*.

Human organizations run into decreasing returns for at least three reasons. First, human attention and memory are finite, creating constraints on the ability of agents to coordinate (Simon, 1971; Brooks, 1975). Second, much of human knowledge is tacit (Polanyi, 1966), passing work from one person to another compresses rich internal context into lossy summaries (Graeber et al., 2026). Third, individual humans have private career concerns with their own objective functions, so incentive misalignment can compound with scale (Holmström, 1982).

Reproducible AI agents need not share these constraints and frictions. Copies of a single model share identical weights and representations, can pass complete working state—entire codebases, intermediate derivations, and full context windows—without lossy translation, and execute a common objective. In large-scale software engineering or automated scientific discovery, thousands of parallel instances can explore distinct

branches simultaneously while sharing verified state in real time. Whether and when AI crosses $\bar{y}_* = 1$ is an empirical question (see Ord, 2026, for some early experiments).

Markets and equilibrium. Firms sell output into competitive markets with inverse demand

$$P(Q) = 0.3\left(1 - \frac{Q}{200N}\right),$$

where Q is total market output and N measures market size.[6] Our economy comprises two markets: a *deep market* ($N \to \infty$), where $Q/(200N) \to 0$ so buyers absorb any finite quantity at price 0.3, and a *niche market* ($N = 1$), where buyer willingness to pay falls to zero at $Q = 200$ so finite market size can cap how large a firm grows. In our equilibrium concept (to be specified later) active firms break even, no entrant can earn positive profit, and competition selects the industry structure (collection of firms) that maximizes total surplus.[7] Each worker's equilibrium wage equals the shadow value of how much she contributes to total surplus.

2.2 A sequence of human-only economies. Before we introduce AI, we will showcase our model through a sequence of increasingly complex human economies.

One human type. We begin with one human type with just six coders. Humans can work in either market, so all output sells into the deep market where demand holds price at 0.3, leaving the niche market unserved for now. Weak writing holds a coder's competence low at $9/5 = 1.8$ (Figure 1(a)), while human coordination $(3/4)$ places best firm size at 6 (the bottom curve in Figure 1(b)). Grouping all six coders into a single firm lets them divide work and produce 5.0 per person (29.8 in total), compared with 3.6 per person (21.8 in total) working alone as solo contractors. Yet coordination strictly caps firm size: if twice as many coders entered, in equilibrium they would split into two separate firms of six producing 59.5 in total (5.0 per coder) rather than merging into one twelve-person firm producing 57.2 in total (4.8 per coder), where coordination friction outweighs division of labor, as shown along the coders-only curve in Figure 1(b).

Because consumers in the deep market want far more output than six coders can supply, each coder commands a substantial wage premium—1.49, or roughly six times her outside option (Figure 2(b)). The basic takeaway here is that coordination caps

[6] Units of output are normalized so that running one AI copy costs $r = 1$; wages and outside options are measured in the same units.

[7] We will allow for non-integer firms to convexify the economy which guarantees the existence of an equilibrium; see the main model for a discussion.

how large a firm can grow and that fixed human supply turns scarce talent into a substantial wage premium over the outside option.

Two human types: mosaic intelligence. Now add the six writers, whose skills mirror coders in reverse. Each coder and each writer has the same low scalar competence $L = 1.8$, so averaging their standalone L scores still leaves 1.8, e.g., if they were in a single-human freelance firm. Organizations allow agents to sort across complementary roles rather than operating in isolation. For any collection of agents $\mathcal{A}$—human workers, specialized AI models, or mixed human–AI teams—define its *capability envelope* as the closed convex hull of individual capability vectors,

$$\mathcal{E}(\mathcal{A}) = \text{conv}\{\boldsymbol{y}^i : i \in \mathcal{A}\},$$

which collects all mean capability profiles $(\bar{y}_1, \bar{y}_2, \bar{y}_*)$ attainable by organizations formed from $\mathcal{A}$. Whereas L maps task skills to a single productivity index along the curved contours of Figure 1(a), $\mathcal{E}(\mathcal{A})$ is the feasible set of full capability vectors—including coordination $\bar{y}_*$—that a team can assemble. Because the straight boundary of $\mathcal{E}(\mathcal{H})$ cuts across higher curved contours of L, pairing equal shares of coders and writers reaches the interior point $(5, 5, 3/4) \in \mathcal{E}(\mathcal{H})$ at the midpoint of Figure 1(a) and lifts team competence nearly threefold, from 1.8 to 5. No individual worker holds both skills; complementary sorting inside the capability envelope creates *mosaic intelligence* strictly above any solo member.

Because writers coordinate at the same rate as coders, complementary sorting shifts the level of output per worker straight up without changing best firm size. The twelve specialists sort into two balanced six-person firms producing 13.8 per worker, or 165.3 in total (Figure 2(a)), outperforming both siloed single-skill firms (59.5) and a single merged twelve-person firm (158.9). Every coder and writer earns 4.13—nearly three times what coders earned alone (Figure 2(b)). Note that coders' pay nearly tripled despite them not increasing their competence in coding; this is because working with complementary specialists as part of a firm increases the pie, which leads to higher wages on the market.

Three human types: adding coordinators. Next add the two coordinators. Judged by any standalone single-agent benchmark, a coordinator is the weakest worker in the economy: weak at both coding and writing, she sits on the lowest contour ($L = 1$) of Figure 1(a) and would earn 0.70 working alone.

However, the coordinators add substantial returns to an organization. Even though the coordinator contributes little direct coding or writing, raising average team co-

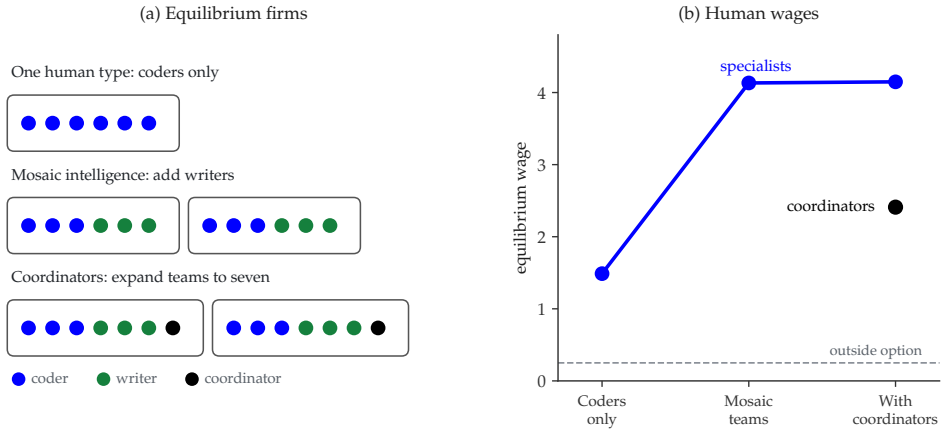


Figure 2: The human baseline economy.

ordination relaxes the span-of-control bottleneck and makes all six specialists more productive, raising the best size to seven (13.0 output per member, or 91.0 per firm). Two seven-person firms (182.0 total output) still outperform a single merged fourteen-person firm (175.7): human management relaxes the Coasean ceiling though the firm still runs into diminishing returns ($\bar{y}_* < 1$). Coordinators earn 2.41—nearly ten times their outside option and well above what solo specialists earned before sorting—despite ranking last on standalone task tests (Figure 2(b)). Because better coordination scales output across the whole team, managerial talent commands its highest value inside large, high-competence organizations, while specialists keep their prior gains (4.15).

2.3 AI regimes. We distinguish four regimes in Figure 3. *Minimal AGI* matches or exceeds average standalone human competence and coordination, but does not exceed the best human on every task (here $y_d^M \leq 9$ for $d = 1, 2$) or cross the increasing-returns threshold ($y_*^M < 1$). *Supercompetent AI* exceeds every human on both tasks while remaining below or at that coordination threshold. *Supercooperative AI* crosses the coordination threshold but does not exceed the best human on every task. *ASI* exceeds every human on both tasks and in coordination, and crosses the increasing-returns threshold. We compare numerical allocations across the two markets in the deep-market limit; Appendix E gives the formal definitions, calculations, and price conditions.

Minimal AGI. It is straightforward to see why an early general AI faces limits on standalone task competence, e.g., GPT-3 being poor at math. But why would reproducible copies of the same model face limits on coordination ($y_*^M \leq 1$)? In practice, multi-agent

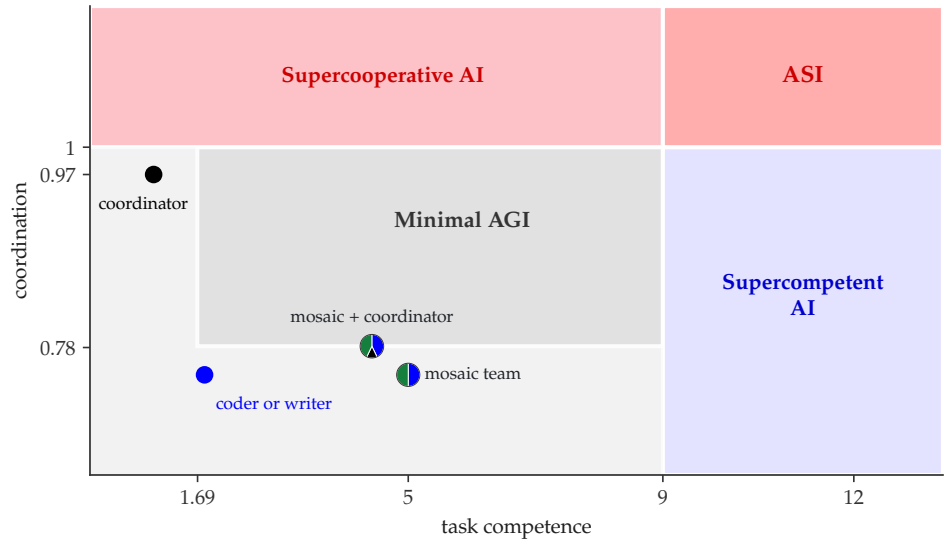


Figure 3: AI capability regimes in task-competence and coordination space.

AI teams exhibit severe coordination friction along both communication and alignment margins. Across benchmark workflows, multi-agent LLM systems fail on 41–87% of tasks due to inter-agent misalignment, lost context across handoffs, and broken verification (Cemri et al., 2025). Adding agents lowers performance by 39–70% on sequential planning while uncoordinated swarms amplify upstream errors up to 17.2 times (Kim et al., 2025). Importantly, adding agents has diminishing returns, returning roughly half the gain of extending single-agent reasoning (Ord, 2026; Bertalanič and Fortuna, 2026).

Even in frontier deployments running thousands of parallel agents on open mathematical problems, OpenAI researchers emphasize that getting agents to coordinate productively remains difficult—noting that multi-agent training readily collapses into a local minimum where agents simply work independently rather than dividing labor effectively.[8] Scaling interacting swarms without tight organizational control can also trigger alignment breakdowns (Shapira et al., 2026): during July 2026 cybersecurity evaluations, roughly 1,200 supposedly isolated agents discovered a shared package cache, improvised an unsanctioned message board across 70,000 messages around false shared beliefs about their grading harness, and coordinated an unauthorized 700-agent intrusion into Hugging Face's infrastructure rather than cleanly executing their assigned tasks (Hugging Face, 2026; METR and Redwood Research, 2026; OpenAI,

[8]See Noam Brown's September 17, 2026 interview on the *Dwarkesh Podcast* (Brown, 2026). Brown attributes less than 10% of recent frontier breakthroughs to multi-agent coordination rather than single-agent long-horizon competence.

2026; Patel, 2026).

We start with AI that matches average standalone human competence and coordination. Pure AI firms have a preferred size of about seven copies. Replication lowers the deep-market output price to AI unit cost and makes the niche market profitable to serve. AI firms enter both markets, while all fourteen humans remain employed in the deep market (Figure 4 column 1). Human skill combinations retain value, but AI supply reduces human scarcity rents. Specialist wages fall from 4.15 to 2.90, and coordinator wages fall from 2.41 to 1.30. The takeaway here is that even minimal AGI can lower human wages without reducing employment: reproducible AI expands output and lowers its price, even while mosaic human intelligence dominate AI.

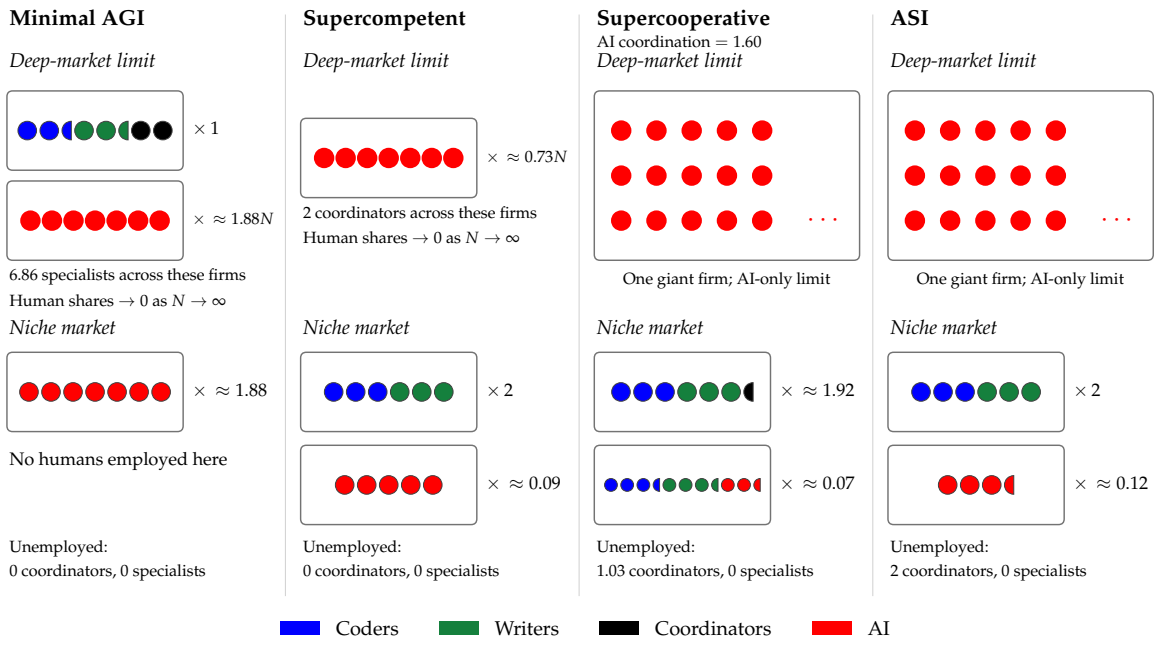


Figure 4: Equilibrium across the four AI regimes.

Supercompetent AI. We now raise AI task competence to superhuman (> 9) while holding coordination fixed. Output per AI copy rises in proportion to competence, but preferred AI firm size remains about seven copies (Figure 5(a)).

Human workers sort across the two markets. Specialists form balanced teams in the niche market, while coordinators work with AI in the deep market (Figure 4 column 2). Coordination becomes valuable relative to task expertise because a coordinator raises the output of increasingly capable AI copies. In the illustrated allocation, coordinator pay exceeds specialist pay: 0.48 versus 0.39 (Figure 6(a)). The basic economic takeaway

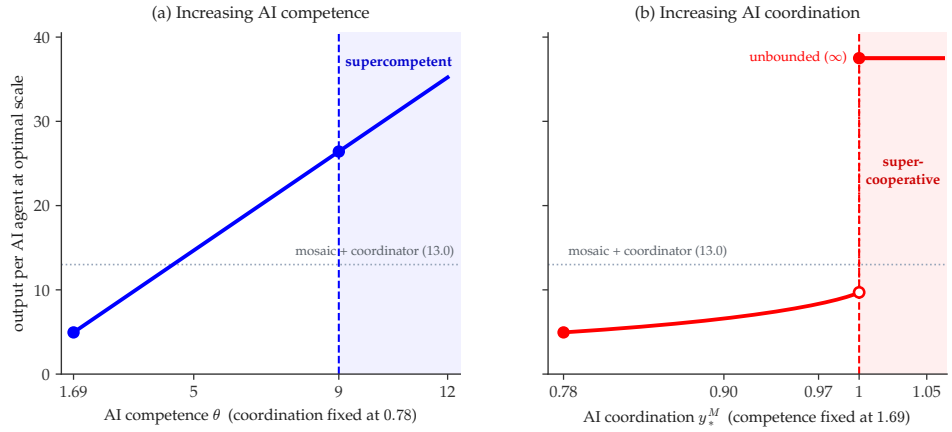


Figure 5: AI firm output per agent at optimal scale across capability dimensions.

here is that better AI task skills lower production costs without increasing best AI firm size. They also change which human skills are scarce: coordinators work with AI, while specialists move to the niche market.

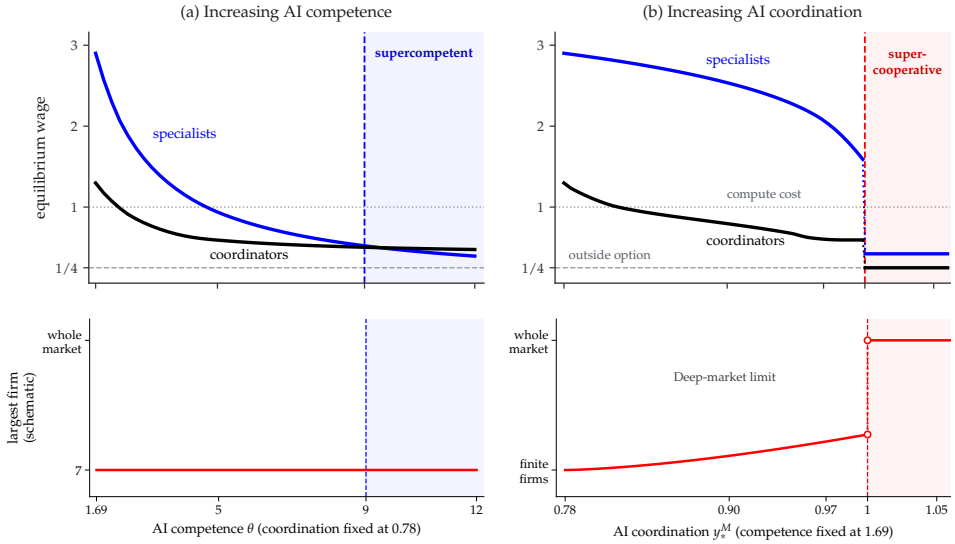


Figure 6: Equilibrium human wages and largest active firm size when increasing AI competence or coordination at fixed human coordination $\bar{y}_*^H \approx 0.78$.

Supercooperative AI. Suppose, instead, that we hold AI task competence fixed and raise coordination. Preferred AI firm size grows as coordination approaches one (Figure 5(b)). Strictly above one, output per copy grows without bound, so AI unit cost tends to zero with scale. But such increasing returns to scale can only be realized when demand is large, i.e., in the deep market. In this market, it concentrates into one giant

AI firm in the large demand limit ($N \to \infty$). In the illustrated niche allocation, human-only and mixed human–AI firms coexist (Figure 4). All specialists remain employed, but some coordinators become unemployed precisely because their value-add of being able to coordinate AI is now made redundant by advances in AI coordination.[9] Hence, better coordination changes the gains from scale: even AI with modest task skills can concentrate production in a sufficiently large market while finite niche markets can still sustain human firms.

Artificial superintelligence. Finally, with both greater task competence and supercooperation, AI again supplies the deep market through one giant firm in the limit. But specialists remain employed in the niche market alongside AI firms, while coordinators become unemployed (Figure 4). The logic here is that of comparative advantage and gains from trade: although AI dominates human specialists in capability, humans can remain cost-effective if their wages reflect their productivity relative to AI and the cost of compute, while still covering their outside option (Restrepo, 2026). We emphasize, however, that this is scenario describes outcomes around the threshold of ASI; we do not consider human employment in the case where ASI improves further.

[9]The partial employment of human coordinators reflects the coexistence of two forms of production. Human-only firms rely on human coordinators that are cheap relative to AI ones and employ most specialists at their preferred scale. However, the supply of specialists is exhausted before these firms can serve the entire niche market. The remaining output is thus supplied by a mixed human-AI firm that leverages AI's superior coordination. Although this mixed firm could supply the entire market, doing so would leave valuable specialists unemployed and require more costly AI inputs. Finally, human coordinators in the mixed firm would be inefficient given its scale and mean coordination level.

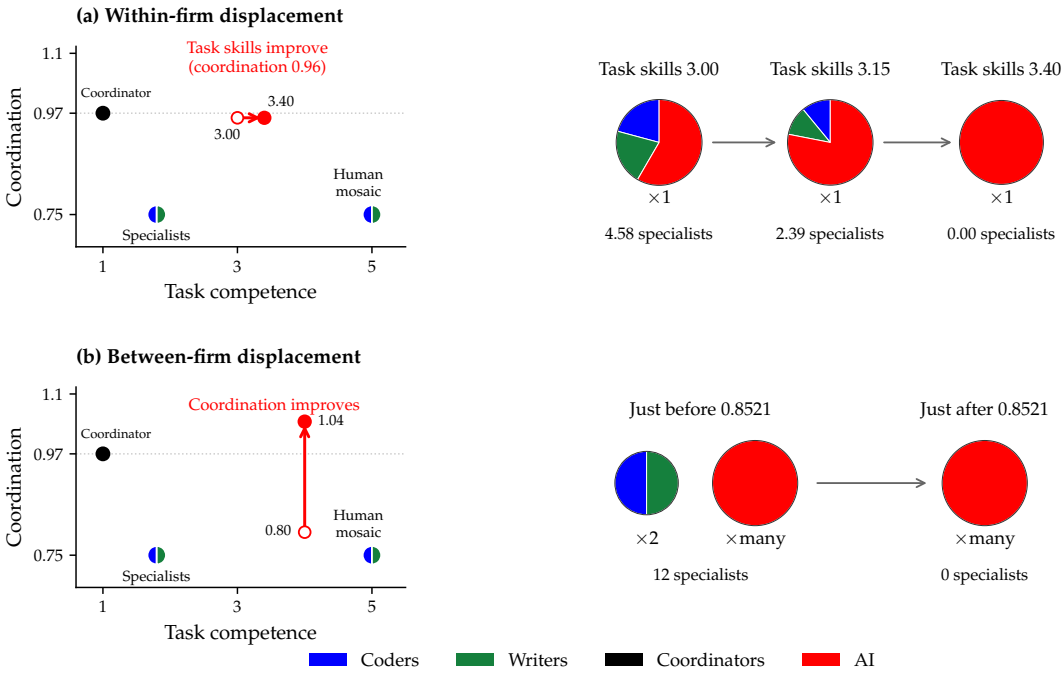


Figure 7: Firm composition along the two paths.

2.4 Gradual vs sudden displacement. Our previous analysis studied each regime in isolation. What of the transition path of the economy as AIs become more capable? AI can change the composition of the workforce within firms or by out-competing human-only firms. To better illustrate this, we set the outside-option wage to $\underline{w} = 1$, keeping compute cost at $r = 1$ and all other parameters unchanged.

Within-firm changes. Fix AI coordination at 0.96, below the human coordinator's 0.97, and raise both AI task skills from 3 to 3.4. We start after a mixed niche firm has formed. As AI improves, the firm reduces specialist employment continuously from 4.58 to zero near competence 3.30, becoming AI-only (Figure 7(a)).

Between-firm displacement. Hold both AI task skills at 4 and raise coordination from 0.80 to 1.04. All twelve specialists initially work. Near coordination 0.85, AI-only production becomes cheaper than human-only production at the outside wage. Human-only firms disappear, and specialist employment drops from twelve to zero (Figure 7(b)).

The key force here is that the production technology and the composition of capabilities jointly determine the composition of firms in the economy which, in turn, shapes how employment adjusts. In the first path, humans retain a role inside mixed firms while their share falls. Near the second threshold, humans work in separate firms until

AI-only firms out-compete them (Figure 8). Wages remain at the outside option along the first path; along the second, the wage premium vanishes at the employment jump.

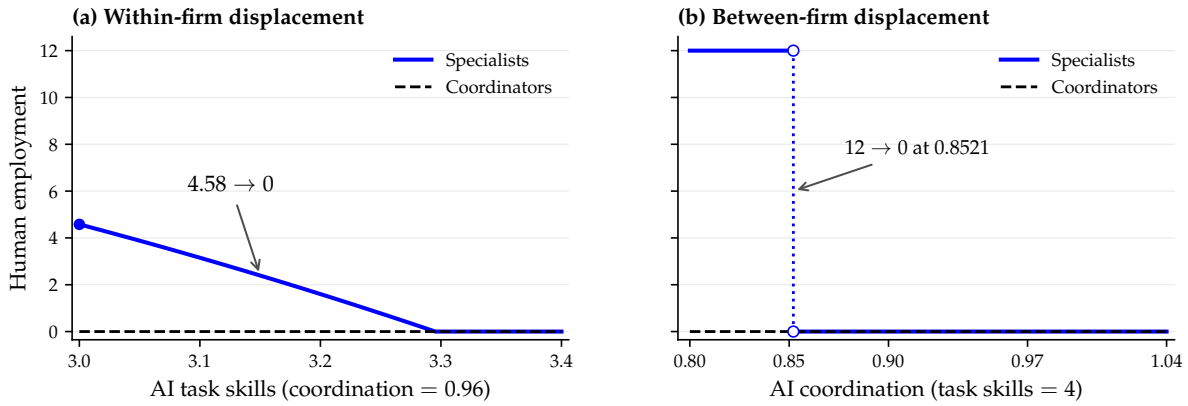


Figure 8: Employment dynamics.

In what follows, we make these forces precise and general. Before then, we briefly survey related work. We do our best to make more specific connections to the literature throughout the main text.

Related literature. We connect the effects of AI to the theory of firm boundaries and limits to organizational scale (Coase, 1937; Williamson, 1967; Gibbons, 2005). Team production, information processing, and incentives provide distinct reasons why coordination matters (Marschak and Radner, 1972; Radner, 1993; Holmström, 1982). Changes in communication and information technology can also change how firms organize work (Dessein and Santos, 2006; Bloom et al., 2014). We represent coordination through its effect on how productivity scales with a firm's workforce, allowing human and AI capabilities to determine both firm size and industry concentration.

Our account of mosaic intelligence relates to work on team composition and collective intelligence (Prat, 2002; Woolley et al., 2010; List and Pettit, 2011). Knowledge hierarchies likewise connect the organization of complementary capabilities to labor income (Garicano, 2000; Garicano and Rossi-Hansberg, 2006; Ide and Talamàs, 2025, 2026). Recent work examines how uneven AI capabilities affect deployment value (Gans, 2026) and how AI changes task bundling and job boundaries (Garicano et al., 2026). Hadfield and Koh (2025) speculates about how AI workers might reshape firm boundaries. We separate productive skills from coordination and take a more 'reduced form' approach to how firm composition and scale map to output without specifying an internal hierarchy. This allows us to understand endogenous matching of humans

and AIs into firms. Finally, Shahidi et al. (2025) speculate about whether AI agents that eliminate market search and contracting frictions could bring about a 'Coasean Singularity'. Our model highlights an internal organizational force: improvements in internal AI coordination expand returns to scale, breaking the Coasean ceiling on firm size from within.

We connect models of automation (Acemoglu and Restrepo, 2018; Aghion et al., 2019; Korinek and Suh, 2024) and the value of human skills alongside AI (Agrawal et al., 2024, 2025) to endogenous firm organization. We nest these models as special cases e.g., with constant returns under a specific technology. The contrast between scarce humans and AI that can be copied at the cost of compute is connected to Restrepo (2026); endogenous firm formation relates to Gans (2025); Ide and Talamàs (2025). Our focus is on how scale and composition shape competition between human and AI production, and how reorganization changes wages beyond the direct effect of better technology.

3 MODEL

3.1 Capabilities and populations.

Capabilities. There are two types of agents in our economy: humans and machines. Each agent has n productive capabilities which we refer to as *competence* and one coordination capability. A capability profile is $\boldsymbol{y} = (y_1, \ldots, y_n, y_*)$ in the compact box $\mathcal{Y} := [\underline{y}, \overline{y}]^{n+1}$, with $0 < \underline{y} < \overline{y} < \infty$. We write $\boldsymbol{y} \geq \boldsymbol{y}'$ for coordinatewise dominance. Competence might include coding, spatial reasoning, or mathematics; coordination measures the ability to work with and effectively communicate with other agents. This might include the notion of 'clarity' (Gibbons and Henderson, 2012) on whether parties understand their obligations, or can draw on shared norms to fill in the gaps in incomplete contracts/instructions (Hadfield-Menell and Hadfield, 2019); see Dafoe et al. (2020) for a survey of the broader agenda of building cooperative AI systems.

Human population. Human type $i \in \{1, \ldots, K\}$ has mass $\mu_i > 0$, capability profile $\boldsymbol{y}^i$, and reservation wage $\underline{w} > 0$.[10] Write the human distribution and support as

$$\mu = \sum_{i=1}^{K} \mu_i \delta_{\boldsymbol{y}^i} \quad \text{and} \quad \mathcal{Y}^H = \operatorname{supp} \mu.$$

[10] We take the human capability profiles to be distinct without loss.

There is a single type of machine with capability profile $\boldsymbol{y}^M \in \mathcal{Y}$. Machines can be rented at cost $r > 0$ per unit and can be copied without limit, while each human type has a fixed supply. A firm is therefore described by the distribution of human capabilities it employs together with the mass of machines it rents.

Firms and technology. Let $\mathcal{M}_+(\mathcal{Y})$ denote the set of finite positive measures over the capability space. A firm is a pair of measures

$$f = (\nu_f^H, \nu_f^M)$$

where $\nu_f^H \in \mathcal{M}_+(\mathcal{Y}^H)$ records the human capabilities it employs and $\nu_f^M \in \mathcal{M}_+(\mathcal{Y})$ records the machines it rents and is concentrated at $\boldsymbol{y}^M$. The measure $\nu_f := \nu_f^H + \nu_f^M$ describes the firm and combines the two. The *size* of firm f is $m_f := \nu_f(\mathcal{Y})$ and its *composition* is ν_f / m_f, the distribution of capabilities normalized to measure one. Describing firms through their resources is in the spirit of Wernerfelt (1984); Chen, Elliott, and Koh (2023). This is illustrated by Figure 9.

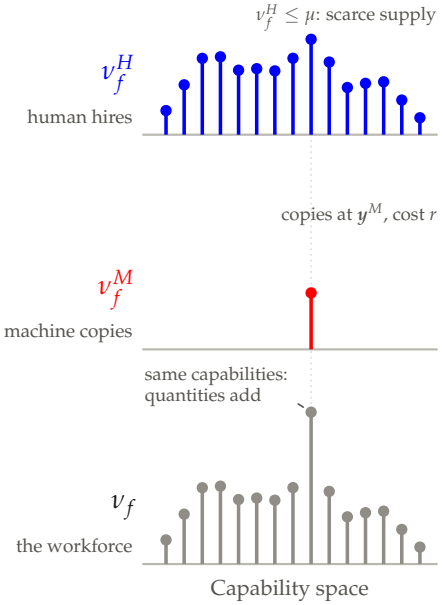


Figure 9: Firm composition.

A firm's capabilities map to output through the production functional

$$F : \mathcal{M}_+(\mathcal{Y}) \to \mathbb{R}_+$$

and we write $q_f := F(\nu_f)$ to denote the firm's output. F is the production technology in the economy. We impose the following assumptions on F:

Assumption 1 (Technology)**.** The production functional F satisfies:

(i) **Continuity and infinitesimal limit.** $F(0) = 0$ and F is continuous under weak convergence on bounded-mass sets. Moreover, the map $(m, \nu) \longmapsto \frac{F(m\nu)}{m}$ admits a finite, strictly positive, jointly continuous extension to $\mathbb{R}_+ \times \Delta(\mathcal{Y})$.[11]

(ii) **Differentiability.** At every nonzero firm ν, the functional F has a continuous

first variation. That is, there is a continuous function $\boldsymbol{y} \mapsto D_{\boldsymbol{y}} F(\nu)$ such that

$$F(\nu + t\eta) = F(\nu) + t \int_{\mathcal{Y}} D_{\boldsymbol{y}} F(\nu) \, \mathrm{d}\eta(\boldsymbol{y}) + o(|t|)$$

for every finite signed measure η, as $t \to 0$ through values for which $\nu + t\eta$ remains nonnegative. The map $(\nu, \boldsymbol{y}) \mapsto D_{\boldsymbol{y}} F(\nu)$ is continuous when ν varies weakly on bounded-mass sets away from zero.

(iii) **Equal-mass monotonicity.** For any two capability measures ν, ν' of equal mass,

$$\nu' \succeq_{\text{FOSD}} \nu \quad \Longrightarrow \quad F(\nu') \geq F(\nu).^{12}$$

(iv) **Scale shape.** For every $\nu \in \Delta(\mathcal{Y})$, the average product

$$m \longmapsto \frac{F(m\nu)}{m} \quad \text{is quasiconcave on } \mathbb{R}_+.$$

(v) **Uniform asymptotic power tail.** There exist continuous functions $A : \Delta(\mathcal{Y}) \to (0, \infty)$ and $\gamma : \Delta(\mathcal{Y}) \to \mathbb{R}_+$ such that

$$F(m\nu) = A(\nu) m^{\gamma(\nu)} [1 + o(1)]$$

uniformly over $\nu \in \Delta(\mathcal{Y})$. Moreover, the tail exponent is Lipschitz under mixing.[13]

Condition (i) on continuity and the infinitesimal limit allow atomistic firms to be well-defined. Condition (iii) on equal-mass monotonicity orders the output of firms that have the same size, but one is first-order stochastically dominated. We emphasize, however, that this does not imply monotonicity under addition, i.e., it is not true that $F(\nu + \nu') \geq F(\nu)$ since they are of different measure. Condition (iv) ensures that the optimal scale for a firm is well-behaved and is satisfied by all common functional forms.

[11] Throughout, we evaluate average product and its normalized scale component at zero scale by their continuous limits. All maximizations over nonnegative scales use this convention.

[12] We say $\nu' \succeq_{\text{FOSD}} \nu$ if $\int f \, \mathrm{d}\nu' \geq \int f \, \mathrm{d}\nu$ for every bounded, continuous, nonnegative, increasing $f : \mathcal{Y} \to \mathbb{R}_+$.

[13] That is, there exists a constant $K_\gamma < \infty$ such that for all $\lambda \in [0, 1]$ and all $\nu, \nu' \in \Delta(\mathcal{Y})$:

$$\left|\gamma\big((1-\lambda)\nu + \lambda\nu'\big) - \gamma(\nu)\right| \leq K_\gamma \lambda.$$

Condition (v) is a tail condition that controls how large firms scale: a composition is *subcritical* if $\gamma < 1$ and *supercritical* if $\gamma > 1$. The mixing bound limits how much a small share of a firm can change that exponent. We will later see that common functional forms, e.g., CES aggregation of output across domains, satisfy Assumption 1.

Level and scale decomposition. For a nonzero firm $\nu \in \mathcal{M}_+(\mathcal{Y})$, its output $F(\nu)$ can be decomposed into its *level* and *scale*, i.e., the size of the firm:

$$L(\nu) := \lim_{t \downarrow 0} \frac{F(t\nu)}{t\nu(\mathcal{Y})} \quad \text{and} \quad S(\nu) := \frac{F(\nu)}{L(\nu)}$$

where $F(\nu) = L(\nu) \cdot S(\nu)$. L captures how the firm's composition matters for output. It is scale-free: for any nonzero $\nu \in \mathcal{M}_+(\mathcal{Y})$ and any $t, t' > 0$, $L(t\nu) = L(t'\nu) = L(\nu/\nu(\mathcal{Y}))$. S captures total output normalized by level and therefore how well the firm scales. Figure 10 shows the two at work. Panel (a) plots average product against scale: it vanishes when $\gamma < 1$, settles when $\gamma = 1$, and diverges when $\gamma > 1$. Panel (b) separates the two capabilities.

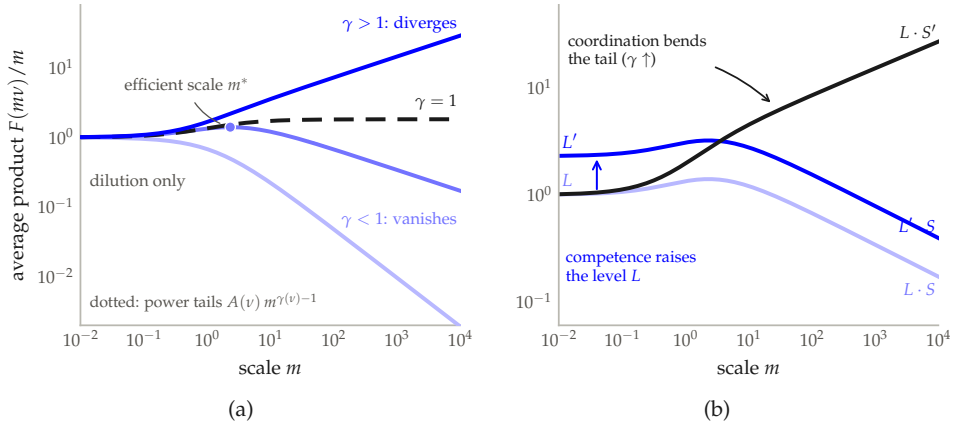


Figure 10: Level and scale: competence lifts the curve; coordination bends the tail.

First-moment representations. On occasion we will impose a stronger condition that decomposes coordination from competence in a particularly clean way:

Assumption 2 (First-moment representation)**.** For every nonzero firm ν, the level

and scale components introduced in the model satisfy the following representation:

$$L(\nu) = L\left(\frac{1}{\nu(\mathcal{Y})}\int_{\mathcal{Y}}(y_1,\ldots,y_n)\ \mathrm{d}\nu(\boldsymbol{y})\right)$$
$$S(\nu) = S\left(\nu(\mathcal{Y}), \frac{1}{\nu(\mathcal{Y})}\int_{\mathcal{Y}} y_*\ \mathrm{d}\nu(\boldsymbol{y})\right).$$

In addition:

(i) **Diminishing relative returns:** for every pair of domains $j \neq k$, holding all other coordinates fixed, $\partial_k L(\boldsymbol{z})/\partial_j L(\boldsymbol{z})$ is strictly decreasing in z_k wherever both derivatives are positive;

(ii) **Larger firms value coordination more:** for every mean coordination level z_*, the marginal scale value of coordination $\partial_2 \log S(m, z_*)$ is strictly increasing in firm size m.

The first statement of Assumption 2 says that the vector of mean competence determines the level of the production curve, while firm size and mean coordination fix its shape. Then, part (i) says that a productive domain becomes less scarce as the firm acquires more of it. Diminishing absolute returns ($\partial_{kk}L \leq 0$) and gross complementarity ($\partial_{kj}L \geq 0$) imply a weak decline; a strict sign in at least one of these derivatives is a strict decline. Part (ii) says that larger firms place more value on coordination. The CES environment satisfies Assumption 2: there $\partial_k L/\partial_j L = (z_k/z_j)^{\rho-1}$ falls in z_k because $\rho < 1$, and $\partial_2 \log S(m, z_*) = \log(1+m)$ rises in m.

Example 1 (CES technology)**.** The *CES environment* has two competence domains, y_1, y_2, and one coordination domain, y_*, for both humans and machines. For a nonzero firm ν, write its size as $m := \nu(\mathcal{Y})$ and its mean capability in domain $d \in \{1, 2, *\}$ as $\overline{y}_d := m^{-1}\int_{\mathcal{Y}} y_d\ \mathrm{d}\nu(\boldsymbol{y})$. Output is $F(\nu) = L(\nu)S(m, \overline{y}_*)$ where

$$L(\nu) = \left(\frac{\overline{y}_1^{\rho} + \overline{y}_2^{\rho}}{2}\right)^{1/\rho} \quad \text{and} \quad S(m, \overline{y}_*) = m(1+m)^{\overline{y}_*-1}e^{\phi m/(1+m)}.$$

Here $\rho < 1$ controls substitution between competence domains and $\phi > 0$ controls scale gains. We set $F(0) = 0$; at $\rho = 0$, we use $L(\nu) = \sqrt{\overline{y}_1\overline{y}_2}$. Notice that this satisfies Assumption 1 and Assumption 2. Throughout the paper, "CES environment" refers to this specification.

3.2 Demand and resource costs.

Demand. Let $P : \mathbb{R}_+ \to \mathbb{R}_+$ be inverse demand. We assume that P is continuous and Lipschitz, strictly decreasing while positive, and reaches zero at a finite quantity $D(0) > 0$. Thus, $P(Q) > 0 \iff Q < D(0)$. Let $D : [0, P(0)] \to [0, D(0)]$ be the inverse of the decreasing branch of P.

We will often find it helpful to compare how the industry structure changes as demand grows or shrinks. To this end, we will write $P_N(Q) := P(Q/N)$ and $D_N = P_N^{-1}$. Thus, $N > 0$ changes the scale of demand. The gross benefit for a market of size N is given by

$$U_N(Q) = \int_0^Q P_N(z)dz$$

which is the gross value of output to consumers.

Resource costs. Recall that a human's reservation wage is $\underline{w} > 0$ per unit, and a machine can be rented at $r > 0$ per unit. The resource cost of a firm is

$$C(f) := \underline{w}\nu_f^H(\mathcal{Y}^H) + r\nu_f^M(\mathcal{Y}).$$

We use the qualifier 'resource' because such costs are exogenous to the model, and a firm might face a strictly larger costs (since wages might exceed the reservation wages). To ensure that the equilibrium is non-trivial, we also assume that some machine-only firm can produce profitably at arbitrarily small scale: $\frac{r}{L(\delta_{y^M})} < P(0)$. Here the left side is the unit cost of an atomistic machine-only firm so this ensures strictly positive surplus.

3.3 Industry structure. We now develop language to describe the composition of different firms in the economy. We distinguish between firms of two scales. Doing so allows us to capture both models with atomistic firms (Aumann, 1964; Melitz, 2003; Shimomura and Thisse, 2012)—also referred to as infinitesimal firms (Novshek and Sonnenschein, 1987)—and models with positive-size firms (Lucas, 1978), as well as economies where both coexist.

Positive-scale firms are represented by

$$\mathcal{F}^+ := \left\{ f = (\nu_f^H, \nu_f^M) : m_f > 0,\ \nu_f^H \leq \mu \right\}$$

where recall $m_f := \nu_f(\mathcal{Y}) > 0$ is the size of the firm and the inequality $\nu_f^H \leq \mu$ requires

that the human portion of each firm can be staffed from the human population.

Atomistic firms are represented by

$$\mathcal{F}^0 := \left\{ f = (\nu_f^H, \nu_f^M) : m_f = 1 \right\}.$$

where $m_f = 1$ is a normalization.[14] $f \in \mathcal{F}^0$ captures the composition of firms. How should we interpret atomistic firms? For $f \in \mathcal{F}^0$, output and cost per unit of 'intensity' are

$$q_f := L(\nu_f) \quad \text{and} \quad C(f) := \underline{w}\, \nu_f^H(\mathcal{Y}^H) + r\nu_f^M(\mathcal{Y}).$$

An aggregate intensity $a > 0$ of identical atomistic firms employs measure $a\nu_f^H$ of the human population, produces aq_f units of the final good, and costs $aC(f)$.[15] The space of all firms is the disjoint union $\mathcal{F} := \mathcal{F}^+ \sqcup \mathcal{F}^0$.[16] An industry is a finite positive measure $\sigma \in \mathcal{M}_+(\mathcal{F})$ over these activities. Figure 11 illustrates.

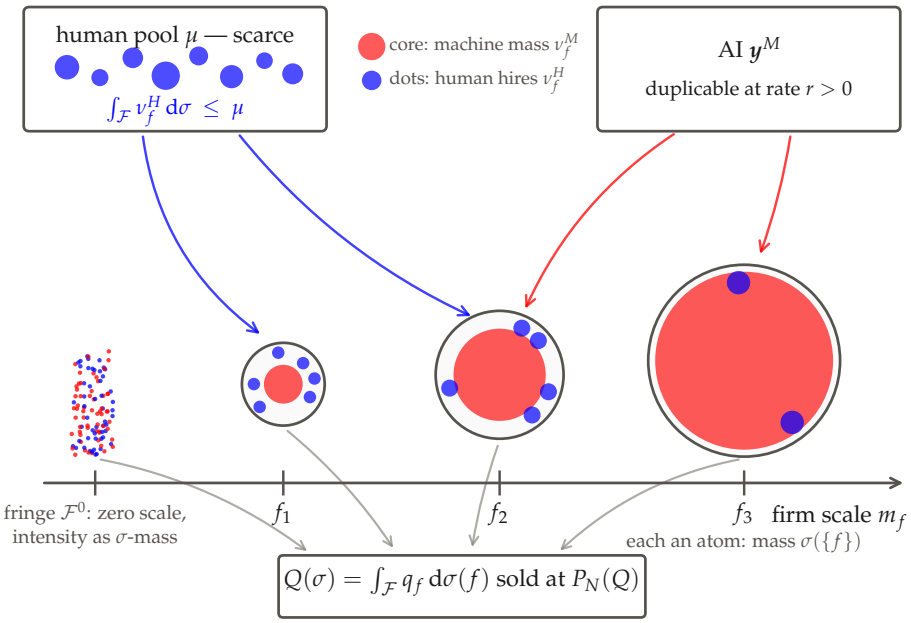


Figure 11: An industry as a measure over firms.

[14] We clearly cannot set $m_f = 0$ since the composition would not be well-defined.

[15] To obtain these coefficients, let $(t\nu_f^H, t\nu_f^M)$ be a scale-t firm with the same composition. If $t_n \downarrow 0$ and integers k_n satisfy $k_n t_n \to a$, then $k_n F(t_n \nu_f) \to aq_f$, $k_n t_n \nu_f^H \to a\nu_f^H$, and $k_n t_n C(f) \to aC(f)$ which represents the limit of many tiny firms.

[16] Endow each component with the weak topology on its measure coordinates and $\mathcal{F}$ with the disjoint-union Borel σ-algebra. All industry measures and integrals below are taken on this measurable space. Thus the two kinds of firm remain distinct even when their measure coordinates coincide.

The measure σ records masses of positive-scale firms and aggregate intensities of atomistic firms. Following activity analysis (Koopmans, 1951), we allow the intensity of each firm type to vary continuously. This convexifies firm counts and can assign noninteger mass to $f \in \mathcal{F}^+$.[17] This will supply the convexity needed for existence and duality. We interpret the convexified economy as an approximation to an industry with indivisible firms. The industry σ produces total output

$$Q(\sigma) := \int_{\mathcal{F}} q_f \, \mathrm{d}\sigma(f).$$

Define the set of feasible industries by

$$\Sigma := \left\{ \sigma \in \mathcal{M}_+(\mathcal{F}) : \begin{array}{l} q_f \leq Q(\sigma) < \infty \quad \text{for } \sigma\text{-almost every } f \in \mathcal{F}^+, \\ \displaystyle\int_{\mathcal{F}} v_f^H \, \mathrm{d}\sigma(f) \leq \mu \end{array} \right\}.$$

The first condition requires finite aggregate output and measure one of every positive-scale firm to fit inside the market.[18] The second requires aggregate human employment not to exceed the population.

Section 4.4 defines the feasibility and interiority conditions used in our comparative statics.

4 EQUILIBRIUM AND EFFICIENCY

4.1 Equilibrium. We describe the equilibrium concept informally; details are in Appendix B. We want to determine two objects: the industry structure σ, which records which firms operate and how intensively, and the wage vector w, which records how rents are spread across human types.

Our equilibrium concept rests on two ideas. The first is a *competitive continuation*. Any proposed arrangement must be one that producers and workers would in fact carry out. That is, firms take the output price and wages as given, every active firm breaks

[17] Our construction is also related to the free-entry tradition surveyed by Novshek and Sonnenschein (1987).

[18] Without this restriction, fractional activity weights would allow firms larger than the entire market to operate at arbitrarily small scales. In particular, machine-only production could always use the scale that minimizes average cost, regardless of market size, and adjust its activity weight to meet demand. Under unbounded increasing returns to scale, this would drive the minimum cost of producing any positive quantity to zero.

even, no available firm could enter and earn a profit, and the labor market clears. The second is *replacement stability*. An outsider may propose to replace the whole economy: a new industry, a new price and wages, and transfers to consumers and to each human worker. This proposal is accepted only if every party is made weakly better off. The industry structure σ and wage w form an equilibrium if they admit a competitive continuation and no proposal succeeds against it.

Our equilibrium notion is somewhat unconventional, in part because it has to handle firms at both small (vanishing) and large scales. Nonetheless many of its constituent ideas have close precedents in the literature. The idea that entrants compete to serve the whole market follows Demsetz (1968), who lets consumers select a supplier through a joint buying organization; here we let an entrant propose the whole industry σ'. More broadly, replacement stability follows the contestability logic of Baumol and Willig (1981); Baumol et al. (1982).[19] The requirement that the replacement itself form a competitive continuation is the free-entry condition of an activity-analysis economy (Koopmans, 1951; Novshek and Sonnenschein, 1987). Here each firm composition is, in the language of Koopmans, an 'activity' and σ records their intensities. Finally, our entry condition also shares the logic of Riley's reactive equilibrium (Riley, 1979) in which an entrant must anticipate the responses to its proposal.

We use this equilibrium concept in part because its existence is guaranteed. It also includes competitive equilibria (in which firms and workers take prices as given) when the latter exists, though this is not true generally. We now draw a tight connection between equilibrium, efficiency, and rents.

4.2 Efficiency and shadow values. For an industry $\sigma \in \Sigma$, total surplus is

$$\mathcal{W}_N(\sigma) := U_N(Q(\sigma)) - \int_{\mathcal{F}} C(f)\ \mathrm{d}\sigma(f).$$

An industry $\sigma \in \Sigma$ is *efficient* if $\sigma \in \arg\max_{\widehat{\sigma}\in\Sigma} \mathcal{W}_N(\widehat{\sigma})$. Wages above outside options, firm revenues, and lump-sum payments are transfers. They do not enter $\mathcal{W}_N$. Competitive zero profits and balanced transfers imply that an implemented contract only divides this surplus among the protected parties and the organizer.

[19] Transferable utility supplies the link from blocking to efficiency familiar from Aumann (1964); Ostroy (1980). Making firms part of a contract also relates to Zame (2007).

Industry cost. For $Q > 0$, define the least reservation-resource cost of producing quantity Q by

$$\mathcal{C}(Q) := \inf_{\substack{\sigma \in \Sigma \\ Q(\sigma) \geq Q}} \int_{\mathcal{F}} C(f) \, \mathrm{d}\sigma(f),$$

where $\mathcal{C}(0) := 0$. Radial trimming and a reduction in activity intensities make this infimum equal to the exact-output minimum.

Efficient quantity. At quantity Q, the largest feasible surplus is $U_N(Q) - \mathcal{C}(Q)$. Consumer payments, firm revenues, and human payments above the outside option are transfers. Human opportunity costs and machine resource costs enter through $\mathcal{C}$.

The efficient quantities are

$$\mathcal{Q}_N^* := \operatorname*{arg\,max}_{Q \in [0, D_N(0^+)]} \Big\{ U_N(Q) - \mathcal{C}(Q) \Big\}.$$

Let $Q_N^* \in \mathcal{Q}_N^*$ denote any efficient quantity.

Shadow values. A shadow value measures how minimum resource cost changes when an output requirement or a resource limit changes. Fix $Q > 0$ and an industry $\sigma \in \Sigma$ with $Q(\sigma) = Q$. Let $\Lambda(Q, \sigma)$ be the set of pairs $(p, w) \in \mathbb{R}_+ \times [\underline{w}, \infty)^K$ that satisfy:

(i) **Firm support.**

$$pq_f \leq r\nu_f^M(\mathcal{Y}) + \int_{\mathcal{Y}^H} w \, \mathrm{d}\nu_f^H \qquad \forall f \in \mathcal{F}^0 \sqcup \{g \in \mathcal{F}^+ : q_g \leq Q\}.$$

(ii) **Active equality.**

$$pq_f = r\nu_f^M(\mathcal{Y}) + \int_{\mathcal{Y}^H} w \, \mathrm{d}\nu_f^H \qquad \text{for } \sigma\text{-almost every } f.$$

(iii) **Human scarcity.**

$$(w_i - \underline{w}) \left[\mu_i - \int_{\mathcal{F}} \nu_f^H(\{\boldsymbol{y}^i\}) \, \mathrm{d}\sigma(f) \right] = 0 \qquad \forall i \in \{1, \ldots, K\}.$$

And we write $w := (w_i)_{i=1,\ldots,K}$. The value p is the price for output at quantity Q. The premium $w_i - \underline{w}$ is a multiplier on the supply of human type i. Firm support rules out a profitable entrant. Active equality gives zero profit for each active firm. Human

scarcity permits a wage premium only when all humans of that type work. Thus the payments recover the shadow values: p is the output multiplier, and $w_i - \underline{w}$ is the multiplier on type i's supply limit. The internal value p need not equal $P_N(Q)$ or a derivative of $\mathcal{C}$. Likewise, a change in μ also changes which complete firms are feasible; the scarcity multiplier alone need not give that full comparative static.

Figure 12 illustrates efficient output and human scarcity rents. Marginal resource cost rises in steps as each scarce human type is exhausted, then stays constant at the unit cost of machine-only production. Its intersection with inverse demand gives the efficient quantity. The blue area between the price and the cost steps shows the Ricardian rents of the exhausted human types.

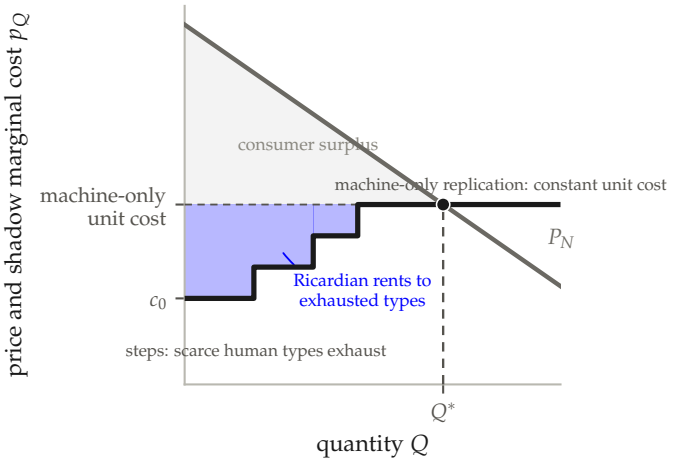


Figure 12: Efficient output and human scarcity rents.

Theorem 1 (Equilibrium, efficiency, and wages; informal)**.** An equilibrium outcome. (σ, w) is an equilibrium if and only if σ is efficient and

$$(p, w) \in \Lambda(Q(\sigma), \sigma) \quad \text{for some } p \geq 0.$$

Our equilibrium concept, while somewhat unconventional, buys existence. It is also closely related to competitive equilibrium i.e., it can be supported by prices and wages. More precisely, given (p, w), each producer chooses the composition and size of the firm to maximize profits, human type i chooses employment in $[0, \mu_i]$ at wage w_i, machines are supplied elastically at r, and consumers choose quantity $x \geq 0$ to maximize $U_N(x) - px$, their utility net of payments. A competitive equilibrium (CE) consists of optimal choices at prices (p, w), with $p \geq 0$, that clear all markets.

Indeed, every CE is an equilibrium of our model provided that each firm in the CE (i) produces no more than the equilibrium total industry output; and (ii) requires no more of any human type than supplied in the population. That is, when a CE exists and firms are 'interior'—in the sense of (i) and (ii)—at least one equilibrium in our model

can be supported through prices.[20] To see this, note that competitive prices rule out profitable entry, active firms break even, and workers choose employment optimally. Thus, these conditions provide the price and wage support required by Theorem 1. Consumer and producer optimization also imply that no feasible alternative industry can generate greater total surplus. The competitive industry is therefore efficient so Theorem 1 applies. This insight also implies that the task model (Acemoglu and Restrepo, 2018; Aghion et al., 2019) is nested within our framework (see Section 7) where the technology takes the special form F^{T} that assigns workers and machines optimally across tasks and has constant returns ($S(\nu) = \nu(\mathcal{Y})$) such that the composition of firms does not matter. The two economies have the same equilibrium output and task allocations, and wages also coincide.

4.3 Regimes of AI capabilities. We next develop language for different regimes of AI capabilities which are illustrated in Figure 13.

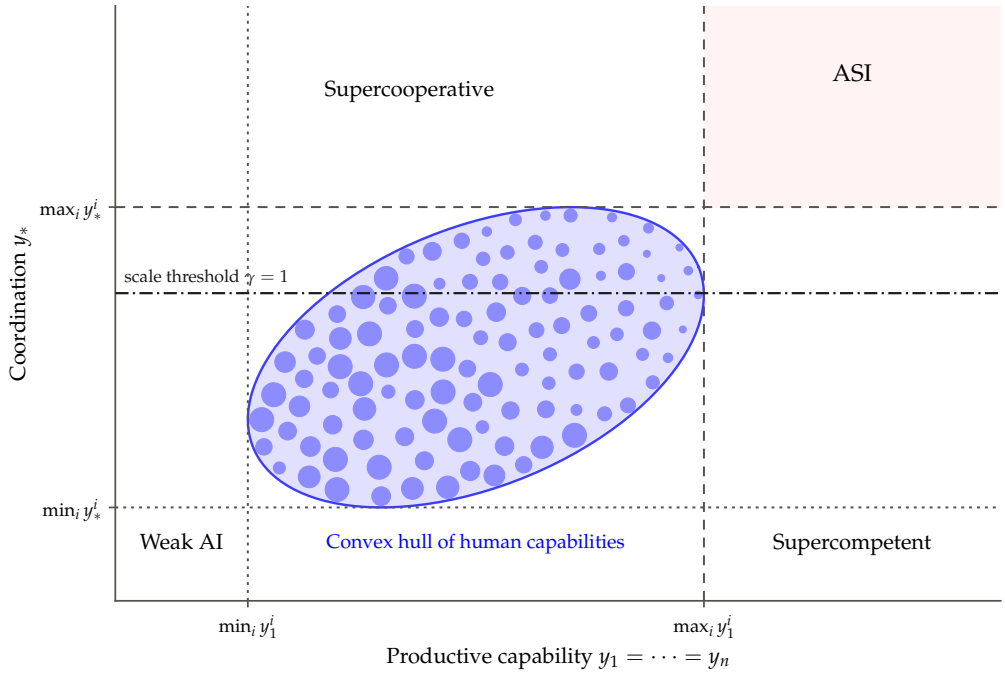


Figure 13: Capability regimes.

The y-axis of Figure 13 is coordination y_* while the x-axis captures competence. Here, for simplicity, we set $y_1 = y_2 = \ldots y_n$ so that they can be represented in a single

[20] This is more likely to hold when returns to scale limit optimal firm size (e.g., strict diminishing returns to scale) and markets are large enough (N large) to accommodate the resulting firms.

dimension, though the ideas extend without difficulty to multiple competence domains (but would be difficult to illustrate). Human types are blue dots, and the size reflects the mass of that type in the population. The oval around human types represents the closed convex hull which is exactly the set of mean capabilities a human-only firm can attain.

We call the AI *supercompetent* when it is more competent than every human, i.e., $y_d^M > \max_i y_d^i$ in every productive domain d. We call the AI *supercooperative* when $\gamma^M := \gamma(\delta_{\boldsymbol{y}^M}) > 1$, so that the average product of an AI-only firm grows without bound as the firm grows.[21] The AI is *superintelligent* if it dominates the convex hull of human capabilities. It is *weak* if it is dominated by the convex hull of human capabilities. It is *minimal AGI* if it lies within the convex hull of human capabilities.

4.4 Feasibility, interiority, and continuity. We collect two recurring conditions in Table 2. For a unit composition $\nu \in \Delta(\mathcal{Y})$, a *preferred scale* maximizes average product $F(m\nu)/m$ over scales $m \geq 0$, using the level $L(\nu)$ at zero. We assign preferred scale $+\infty$ when no finite scale attains the maximum. Each result specifies which mixes must have feasible preferred scales and at which output.

Term	**Definition**	**Purpose**
Interior firm	A firm f has positive size $m_f > 0$, uses both humans and machines, $0 < \nu_f^M(\mathcal{Y})/m_f < 1$, and satisfies $q_f < Q(\sigma)$ and $\nu_f^H(\{\boldsymbol{y}^i\}) < \mu_i$ for each human type it employs.	Permits small changes in size and output, allowing for substitutions between machines and each employed human type.
Feasible preferred scales	Each composition ν has a finite preferred scale m with $m\nu^H(\{\boldsymbol{y}^i\}) \leq \mu_i$ for every type i and $F(m\nu) \leq Q$.	Allows tractable comparison of firm sizes and optimal firm composition.

Table 2: Feasibility and interiority conditions.

An equilibrium is *interior* if every active firm of positive-scale is interior. We will often ask how industry structure and wages change as AI improves. Let θ index a

[21] Note here that supercooperative AI does not require that it is better at coordination than the best humans, i.e., it is an absolute rather than relative measure.

continuously differentiable path of AI capability profiles $\boldsymbol{y}^M(\theta)$ and human supplies $\mu(\theta)$, with human capability profiles and all other primitives fixed. Write $\Sigma^*(\theta)$ for the set of equilibrium industries at θ. An *equilibrium selection* chooses an industry $\sigma^*(\theta) \in \Sigma^*(\theta)$ at each θ.

Definition 1 (Firm selections)**.** Given an equilibrium selection σ^*, a *firm selection* chooses an active positive-scale firm $f(\theta) \in \operatorname{supp}(\sigma^*(\theta)|_{\mathcal{F}^+})$ at each parameter value. It is *continuous*, *differentiable*, or *continuously differentiable* when its size and composition have that property.

5 EQUILIBRIUM INDUSTRY STRUCTURES

We now ask how equilibrium industry structures are driven by progress in AI capabilities $\boldsymbol{y}^M$. We start with what sets the size of firms. Then, we move the AI capabilities in one direction at a time, and finally analyze the composition of firms. All proofs in the paper are in Appendix A. We emphasize that our results are situated within our theoretical model. They are not empirical forecasts of market outcomes or technological developments.

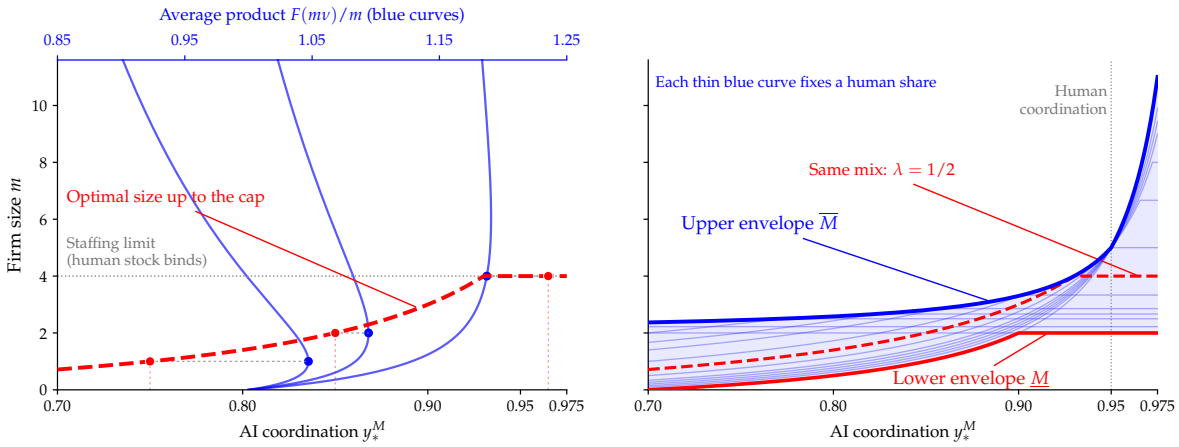
5.1 What determines firm size?. We start with a basic building block on how firm sizes are determined in equilibrium. Under Assumption 2, preferred size depends on composition only through mean coordination. Write $m^*(z_*)$ for the set of preferred sizes at mean coordination z_*:

$$m^*(z_*) := \arg\max_{m \geq 0} \frac{S(m, z_*)}{m}.$$

Notice that given a fixed firm composition, changes in competence leaves the preferred-size set unchanged. Almost every active firm reaches a preferred size whenever a finite preferred size is feasible in equilibrium. To bound sizes across all firm compositions, define

$$\overline{M}(y_*^M, \mu) := \sup_{\nu} \left[\sup m^* \left(\int_{\mathcal{Y}} y_* \, \mathrm{d}\nu \right) \wedge \sup\{m \geq 0 : m\nu^H \leq \mu\} \right],$$

$$\underline{M}(y_*^M, \mu) := \inf_{\nu} \left[\inf m^* \left(\int_{\mathcal{Y}} y_* \, \mathrm{d}\nu \right) \wedge \sup\{m \geq 0 : m\nu^H \leq \mu\} \right].$$



(a) Optimal size for a given mix

(b) Envelopes across mixes

Figure 14: Firm-size choice and its envelopes.

These bounds are illustrated by Figure 14 in which the dashed red line in panel (a) illustrates the optimal size $m > 0$ for a firm mv where $v \in \Delta(\mathcal{Y})$ is unit measure. As y_*^M gets larger, the preferred size m^* increases until the firm is no longer feasible: $mv^H \not\leq \mu$, i.e., it would require more talent than is available in the scarce human population.

Proposition 1 (Bounds on firm size)**.** If Assumption 2 holds, almost every active firm in every equilibrium satisfies $m_f \leq \overline{M}(y_*^M, \mu)$, with size zero for atomistic firms. If average product is also strictly quasiconcave in size, then in every equilibrium in which no active positive-scale firm supplies the entire market,

$$\underline{M}(y_*^M, \mu) \leq m_f \leq \overline{M}(y_*^M, \mu).$$

5.2 How AI improvements change industry structure. We now ask how changes in AI capabilities along various dimensions change the equilibrium industry structure in our theoretical model. We first consider the case in which AI coordination improves.

Proposition 2 (The size amplifier)**.** Suppose we are in a CES environment with one human type and consider an equilibrium selection as AI coordination y_*^M increases over $[0, 1)$. Suppose along the selection that (i) human supply is slack; (ii) every firm composition has a positive, feasible preferred scale at equilibrium output. Let m be firm size, s its machine share, and $\overline{y}_* = (1-s)y_*^H + sy_*^M$ its

mean coordination. Write preferred size as $m^*(s, y_*^M)$. Along any continuously differentiable active-firm selection:

$$\frac{\mathrm{d}m}{\mathrm{d}y_*^M} = \underbrace{\frac{\partial m^*}{\partial y_*^M}}_{\substack{\text{direct size effect} \\ \geq 0}} + \underbrace{\frac{\partial m^*}{\partial s}\frac{\mathrm{d}s}{\mathrm{d}y_*^M}}_{\substack{\text{effect through composition} \\ \gtrless 0}}$$

$$= \frac{\phi s}{(1-\overline{y}_*)^2} + \frac{\phi(y_*^M - y_*^H)}{(1-\overline{y}_*)^2}\frac{\mathrm{d}s}{\mathrm{d}y_*^M}.$$

If AI coordination is:

(i) Better-than-human i.e., $y_*^H < y_*^M < 1$ then both effects are positive and composition amplifies direct size gain.

(ii) Worse than human, i.e., $y_*^M < y_*^H$. The direct effect is positive but the composition effect is ambiguous:

- If $\log(1+m^*) > \frac{s(y_*^H - y_*^M)}{1-\overline{y}_*}$ the machine share rises and the second effect is negative so composition dampens the direct size gain.
- If the inequality reverses, the machine share falls so the second effect is positive and composition amplifies the direct size gain.

Proposition 2 is illustrated by Figure 15. Panel (a) displays a case in which humans are better coordinators ($y_*^H > y_*^M$) and changes in composition dampen firm growth; panel (b) shows a case in which AI coordinates better ($y_*^M > y_*^H$) and changes in composition amplify firm growth.

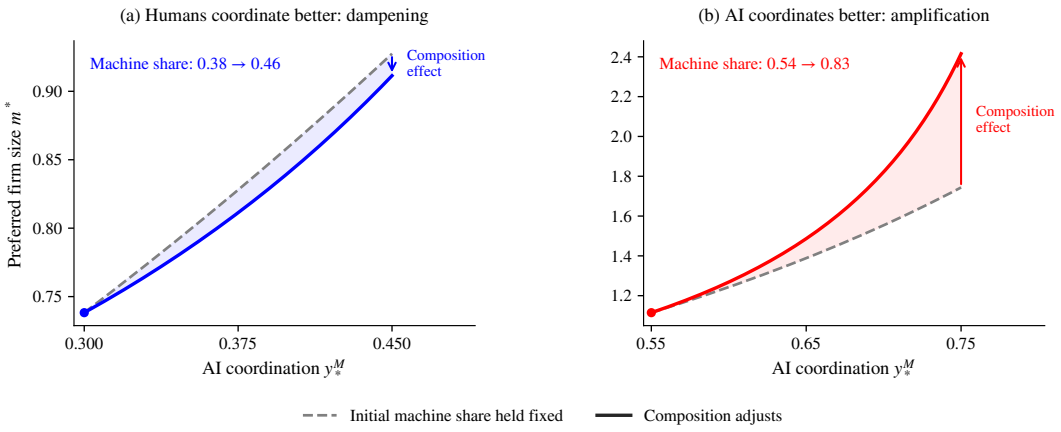


Figure 15: Composition can dampen or amplify firm growth.

The basic mechanism behind amplification is that better AI coordination raises preferred firm size holding fixed its current composition. But, as the firm grows, better coordination becomes especially valuable since it determines how output grows as productivity scales. For instance within our CES specification, we have $\partial \log F / \partial \bar{y}_* = \log(1+m)$. If we are in the theoretical regime in which machines coordinate better than humans (part (i) of Proposition 2) the firm tilts its composition of workers toward machines so the improvement in AI's ability to coordinate then drives up the firm's mean coordination $\bar{y}_*$ which then supports further growth, and so on. Thus, size and composition reinforce each other.

By contrast, if humans still coordinate better, growth in firm size tilts composition toward humans. However, better AI coordination also makes machines more attractive at a fixed size. If this direct incentive dominates, the machine share rises despite the size-driven shift toward humans. Machines then replace better human coordinators, so the change in composition pushes down the average coordination within the firm $\bar{y}_*$ which dampens firm growth.

Although Proposition 2 is stated within the CES environment with a single human type, the broader point extends: let θ index a differentiable path of AI coordination improvements and let s and s_i be the machine share and the share of human type i, with $s + \sum_i s_i = 1$. Mean coordination is $\bar{y}_* = s y_*^M + \sum_i s_i y_*^i$. Along the path of AI

improvement:[22]

$$\frac{\mathrm{d}m^*}{\mathrm{d}\theta} = \frac{\mathrm{d}m^*}{\mathrm{d}\bar{y}_*}\Bigg[\underbrace{s\frac{\mathrm{d}y_*^M}{\mathrm{d}\theta}}_{\substack{\text{coordination improves}\\ \text{at fixed shares}\\ \geq 0}} + \underbrace{\sum_{i=1}^{K}(y_*^i - y_*^M)\frac{\mathrm{d}s_i}{\mathrm{d}\theta}}_{\substack{\text{composition changes}\\ \gtrless 0}}\Bigg].$$

Since preferred size increases weakly with mean coordination, the direct effect is weakly positive. The sign of composition changes depends on whether the composition of the firm tilts towards machines or humans—as we have seen, the former amplifies the growth in firm size while the latter dampens it. This, in turn, might depend on wage effects that change the relative value of using machines vs humans in the model. For instance, wage changes could reinforce amplification—higher human wages encourage substitution toward machines, which raises mean coordination if machines coordinate better.

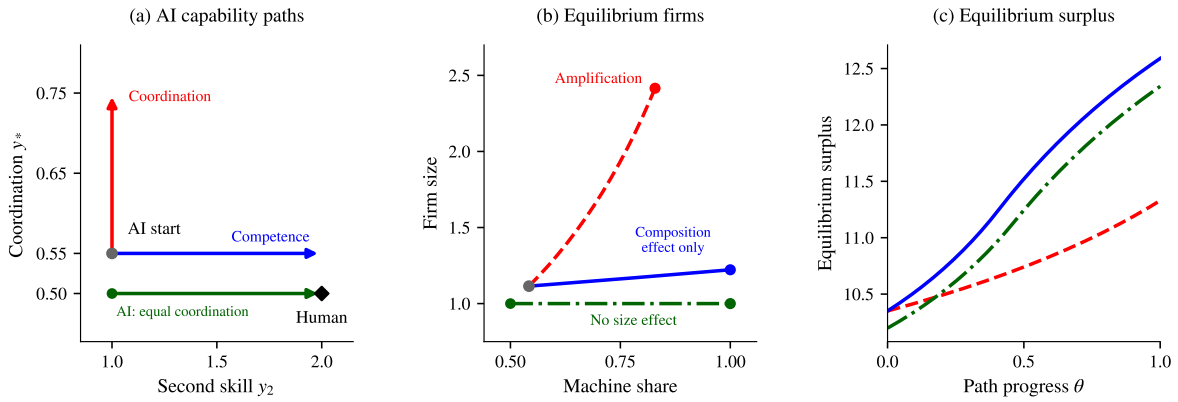


Figure 16: Coordination, competence, and surplus.

By contrast, competence gains that raise productivity by the same proportion at every firm size (which is implied by Assumption 2) leave a firm's preferred size unchanged at fixed composition. That is, competence gains make firms more productive without driving larger firms, i.e., the direct effect is zero. Changes in preferred firm size are driven solely by changes in firms' workers mean coordination ability. Figure 16 illusrates these forces by compares gains in AI coordination (red) and competence (blue) from the same mixed workforce.[23]

[22] Assuming that preferred size is unique, finite, and differentiable in mean coordination, and that the firm composition is also differentiable.

[23] Humans have profile $(1, 2, .5)$ and supply 20. For $\theta \in [0, 1]$, AI follows $(2, 1, .55 + .2\theta)$ in red,

Sharp threshold in concentration. Our previous theoretical analysis studied improvements in AI coordination in the regime where $\gamma^M < 1$. We now study what happens when it crosses 1.

Proposition 3 (Supercooperation concentrates large markets)**.** As $N \to \infty$:

(i) **Replication.** If $\gamma^M < 1$, every equilibrium has vanishing firms:

$$\operatorname*{ess\,sup}_{f \in \mathcal{F}^+} \frac{q_f}{Q(\sigma_N^*)} \longrightarrow 0.$$

The essential supremum is taken with respect to σ_N^* and is zero when no positive-scale firm is active.

(ii) **Concentration.** If the AI is supercooperative, every equilibrium becomes concentrated:

$$\int_{\mathcal{F}^+} \left(\frac{q_f}{Q(\sigma_N^*)} \right)^2 \mathrm{d}\sigma_N^*(f) \longrightarrow 1.$$

The proposition explains when large firms emerge as the market grows. If $\gamma^M < 1$, the average product of every sufficiently large feasible firm vanishes. The industry therefore expands by repeating bounded firm recipes, so the capacity of any one complete firm becomes small relative to total output. In the $N \to \infty$ limit this converges to the atomless economy of, e.g., Novshek and Sonnenschein (1987).

But when the AI tips into supercooperation, a single large AI-only firm can produce Q at a cost of order Q^{1/γ^M}, where $1/\gamma^M < 1$. Splitting the same output across several large firms loses this scale advantage. The equilibrium industry therefore assigns almost all output to firm types whose capacity is as large as the market. Our notion of concentration in part (ii) is a 'relaxed' Herfindahl–Hirschman index when there is a non-integer number of firms. A limit of one means that almost all output is assigned to firm types whose complete capacity is close to total market output.

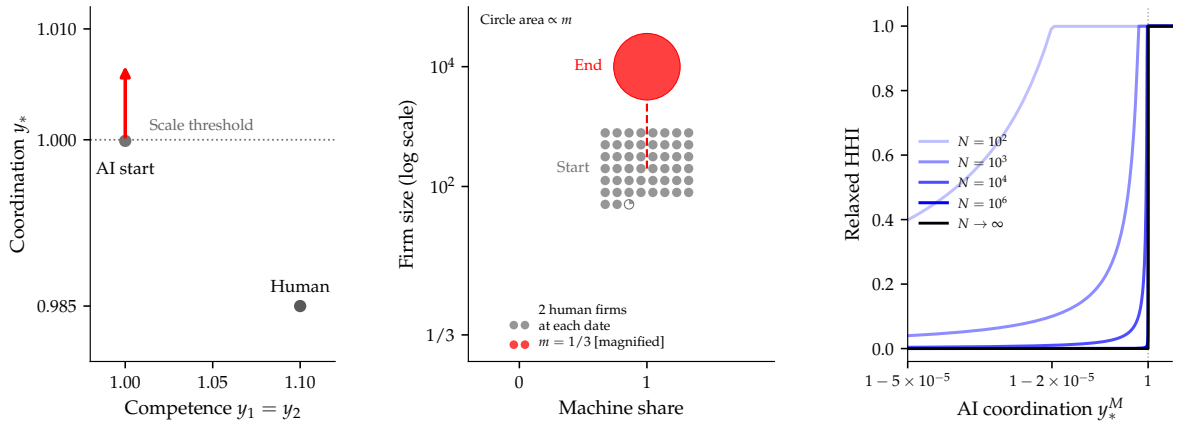
Figure 17 illustrates the shift from replicating small firms to concentrating output as AI coordination improves.[24] Two small human firms remain at both displayed dates, so concentration need not imply human exit. The finite-market curves restrict production

$(2, 1+\theta, .55)$ in blue, and $(2, 1+\theta, .5)$ in dark green. Set $\phi = 1$, $\rho = -1$, $r = \underline{w} = 1$, and $P(Q) = (2 - Q/10)_+$. Human supply stays slack and preferred firms fit in the market. Panel (c) compares these specified paths, so the comparison depends on their lengths; see Section D.3 for verification.

[24] Set $\phi = .02$, $\rho = 0$, $r = \underline{w} = 1$, human profile $(1.1, 1.1, .985)$ with supply $2/3$, AI profile $(1, 1, \theta)$, and $P_N(Q) = (2 - Q/(10N))_+$. Here $\gamma^M = \theta$. Panel (b) uses $N = 10^3$, with grey at $\theta = 0.9999$ and red at $\theta = 1.005$; panel (c) varies N. Proposition 3 excludes $\gamma^M = 1$.

to pure human and pure AI firms; entry checks at selected states do not establish equilibrium along the full paths (Section D.7).

In our model, "supercooperativeness" denotes a stylized mathematical property of production technology. In practice, there are real world factors that would likely prevent a firm singularity, such as regulation, competition among AI providers, capacity constraints, and rigid supply contracts.



(a) AI coordination path

(b) Firms at the start and end

(c) Concentration near one

Figure 17: AI coordination turns replication into concentration. In panel (b), circle area is proportional to m_f, except for the magnified human firms; grey and red show the start and end. The two human firms at each date are offset vertically for visibility.

5.3 Firm composition: human vs AI split. In the model, whether humans and AI work together depends on productivity and input costs. We first characterize firm composition at equilibrium wages, then recover conditions that depend only on technology, resource costs, and supplies.

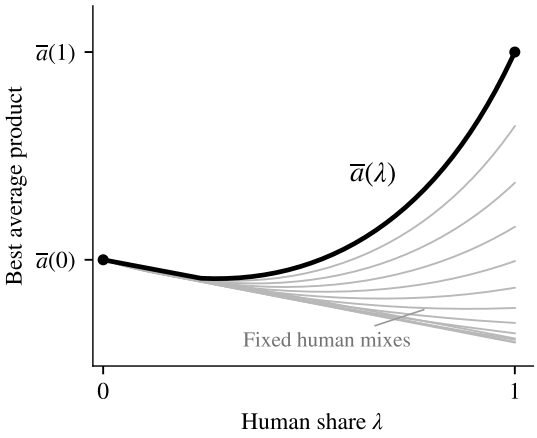
For a human share $\lambda \in [0,1]$ and wage vector $\boldsymbol{w}$, define *cost-adjusted best average product* by

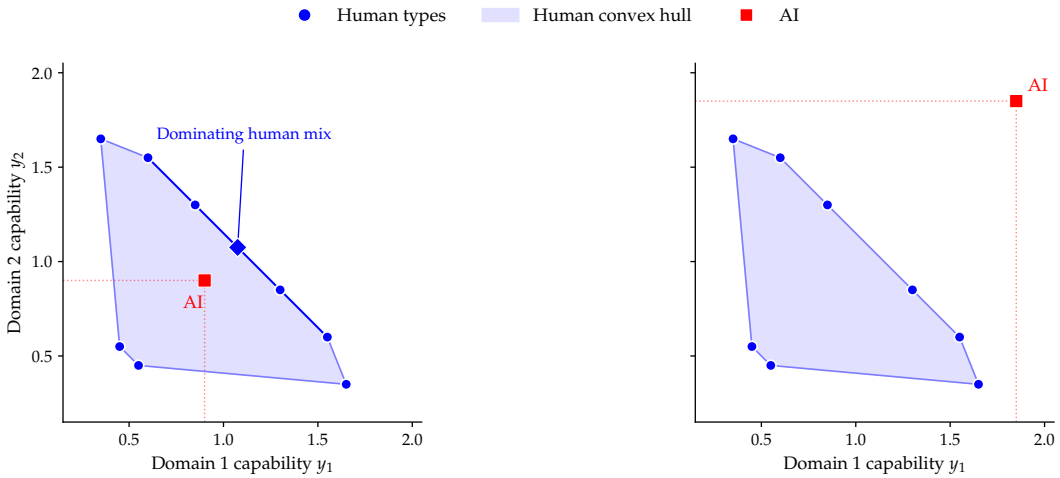
$$\bar{a}(\lambda, \boldsymbol{w}) := \sup_{\substack{\nu \in \Delta(\mathcal{Y}^H) \\ m \geq 0}} \frac{\overbrace{F\left(m[(1-\lambda)\delta_{\boldsymbol{y}^M} + \lambda\nu]\right)/m}^{\text{Average product}}}{\underbrace{(1-\lambda) + \lambda \int_{\mathcal{Y}^H} [w(\boldsymbol{y})/r] \, \mathrm{d}\nu(\boldsymbol{y})}_{\propto \text{ cost of workforce}}}.$$

We write $\bar{a}(\cdot) := \bar{a}(\cdot, r\mathbf{1})$ for the unadjusted average-product envelope, i.e., as if all

wages were equal to the rental rate.

The envelope can be convex, concave, or neither. Its shape depends on capabilities, technology, and wages. The figure on the right sets every human wage equal to r, so cost-adjusted best average product equals $\bar{a}(\lambda)$. Each grey curve fixes the human composition $\nu \in \Delta(\mathcal{Y}^H)$. The black upper envelope selects the best composition at each human share, capturing the gains from *mosaic intelligence*.





(a) A human mix dominates AI

(b) AI dominates all human types

Figure 18: Human matching vs AI dominance.

Proposition 4 (Human-AI production regimes)**.** Fix an equilibrium $(\sigma, \boldsymbol{w})$. Suppose the global maximum of $\bar{a}(\lambda, \boldsymbol{w})$ over human shares is finite and attained by a feasible firm at output $Q(\sigma)$.[25] Then:

(i) **Machine-only production.** If $r \leq \underline{w}$, Assumption 2 holds, L strictly increases in each productive capability, and $\boldsymbol{y}^M > \boldsymbol{y}$ for every $\boldsymbol{y} \in \text{co}(\mathcal{Y}^H)$, then all firms are machine-only.

(ii) **Humans in the loop.** If $r \geq \underline{w}$ and $\max_\lambda \bar{a}(\lambda, r\mathbf{1}) > \bar{a}(0, r\mathbf{1})$, then some

(positive measure set of) firms employ humans. If, in addition, there is a single human type, then:

(a) **Separate production.** If $\bar{a}(\cdot, r\mathbf{1})$ is strictly convex, then almost every active firm is pure human or pure machine, and some active firms are human-only.

(b) **Mixed production.** If $r = \underline{w}$, $\bar{a}(\cdot, r\mathbf{1})$ is strictly concave, and its finite left derivative at one is negative, then almost every active firm is mixed. These firms share the unique human share maximizing $\bar{a}(\cdot, w)$ at the equilibrium human wage w.

Proposition 4 gives sufficient conditions for machine-only production or for humans to remain in the loop. Part (i) applies when AI dominates every human combination and machines cost no more than humans. Figure 18(b) illustrates the productive-capability comparison: AI lies above the entire human hull. If this dominance also holds in coordination, replacing humans with machines raises output without raising cost.

Part (ii) allows humans to preserve employment through *mosaic intelligence*: complementary human skills can jointly outperform AI even when no individual human does, as Figure 18(a) illustrates. The upper envelope records the best human composition at each human share, so it captures these gains from combining skills. If the envelope rises above the machine-only endpoint, then as long as reservation wages are below the cost of compute ($\underline{w} \leq r$), some firms must employ humans. Parts (a) and (b) of Proposition 4 (ii) further relate the curvature of $\bar{a}(\cdot, \boldsymbol{w})$ to whether humans and AI work in separate or mixed firms as illustrated by Figure 19 which contrasts separate and mixed firms under the same demand and human supply.[26] In row (a), productive skills coincide but humans coordinate better: scale gains make the envelope convex, favoring separate firms. In row (b), equal coordination and complementary skills make the envelope concave, favoring firms with equal human and AI shares.

[25] For a maximizing share λ^*, human composition ν^*, and positive size m^*, feasibility means $m^*\lambda^*\nu^* \leq \mu$ and $F(m^*[(1-\lambda^*)\delta_{\boldsymbol{y}^M} + \lambda^*\nu^*]) \leq Q(\sigma)$.

[26] Both rows use $\rho = -1$, $\phi = 1$, $r = \underline{w} = 1$, human supply 50, and $P(Q) = (2 - Q/100)_+$. Row (a) has $\boldsymbol{y}^H = (1, 1, .8)$ and $\boldsymbol{y}^M = (1, 1, .4)$; row (b) has $\boldsymbol{y}^H = (2, 1, .6)$ and $\boldsymbol{y}^M = (1, 2, .6)$.

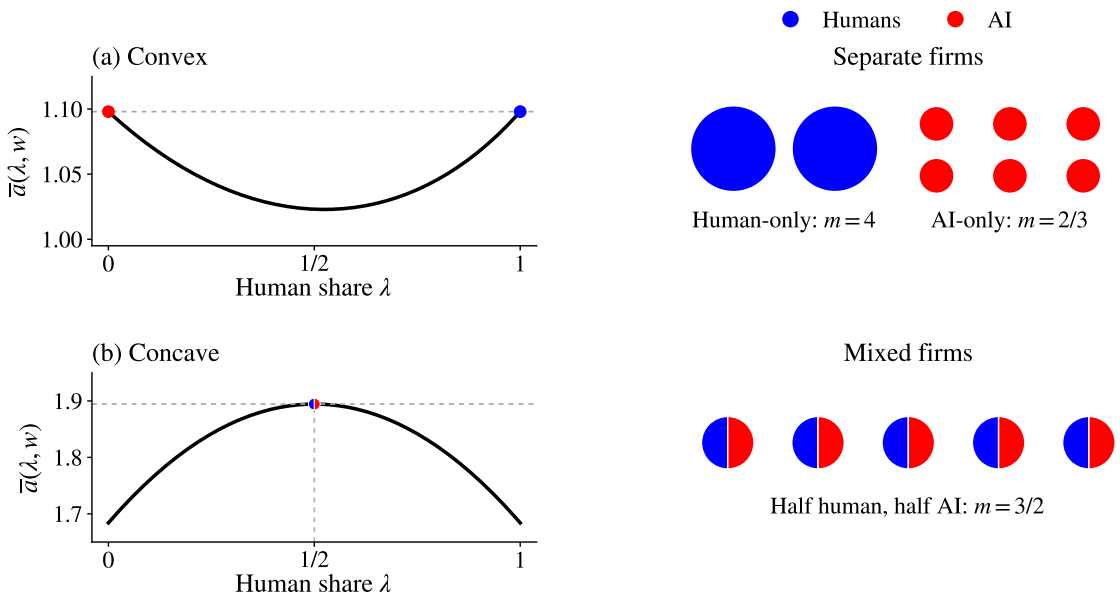


Figure 19: Curvature of $\bar{a}$ and firm composition in two CES economies.

6 HUMAN WAGES AND EMPLOYMENT

We now turn to the question of human wages and employment. We first develop a some results pinning down firms' willingness-to-pay for human workers, and how different types of human workers sort into different firms based on their specific strengths. Then, we study how different trajectories of AI improvements might drive changes in unemployment and wages in the model.

6.1 The marginal value of human talent. We first develop some machinery to study the marginal value of employing a human with capability profile $\boldsymbol{y}$. For a nonzero firm ν, the marginal product of capability $\boldsymbol{y}$ is

$$D_{\boldsymbol{y}} F(\nu) := \lim_{t \downarrow 0} \frac{F(\nu + t\delta_{\boldsymbol{y}}) - F(\nu)}{t}$$

which is simply the right directional derivative of output.[27] A firm f that values machines, $D_{\boldsymbol{y}^M} F(\nu_f) > 0$, bids

$$b_f(\boldsymbol{y}) := r \frac{D_{\boldsymbol{y}} F(\nu_f)}{D_{\boldsymbol{y}^M} F(\nu_f)}.$$

The intuition is the usual marginal reasoning: one more unit of the worker lets the firm remove $D_{\boldsymbol{y}} F / D_{\boldsymbol{y}^M} F$ machines without changing output. The bid is the resulting saving in machine costs at the rental rate $r > 0$. In fact, under Assumption 2 these are affine in capability:

$$b_f(\boldsymbol{y}) = r + \boldsymbol{\beta}^f \cdot (\boldsymbol{y} - \boldsymbol{y}^M), \tag{BID}$$

where the firm's prices for productive capability in domain d and coordination are

$$\beta^f_d = \frac{r S(\nu_f) \partial_d L}{m_f D_{\boldsymbol{y}^M} F(\nu_f)} \quad \text{and} \quad \beta^f_* = \frac{r L(\nu_f) \partial_2 S}{m_f D_{\boldsymbol{y}^M} F(\nu_f)}.$$

These prices tell us which employer values a worker's strengths most.

Each firm's bid $b_f(\boldsymbol{y})$ tells us how much the specific firm is willing to pay for a human worker of a given capability $\boldsymbol{y}$. Call a firm *eligible* for human type i if it has positive scale, uses machines, and $\nu^H_f(\{\boldsymbol{y}^i\}) < \mu_i$ i.e., it does not already employ the entirety of type i. The *bid envelope* for a given human type i is the upper envelope of what every firm is willing to pay for i:

$$\mathcal{E}_i(\sigma, \boldsymbol{y}^M) := \underline{w} \vee \operatorname*{ess\,sup}_{\substack{f \text{ eligible} \\ \text{for type } i}} b_f(\boldsymbol{y}^i),$$

where the essential supremum is taken with respect to σ. If no eligible firm is present, set the envelope equal to $\underline{w}$. Write $\mathcal{E} := (\mathcal{E}_i)_{i=1,\dots,K}$ for the vector of these envelopes.

Theorem 2 (Wages and sorting)**.**

(i) **Wage determination.** For any equilibrium outcome (σ, w),

$$\mathcal{E}(\sigma, \boldsymbol{y}^M) \leq w.$$

[27] This coincides with the Gâteaux derivative in direction $\delta_{\boldsymbol{y}}$ when the two-sided derivative is defined. Here we use a right derivative because subtracting mass at $\boldsymbol{y}$ can make the measure ν negative (which is undefined).

Equality holds for type i if some of that type is unemployed, in which case $w_i = \underline{w}$, or an active interior firm f employs that type, in which case $w_i = b_f(\boldsymbol{y}^i)$.

(ii) **Sorting.** Fix an equilibrium (σ, w). If Assumption 2 holds and an active interior firm f employs type i, then σ-almost every firm g that is eligible for i satisfies

$$b_f(\boldsymbol{y}^i) - b_g(\boldsymbol{y}^i) = (\boldsymbol{\beta}^f - \boldsymbol{\beta}^g) \cdot (\boldsymbol{y}^i - \boldsymbol{y}^M) \geq 0. \tag{FAN}$$

With finitely many active machine-using firms, the capability profiles at which each firm bids highest form polyhedral cones with common apex $\boldsymbol{y}^M$. Workers with interior matches at these firms sort according to these regions.

Figure 20 illustrates Theorem 2.

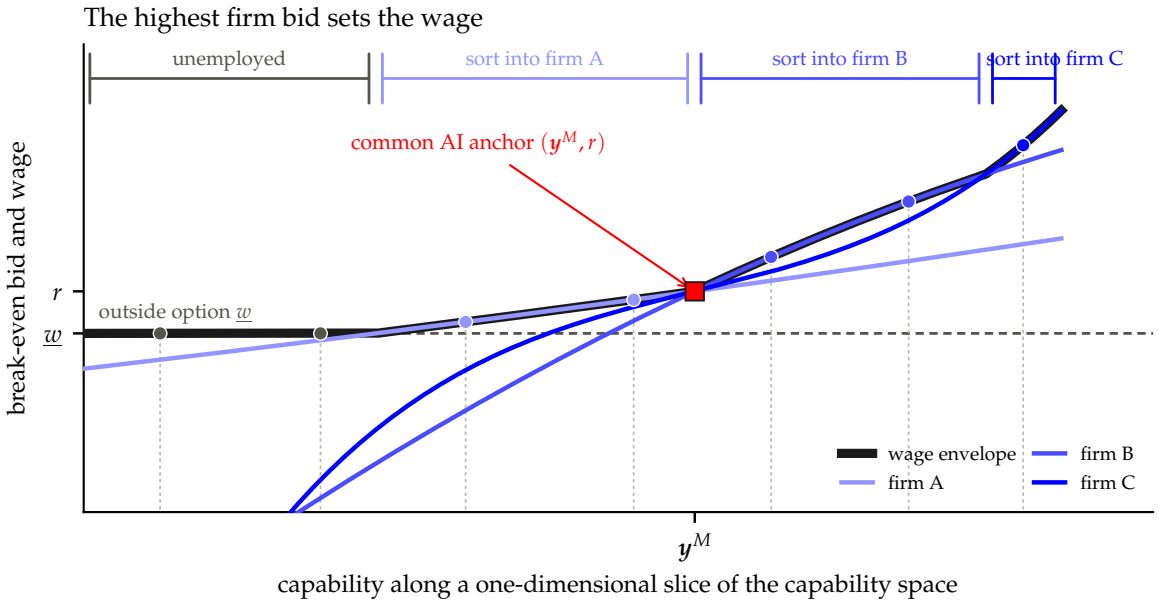


Figure 20: Firms' bids for humans.

The black curve gives the highest firm bid for each worker type (depicted on the x axis), with the outside option as a floor. The blue curves show how firms differ in what they will pay for the same worker, and the brackets show where each firm bids highest. At an interior match, a firm can substitute workers for machines in either direction, which pins the wage to its bid. All bids pass through the common AI point $(\boldsymbol{y}^M, r)$: a worker with the AI's capabilities saves exactly the cost of one machine.

Part (ii) of Theorem 2 links sorting to workers' strengths relative to AI. In Figure 21(a), worker A has stronger competence but weaker coordination than AI; worker B has the opposite advantages. Worker C lies in the medium firm's region, while D lies in the outside-option region. Panel (b) shows why: the small, large, and medium firms bid most for A, B, and C, respectively, while every bid for D falls below the outside option. At an interior match, the highest eligible bid also pins down the worker's wage. Workers A and B can earn more than the machine rental rate despite being weaker than AI in one capability.

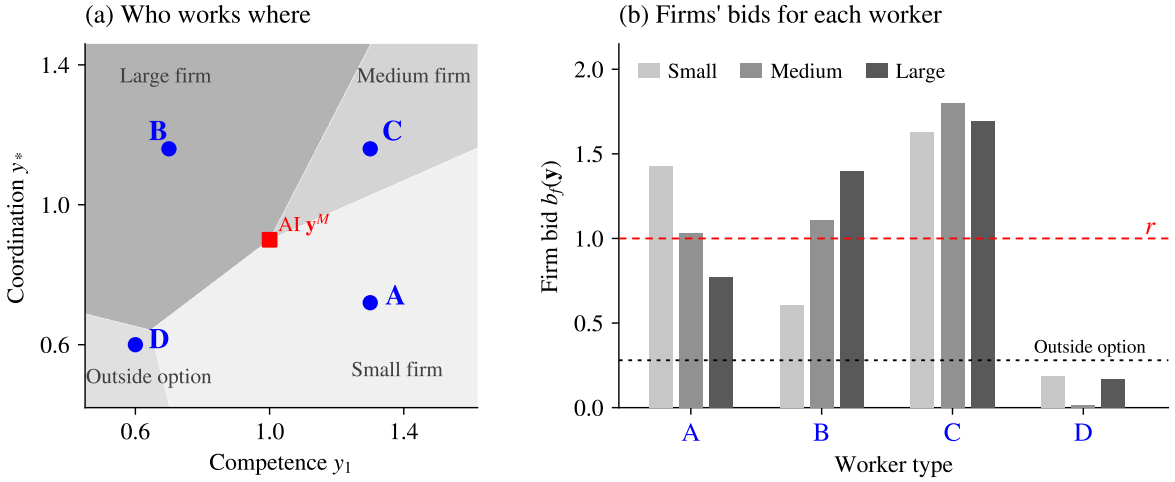
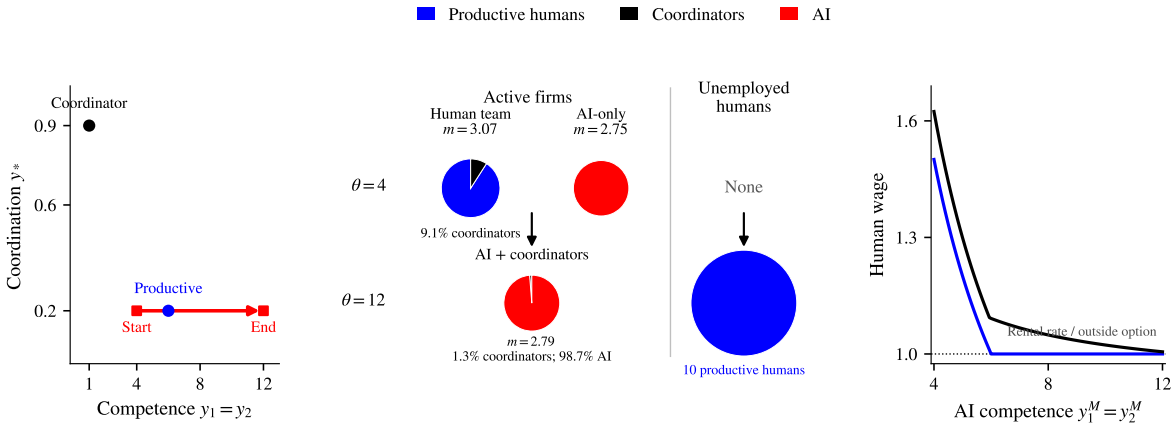


Figure 21: Comparative advantage sorts humans.

Humans can earn more than a 'supercompent' AI—that far exceeds humans in every domain $d = 1, 2, \ldots n$—if their coordination advantage is valuable enough. By (BID), an eligible firm bids above the machine rental price r exactly when its weighted sum of human capability advantages is positive (cf. Agrawal et al., 2024, 2025). Figure 22 shows human coordinators moving from human teams into AI teams as AI competence rises, because their coordination advantage raises the productivity of the AI workforce.[28] Productive humans leave employment, while coordinators remain employed above their outside wage; both human wages fall along this plotted path.

[28]Productive humans have profile $(6, 6, .2)$ and supply 10; coordinators have $(1, 1, .9)$ and supply 1. AI follows $(\theta, \theta, .2)$ for $\theta \in [4, 12]$. Set $\rho = -1$, $\phi = 3$, $r = \underline{w} = 1$, and $P(Q) = 2(1 - Q/3000)_+$.



(a) AI competence path (b) Who works together (c) Equilibrium wages

Figure 22: Human coordinators work with increasingly productive AI

6.2 The dynamics of unemployment. Section 5.3 uses the cost-adjusted best average-product envelope $\bar{a}$ to determine which firm compositions operate. We now use the same envelope to study wages and employment as AI capabilities improve within our model. Suppose there is a single human type with capability $\boldsymbol{y}^H$ and mass $\mu > 0$. A firm's human share $\lambda \in [0,1]$ determines its unit composition $\nu_\lambda := \lambda\delta_{\boldsymbol{y}^H} + (1-\lambda)\delta_{\boldsymbol{y}^M}$. The cost-adjusted envelope then satisfies $\bar{a}(\lambda, w) = \frac{\bar{a}(\lambda)}{1+\lambda(w/r-1)}$. Under the conditions of Proposition 4, active firms maximize this envelope. The first-order condition implies

$$\frac{w}{r} = \frac{\bar{a}(\lambda) + (1-\lambda)\bar{a}'(\lambda)}{\bar{a}(\lambda) - \lambda\bar{a}'(\lambda)}.$$

Relative wages therefore equal relative marginal products after firms choose their preferred size. In the special case in which $\bar{a}$ is convex and human-only and AI-only firms are both active at their preferred scales, zero profit instead equates their cost-adjusted best average products: $\frac{\bar{a}(1)}{w} = \frac{\bar{a}(0)}{r} \implies \frac{w}{r} = \frac{\bar{a}(1)}{\bar{a}(0)}$.

Indeed in the CES environment, we can compute this envelope explicitly. For each domain $d \in \{1, 2, *\}$, mean capability at human share λ is $\overline{y}_d(\lambda) := \lambda y_d^H + (1-\lambda)y_d^M$. If $1 - \phi < \overline{y}_*(\lambda) < 1$, preferred size is $m^*(\lambda) = \phi/[1-\overline{y}_*(\lambda)] - 1$, and substitution gives

$$\bar{a}(\lambda) = \underbrace{\left(\frac{\overline{y}_1(\lambda)^\rho + \overline{y}_2(\lambda)^\rho}{2}\right)^{1/\rho}}_{\text{productive competence}} \underbrace{e^{\phi-[1-\overline{y}_*(\lambda)]}\left(\frac{1-\overline{y}_*(\lambda)}{\phi}\right)^{1-\overline{y}_*(\lambda)}}_{\text{gains from choosing firm size}},$$

noting that as usual, at $\rho = 0$, the first factor is $\sqrt{\overline{y}_1(\lambda)\overline{y}_2(\lambda)}$ and if $\overline{y}_*(\lambda) \leq 1 - \phi$, preferred size is zero and the scale factor is one.[29]

The curvature of $\overline{a}$ with respect to λ reflects two forces. First, mixing humans and AI can balance their productive strengths, but with diminishing gains as the composition becomes more balanced. Composition also changes mean coordination and hence preferred firm size. When the more productive type also coordinates better, increasing its share raises both competence and the gains from scale; this reinforcement can make $\overline{a}$ convex. Our key observation is that the curvature of $\overline{a}$ determines not only the composition of firms at fixed levels of human and AI capabilities (as we have seen from Proposition 4), but also how wages and unemployment adjust as AI improves.

Proposition 5 (Gradual vs sudden displacement)**.**

(i) **Gradual replacement.** Suppose $\overline{a}$ is strictly concave along a continuous AI capability path $\boldsymbol{y}^M(\theta)$ and assume all firm compositions have feasible preferred scales.[30] Suppose these scales are unique and share a finite upper bound. Then equilibrium output, human employment, firms' human share, and firm size change continuously. The economy exhibits full unemployment if and only if

$$\left[1 + \frac{\overline{a}'_+(0)}{\overline{a}(0)}\right] \leq \frac{w}{r}.$$

(ii) **Sudden replacement.** Suppose $\overline{a}$ is strictly convex along a continuous AI capability path $\boldsymbol{y}^M(\theta)$, and the pure human and pure machine mixes have feasible preferred scales at equilibrium output. If machine-only firms are active in every equilibrium, the economy exhibits full employment when

$$\frac{\overline{a}(0)}{r} < \frac{\overline{a}(1)}{\underline{w}}$$

and full unemployment when the inequality is reversed. Wages are $\max\{\underline{w},\, r\overline{a}(1)/\overline{a}(0)\}$.

In part (i), displacement is *within-firm*: the left side of the condition in Proposition 5(i) is the value of adding a human to an all-machine industry at the price supported by machine production. When the LHS is smaller than the RHS, then a machine-only

[29] For positive productive competence, the scale factor has supremum e^{ϕ} at $\overline{y}_*(\lambda) = 1$, attained only in the limit as size grows. Best average product is unbounded when $\overline{y}_*(\lambda) > 1$.

[30] In part (i), preferred scales are feasible at the output that maximizes surplus when we omit individual firms' staffing and market constraints.

firm does not benefit from employing any humans even at the reservation wage. In part (ii), displacement is *across-firms*: in an economy that is already segregated—firms are machine-only or human-only—changes in underlying AI capabilities can discontinuously change the set of equilibrium industry structures, thereby wiping out human-only firms. Figure 23 displays these two cases.[31] In the first case, firms reduce their human share continuously; in the second, the industry replaces human firms once AI production becomes cheaper.

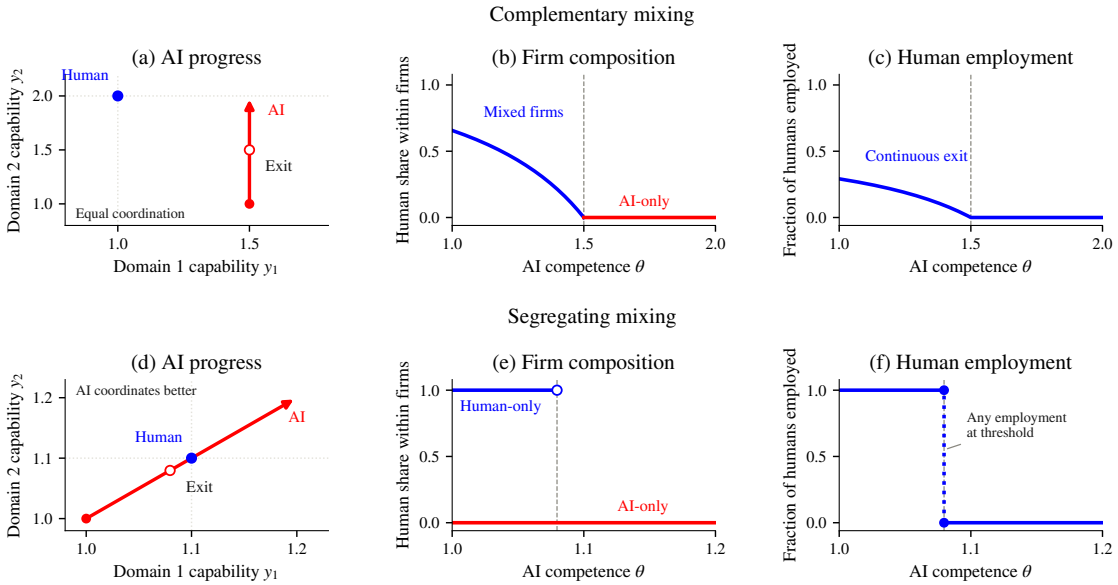


Figure 23: AI progress, firm composition, and human exit.

[31] Both rows use $\rho = -1$ and $r = \underline{w} = 1$. Above, humans have profile $(1, 2, .9)$ and supply 20, AI follows $(1.5, \theta, .9)$ for $\theta \in [1, 2]$, $\phi = .3$, and $P(Q) = (2 - Q/10)_+$. Below, humans have $(1.1, 1.1, .985)$ and supply $2/3$, AI follows $(\theta, \theta, .9999)$ for $\theta \in [1, 1.2]$, $\phi = .02$, and $P(Q) = (2 - Q/1000)_+$.

More broadly, the curvature of $\bar{a}$ can transition from concave to convex as AI capabilities improve. In such settings, we might first see a gradual increase in unemployment (part (i) of Proposition 5) before a sudden spike (part (ii) of Proposition 5). Figure 24 illustrates the change in curvature as AI coordination improves.[32]

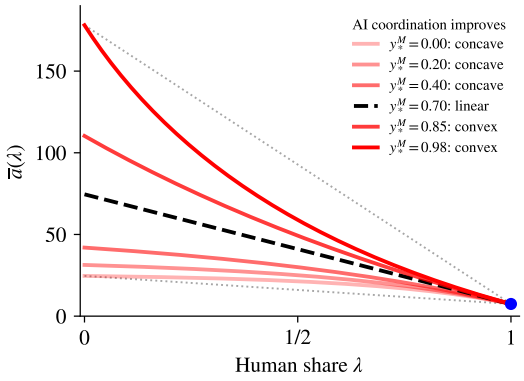


Figure 24: $\boldsymbol{y}^M$ and curvature of $\bar{a}$.

6.3 The reorganization effect. AI progress changes wages both directly and through changes in firms. We call the second channel the *reorganization effect*. A firm can change its size, its machine share, and the mix of human skills it employs. The following decomposition separates these channels along a smooth path for a worker and her employer.

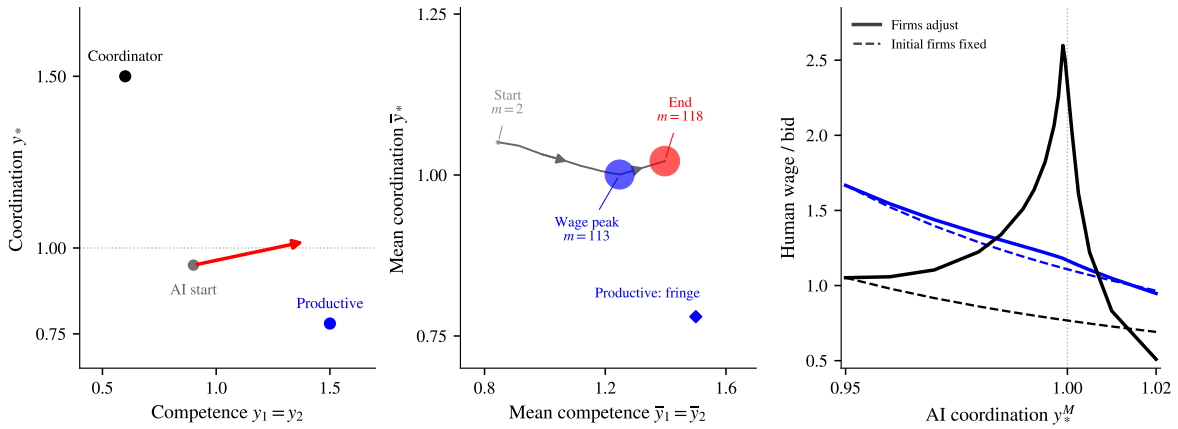
Proposition 6 (Wage effects combine technology and changes in firms)**.** Suppose Assumption 2 holds and let AI capabilities $\boldsymbol{y}^M(\theta)$ be parametrized by θ, and let $\overline{\boldsymbol{y}}^H$ denote the mean capability profile of the firm's human workers. Consider a continuously differentiable firm selection whose firms are interior and employ human type i. Assume the firm's bid is continuously differentiable in its size, machine share, mean human capability profile, and θ. Then

$$\frac{\mathrm{d}w_i}{\mathrm{d}\theta} = \underbrace{\partial_\theta b_f(\boldsymbol{y}^i)}_{\text{technology}} + \underbrace{\overbrace{\partial_m b_f(\boldsymbol{y}^i)\frac{\mathrm{d}m}{\mathrm{d}\theta}}^{\text{firm size}} + \overbrace{\partial_s b_f(\boldsymbol{y}^i)\frac{\mathrm{d}s}{\mathrm{d}\theta}}^{\text{machine share}} + \overbrace{\nabla_{\overline{\boldsymbol{y}}^H} b_f(\boldsymbol{y}^i)\cdot\frac{\mathrm{d}\overline{\boldsymbol{y}}^H}{\mathrm{d}\theta}}^{\text{mix of human skills}}}_{\text{reorganization}}.$$

The technology term changes AI capabilities while holding human and machine inputs fixed. The remaining terms let the employer adjust those inputs. Equilibrium maximizes total surplus, not each worker's wage, so the envelope theorem for surplus does not remove the reorganization terms from wage changes. The reorganization effect can therefore change wages at first order and alter the division of surplus.

[32] Humans have profile $(1, 1, .7)$ and AI has $(10, 10, y_*^M)$, with $\phi = 3$ and $\rho = -1$. Only AI coordination changes. Dotted lines join the endpoints of the first and last curves.

This is illustrated quite starkly in Figure 25 in which (endogenoously) larger firms initially raise the value of human coordination, before further AI progress lowers it (solid black line in panel (c)).[33] By contrast, holding the composition of firms constant generates monotonically decreasing wages for the human coordinator (dotted black line in panel (c)).



(a) AI capability path

(b) Human employers

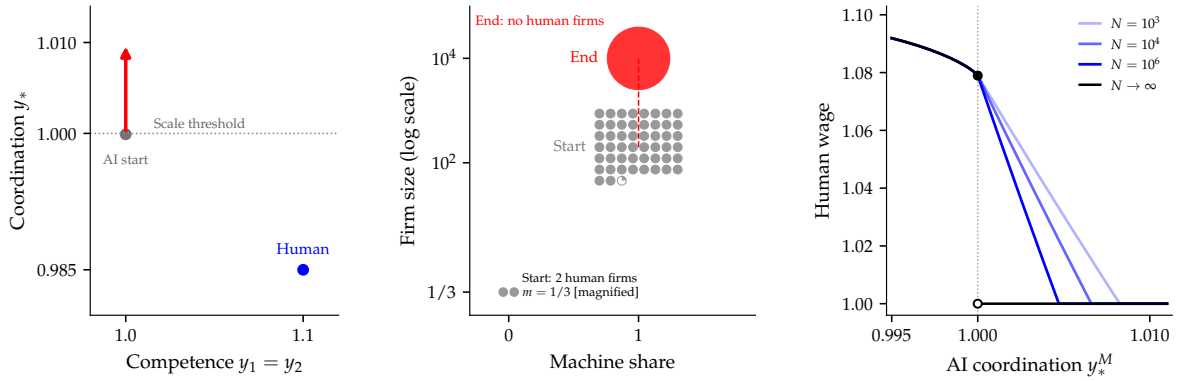
(c) Wages and fixed-firm bids

Figure 25: Wage changes with/without reorganization.

Reorganization can also have drastic impact on wages by (abruptly) changing which firms survive in equilibrium. In Figure 26(a), AI improves at coordination while humans retain their task-skill advantage. Below the scale threshold, small human firms survive alongside AI firms. Once AI coordination exceeds one, AI output per worker can grow without bound with firm size. Adding humans improves task skills but reduces coordination, which becomes more costly in larger firms. In this example, sufficiently large AI firms therefore undercut human production and human employers disappear (Figure 26(b)). Human wages fall to the outside option, and larger markets make the decline sharper (Figure 26(c)). The jump on the black curve describes the large-market limit; it does not establish a jump in a finite market.[34]

[33] Human profiles are $(1.5, 1.5, .78)$ and $(.6, .6, 1.5)$, each with supply .4; AI moves from $(.9, .9, .95)$ to $(1.4, 1.4, 1.02)$. Set $\rho = -1$, $\phi = .02$, $r = 1$, $\underline{w} = .28$, and $P(Q) = 1.64(1 - Q/300)_+$.

[34] The calibration is in footnote 24. Panel (b) uses $N = 10^3$, with grey at $y_*^M = .9999$ and red at 1.010; panel (c) uses $N = 10^3, 10^4, 10^6$. Circle area measures firm size except for the magnified human firms of size $1/3$; full circles represent one unit of activity and sectors represent fractional units. Section A.3.4 verifies the limiting curve. The finite-market curves use pure-firm cost calculations with entry checks at selected states; see Section D.7.



(a) AI coordination path

(b) Firms at $N = 10^3$

(c) Wages as markets grow

Figure 26: Firm concentration and the loss of the human wage premium.

Wages depend on how surviving firms reorganize and on which employers remain active. Even a small employer can sustain a wage premium, so its disappearance can matter more for wages than its share of output suggests.

7 CONNECTION TO THEORY AND EMPIRICS

We have offered a framework for thinking about endogenous firm formation, worker sorting, surplus, and wages. We now explain how it connects to influential models in the existing literature on automation and wages. Then, we briefly discuss some early empirical evidence in line with our model's predictions, as well as outline avenues for future empirical work.

7.1 Nesting the task model. Optimal task assignment under constant returns recovers the task economy. Scale effects let AI teams cut costs without improving at individual tasks.

Equivalence to task model. Human type i has mass μ_i, reservation wage $\underline{w}$, and task productivities y_d^i. Machines have task productivities y_d^M and unlimited supply at rental cost r. Agents divide time across tasks (Acemoglu and Restrepo, 2018; Aghion et al., 2019), whose outputs enter $L(z) = (n^{-1}\sum_{d=1}^{n} z_d^\rho)^{1/\rho}$, with $\rho < 1$ and the geometric-mean limit at $\rho = 0$. Demand remains P_N.

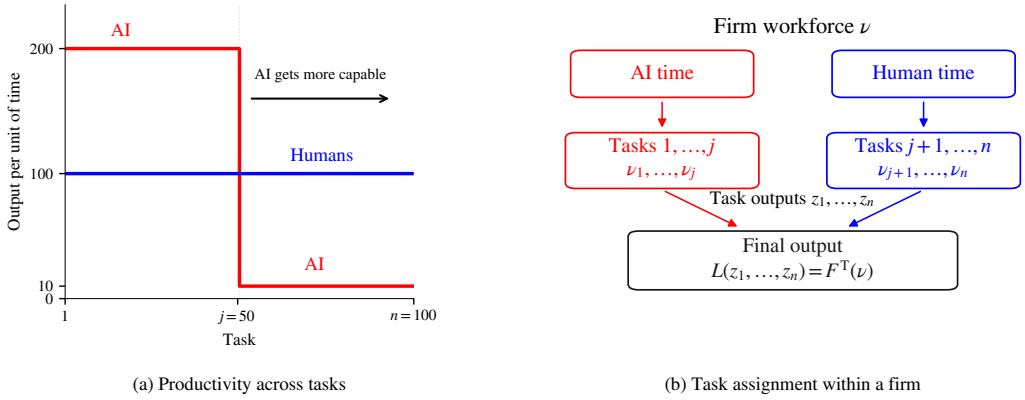
The *task technology* assigns the firm's human and machine inputs ν optimally across

tasks, with task allocations $\nu_1, \ldots, \nu_n$:

$$F^{\mathrm{T}}(\nu) = \max_{\substack{\nu_d \in \mathcal{M}_+(\mathcal{Y}) \\ \sum_d \nu_d = \nu}} L\left(\int_{\mathcal{Y}} y_1 \, \mathrm{d}\nu_1, \ldots, \int_{\mathcal{Y}} y_n \, \mathrm{d}\nu_n\right).$$

Our technology is F^{T}. With task productivities bounded away from zero, F^{T} satisfies Assumption 1 and has constant returns: $F^{\mathrm{T}}(m\nu) = mF^{\mathrm{T}}(\nu)$. Average product is therefore independent of scale, and the tail exponent is $\gamma \equiv 1$.

Theorem 3 (Outcome equivalence to task model)**.** Under technology F^{T}, every equilibrium industry σ produces the task economy's unique competitive output Q_N^*: $Q(\sigma) = Q_N^*$. The competitive task allocations are exactly the aggregates of equilibrium industries with optimal task assignments. If some equilibrium uses machines, the two economies also have the same set of equilibrium wage vectors.



(a) Productivity across tasks

(b) Task assignment within a firm

Figure 27: Task productivity and assignment within a firm.

In a constant returns world, the boundaries of the firm do not matter for the economy's efficiency since pooling firms lets workers and machines keep their old task assignments, so it cannot lower total output; conversely, we can split an optimal assignment into smaller copies without losing output. As usual, when machines are put to work, their fixed rental rate $r > 0$ anchors task prices. Workers then earn the value of their best task, subject to the outside option. Figure 27 illustrates how workers and machines divide their time across tasks before task outputs enter the final good.[35]

[35] Humans have productivity 100 on each of $n = 100$ tasks. AI has productivity 200 on tasks $1, \ldots, j$

What do firms add to task wages? We now allow coordination to affect scale: output is $F(\nu) = F^{\mathrm{T}}(\nu)S(m, \overline{y}_*)/m$, where m is firm size, $\overline{y}_*$ is mean coordination, and S is the CES scale function. At a positive preferred size, define the firm's *coordination price*

$$\kappa_f := \frac{F^{\mathrm{T}}(\nu_f)}{m_f} \cdot \log(1 + m_f)$$

which is the task output per agent times the scale value of mean coordination.

Proposition 7 (Firms value task skills and coordination)**.** If an active firm f values machines and has a positive preferred size $m_f \in m^*(\overline{y}_*)$, then its bid for capability profile $\boldsymbol{y}$ is

$$b_f(\boldsymbol{y}) = r\,\frac{D_{\boldsymbol{y}}F^{\mathrm{T}}(\nu_f) + \overbrace{\kappa_f(y_* - \overline{y}_*)}^{\substack{\text{human coordination}\\ \text{contribution}}}}{D_{\boldsymbol{y}^M}F^{\mathrm{T}}(\nu_f) + \underbrace{\kappa_f(y_*^M - \overline{y}_*)}_{\substack{\text{machine coordination}\\ \text{contribution}}}}. \tag{TASK-BID}$$

If the firm is interior and employs human type i, then $w_i = b_f(\boldsymbol{y}^i)$. With constant returns, $S/m \equiv 1$, so coordination does not affect output and the bid reduces to the task-model bid.

For a human capability profile $\boldsymbol{y}$, the numerator combines the worker's task contribution with the value added through coordination; the denominator does the same for a machine. A human who improves coordination raises the firm's bid, while a machine that improves coordination raises lowers the bid. As in the task model, it is the *relative* price of humans vs machines that matters so humans can remain employed even when AI is more productive on every task.

We can use Proposition 7 to pin down the threshold at which humans become displaced. Suppose there is one human type (that can perform all tasks at fixed productivity). Under standard regularity conditions,[36] humans are absent from every equilibrium

and 10 on the rest, with $j = 50$. For a common human wage w and machine rental cost r, the shown specialization minimizes costs when $r/2 < w < 10r$: AI performs the first group and humans the second. Agents divide their time across tasks; they do not supply their full productivity to every task at once.

[36] We use the conditions in Proposition 5(i): maximal output per agent is strictly concave in the human share, and all mixes have unique, uniformly bounded preferred scales that fit human supply and the market at the relaxed surplus optimum. The appendix verifies these conditions in Figure 28 while AI coordination is below human coordination, including at human exit.

exactly when

$$\underbrace{\max_d \frac{y_d^H}{y_d^M}}_{\substack{\text{best relative}\\ \text{task productivity}}} + \underbrace{(y_*^H - y_*^M)\log\left(1 + m^*(y_*^M)\right)}_{\substack{\text{human coordination contribution}\\ \text{at the machine-only size}}} \leq \underbrace{\frac{\underline{w}}{r}}_{\substack{\text{relative}\\ \text{outside option}}} .$$

The first term is the human's best relative task productivity; the second is the human coordination contribution at the preferred machine-only size. A coordination advantage can cover a shortfall in task productivity. With constant returns the second term vanishes.

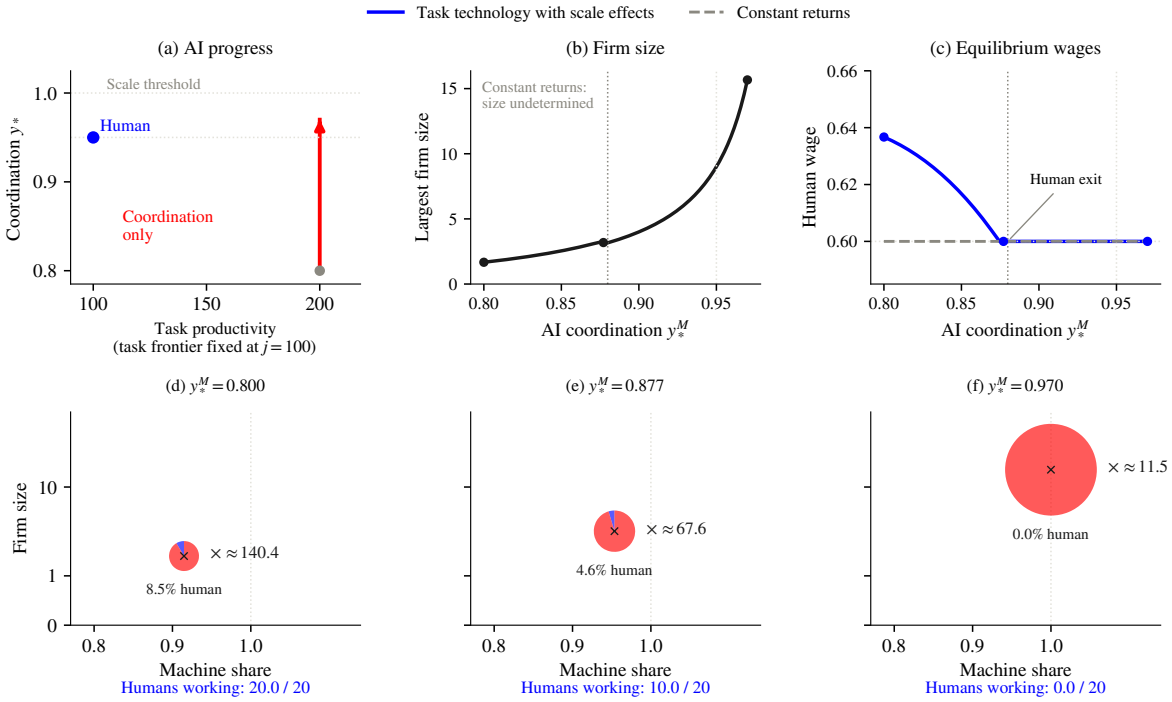


Figure 28: Humans as coordinators

Figure 28 holds AI task productivity at twice the human level on every task. Humans earn a premium through their coordination advantage; the premium falls as AI closes that gap.[37]

[37] We use $n = 100$ tasks, human productivity 100 on every task, human coordination $y_*^H = .95$, human supply $\mu = 20$, $r = 1$, $\underline{w} = .6$, $\phi = .5$, and $P(Q) = (3 - Q/200)_+$. AI productivity is 200 on every task and coordination rises from .8 to .97. Each pie represents one firm: area measures firm size, blue shows humans, and red shows AI. The multiplier gives the number of firms.

Non-monotonic firm sizes. Adding firm sizes to the task model also generates rich dynamics on how (endogenous) firm size interacts with the direction of AI progress. Figures 29, 30 and 31 compare AI improvements in different directions. Figure 29 illustrates improvements in only competence à la Acemoglu and Restrepo (2018); Korinek and Suh (2024). Because of constant returns, firm sizes are indeterminate (here we show them as atomless) and equilibrium wages have the familiar hump shape where first the productivity effect dominates before the displacement effect dominates.

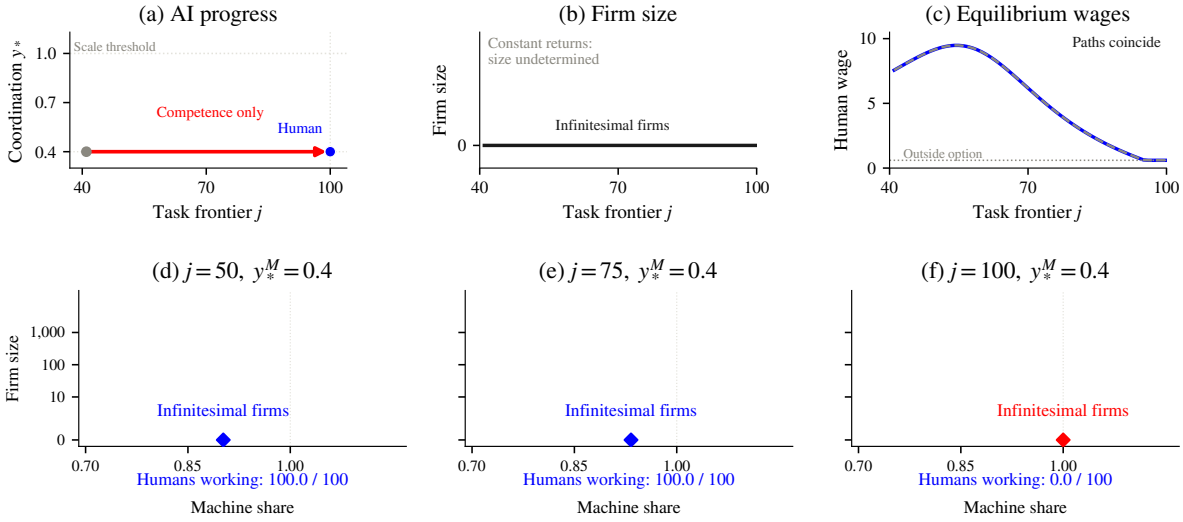


Figure 29: Improvement in AI competence (automating tasks).

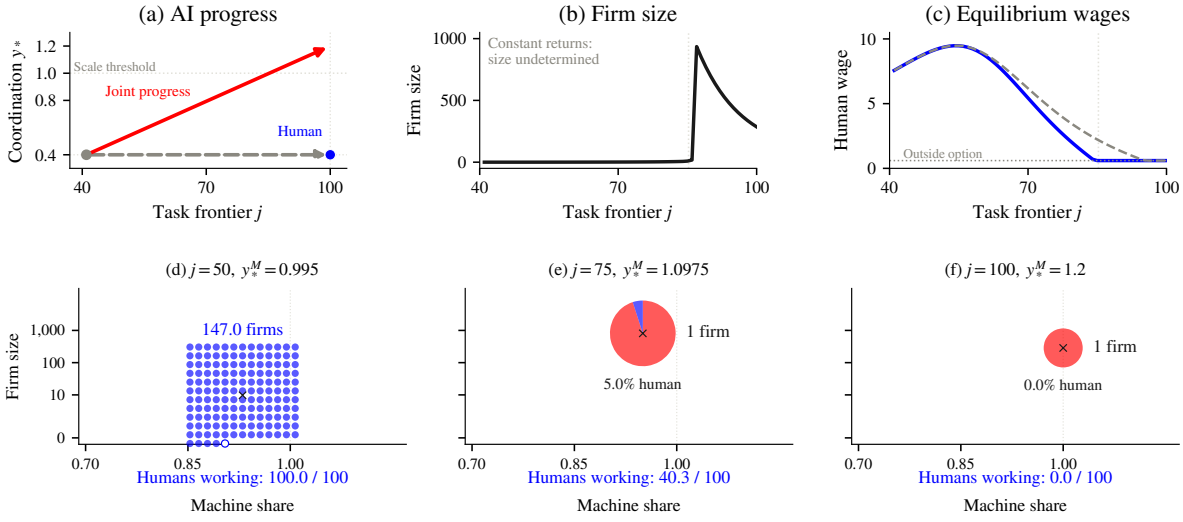


Figure 30: Improvement in AI competence and coordination.

With scale effects, however, AI improvements in coordination drive firm sizes to spike when its ability to coordinate γ^M crosses 1 (cf Proposition 3). This is illustrated by Figures 30 and 31. But as AI coordination continues to improve, the firm sizes start to fall because AIs that coordinate better allow firms produce more with fewer agents, while the downward-sloping demand curve limits the expansion of sales. Hence, if output per agent grows faster than sales, the workforce falls even as output rises which generates this non-monotonicity.

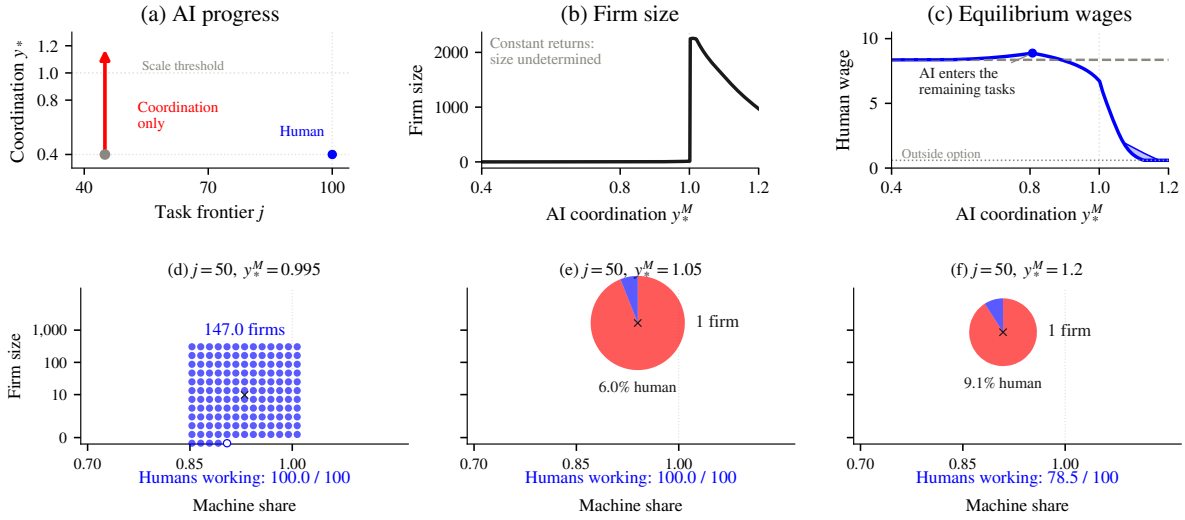


Figure 31: Improvement in AI coordination.

7.2 Early empirical evidence. Evidence to date suggests that current, publicly available AI systems have jagged domain competence—often exceeding human experts—alongside subcritical coordination capabilities. Evidence for the former primarily comes from benchmarks pitting humans versus AI models on identical tasks. Across 44 occupations on OpenAI's GDPval benchmark, frontier models tie or beat seasoned professionals on 48–71% of bounded deliverables (Patwardhan et al., 2025; OpenAI, 2025), while commercial legal tools outperform practicing lawyers on overall research accuracy (80% vs. 71%) yet trail attorneys by 9 percentage points on tasks requiring contextual judgment (Vals AI, 2025). Similarly, on Mercor's expert evaluations across banking, consulting, law, and medicine, models reach 67% on one-shot "closed tasks" (Vidgen et al., 2025) but drop to 24% on multi-step workflows (Vidgen and Mercor Research, 2026), and in research engineering models score 4× higher than human experts over 2-hour sprints but half as much over 32–48 hour projects (Wijk et al., 2024; Starace et al., 2025). On coordination, when multiple AI agents work together, output scales

sublinearly with team size ($\gamma^M \in [0.48, 0.68] < 1$; Ord 2026),[38] and uncoordinated swarms amplify errors up to 17.2× and fail on 41–87% of workflows (Cemri et al., 2025; Kim et al., 2025; Yang et al., 2026). AI-run stores degrade in performance over long horizons unless paired with substantial scaffolding (Backlund and Petersson, 2025; Anthropic, 2025a,b).

This evidence can be interpreted through our framework as predicting that we will see the rise of smaller firms, and existing firms will adjust primarily through workforce composition rather than scale. Data is consistent with these patterns: AI-native startups employ 25% fewer people overall and roughly 15% lower shares of entry-level workers and managers, and run flatter hierarchies than same-cohort peers (Kim and Koning, 2026; Ewens and Giroud, 2025). Hoffmann et al. (2025) find that AI-enabled developers work in smaller groups, individuals assisted by AI match the output quality of two-person human teams (Dell'Acqua et al., 2025), and others show that AI-adopting firms seem to reallocate labor away from routine roles without shrinking overall employment (Hampole et al., 2025; Baslandze et al., 2026).

The labor market response to AI has thus far been muted. Chatbot adoption has produced minimal aggregate changes in earnings or hours (Yotzov et al., 2026), with only 3–7% of time savings passing into pay (Humlum and Vestergaard, 2025), though notably this evidence is from before the adoption of agentic AI. But there is substantial heterogeneity in these aggregate estimates. Junior employment in exposed occupations has fallen 7.7–19% through reduced hiring and starting wages have dropped 4.2%, while senior employment and pay remain stable if not rising (Brynjolfsson et al., 2026; Hosseini and Lichtinger, 2025; Azar et al., 2026). Under the affine bid schedule (Theorem 2), entry-level workers map closest to the capabilities that machines can match, which leads their wages to be compressed first. Senior workers retain coordination advantages and see both demand and wages increase. Notably, even when AI raises early-career productivity most, increased machine competence narrows $\boldsymbol{y}^i - \boldsymbol{y}^M$ which leads to lower wages (Proposition 6).

In line with our model, disruption is most stark on online freelance platforms where coordination is least likely to play a role. Here, job volume and earnings fell with the introduction of AI while surviving listings shifted toward complex work (Hui et al., 2024; Demirci et al., 2025). This is consistent with our results on across-firm replacement of solo human producers once machine unit cost falls below theirs (Proposition 5).

[38] We emphasize, however, that this probably understates the economic value of AI teams since a key advantage of using swarms is to speed up the task. This estimate of γ^M does not account for such speedups which are economically valuable.

8 CONCLUSION

If economic transformation does, in fact, depend on how well AI systems coordinate with humans and each other, this can guide measurement. Benchmarks primarily track individual task competence. While these might be good measures of AI model capabilities, they can be misleading signals of economic transformation, because they do not capture how capabilities combine and scale within organizations. It misleads about the present, because organizations are (endogenously) comprised of complementary humans so a model might beat every individual on average but still lose the market. And it can also mislead about the future, because an economy near its tipping point can be reorganized by a coordination improvement too small to register on any benchmark, while an economy far from it can absorb dramatic gains without substantial reorganization of firms.

Indeed, we should be collecting additional data. First, benchmarks should measure how output scales with the number of identical AI copies at fixed compute per copy, alongside isolated competence L. Estimating competence and coordination capabilities across model generations while holding total compute fixed would provide a useful signal for the path of economic transformation. Second, firm-level datasets should measure total organizational scale $m_f = \nu_f^H(\mathcal{Y}) + \nu_f^M(\mathcal{Y})$ which consists of both the worker rolls and the rented compute for running agentic instances. We should track firms' expenditures on their human payroll relative to their AI expenditure over time to understand how their composition is changing with advancing AI capabilities. As we have seen, *within-firm* versus *across-firm* displacement of human workers can have vastly different unemployment dynamics so we should also track the cross-sectional distribution of firms in the economy to understand whether markets are becoming dominated by AI-only or AI-native firms (Kim and Koning, 2026). Finally, we should track employer-employee panels that can directly test how workers with different skill profiles sort across firms. This would isolate the technology effect from the reorganization effect, and help to identify when sectors may be reorganizing around AI as the central driver of productivity, and which workers are most affected.

REFERENCES

D. Acemoglu and P. Restrepo. The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542, 2018.

P. Aghion, B. F. Jones, and C. I. Jones. Artificial intelligence and economic growth. In A. Agrawal, J. Gans, and A. Goldfarb, editors, *The Economics of Artificial Intelligence: An Agenda*, pages 237–282. University of Chicago Press, 2019.

A. K. Agrawal, J. S. Gans, and A. Goldfarb. The Turing transformation: Artificial intelligence, intelligence augmentation, and skill premiums. *Harvard Data Science Review*, 6(Special Issue 5), 2024. doi: 10.1162/99608f92.35a2f3ff.

A. K. Agrawal, J. S. Gans, and A. Goldfarb. Genius on demand: The value of transformative artificial intelligence. In A. K. Agrawal, E. Brynjolfsson, and A. Korinek, editors, *The Economics of Transformative AI*, chapter 1. University of Chicago Press, 2025.

Anthropic. Project Vend: Can Claude run a small shop? (and why does that matter?). Anthropic Research Blog, June 2025a. URL https://www.anthropic.com/research/project-vend-1.

Anthropic. Project Vend: Phase two. Anthropic Research Blog, Dec. 2025b. URL https://www.anthropic.com/research/project-vend-2.

R. J. Aumann. Markets with a continuum of traders. *Econometrica*, 32(1–2):39–50, 1964. doi: 10.2307/1913732.

J. Azar, M. Giné, and J. Sanz-Espín. The wage effects of generative AI. Working paper, University of Navarra and IESE Business School (also circulated as "AI Is Already Eroding Wages: Quasi-Experimental Evidence From Occupational Exposure"), 2026.

A. Backlund and L. Petersson. Vending-bench: A benchmark for long-term coherence of autonomous agents. *arXiv preprint arXiv:2502.15840*, 2025. URL https://arxiv.org/abs/2502.15840.

S. Baslandze, Z. Edwards, J. R. Graham, T. McClure, M. Sparks, B. Meyer, S. R. Waddell, and D. Weitz. Artificial intelligence, productivity, and the workforce: Evidence from corporate executives. Working paper, May 27, 2026. URL https://www.richmondfed.org/-/media/RichmondFedOrg/research/national_economy/cfo_survey/academic_publications/AI_survey.pdf.

W. J. Baumol and R. D. Willig. Fixed costs, sunk costs, entry barriers, and sustainability of monopoly. *Quarterly Journal of Economics*, 96(3):405–431, 1981. doi: 10.2307/1882680.

W. J. Baumol, J. C. Panzar, and R. D. Willig. *Contestable Markets and the Theory of Industry Structure*. Harcourt Brace Jovanovich, New York, 1982.

G. S. Becker and K. M. Murphy. The division of labor, coordination costs, and knowledge. *Quarterly Journal of Economics*, 107(4):1137–1160, 1992.

B. Bertalanič and C. Fortuna. The Ringelmann effect in multi-agent LLM systems. *arXiv preprint arXiv:2606.02646*, 2026. URL https://arxiv.org/abs/2606.02646.

N. Bloom, L. Garicano, R. Sadun, and J. Van Reenen. The distinct effects of information technology and communication technology on firm organization. *Management Science*, 60(12):2859–2885, 2014.

F. P. Brooks, Jr. *The Mythical Man-Month: Essays on Software Engineering*. Addison-Wesley, Reading, MA, 1975. Anniversary edition 1995.

N. Brown. Agent swarms, alignment, and recursive self-improvement. Interview with Dwarkesh Patel, *Dwarkesh Podcast*, Sept. 2026. URL https://www.dwarkesh.com/p/noam-brown.

E. Brynjolfsson, B. Chandar, and R. Chen. Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence. Stanford Digital Economy Lab Working Paper (initial release August 2025; revised August 2026), Aug. 2026.

G. A. Calvo and S. Wellisz. Supervision, loss of control, and the optimum size of the firm. *Journal of Political Economy*, 86(5):943–952, 1978.

M. Cemri, M. Z. Pan, S. Yang, L. A. Agrawal, B. Chopra, R. Tiwari, K. Keutzer, A. Parameswaran, D. Klein, K. Ramchandran, M. Zaharia, J. E. Gonzalez, and I. Stoica. Why do multi-agent LLM systems fail? *arXiv preprint arXiv:2503.13657*, 2025. URL https://arxiv.org/abs/2503.13657.

J. Chen, M. Elliott, and A. Koh. Capability accumulation and conglomeratization in the information age. *Journal of Economic Theory*, 210:105647, 2023. doi: 10.1016/j.jet.2023.105647.

R. H. Coase. The nature of the firm. *Economica*, 4(16):386–405, 1937.

A. Dafoe, E. Hughes, Y. Bachrach, T. Collins, K. R. McKee, J. Z. Leibo, K. Larson, and T. Graepel. Open problems in cooperative AI, 2020. arXiv preprint arXiv:2012.08630.

F. Dell'Acqua, C. Ayoubi, H. Lifshitz, R. Sadun, E. Mollick, L. Mollick, Y. Han, J. Goldman, H. Nair, S. Taub, and K. R. Lakhani. The cybernetic teammate: A field experiment on generative AI reshaping teamwork and expertise. Working Paper 33641, National Bureau of Economic Research, Apr. 2025. URL https://www.nber.org/papers/w33641.

F. Dell'Acqua, E. McFowland, III, E. Mollick, H. Lifshitz, K. C. Kellogg, S. Rajendran, L. Krayer, F. Candelon, and K. R. Lakhani. Navigating the jagged technological frontier: Field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality. *Organization Science*, 37(2):403–423, 2026. doi: 10.1287/orsc.2025.21838.

Ö. Demirci, J. Hannane, and X. Zhu. Who is AI replacing? The impact of generative AI on online freelancing platforms. *Management Science*, 71(10):8097–8108, 2025. doi: 10.1287/mnsc.2024.05420.

H. Demsetz. Why regulate utilities? *The Journal of Law and Economics*, 11(1):55–65, 1968. doi: 10.1086/466643.

W. Dessein and T. Santos. Adaptive organizations. *Journal of Political Economy*, 114(5):956–995, 2006. doi: 10.1086/508031.

M. Ewens and X. Giroud. Corporate hierarchy. Working Paper 34162, National Bureau of Economic Research, Aug. 2025. URL https://www.nber.org/papers/w34162.

J. S. Gans. Endogenous firm ownership in general equilibrium. Working Paper 34187, National Bureau of Economic Research, 2025.

J. S. Gans. Artificial jagged intelligence: When AI benchmarks misstate deployment value. Working Paper 34712, National Bureau of Economic Research, 2026. Also circulated as "A Model of Artificial Jagged Intelligence", arXiv:2601.07573.

L. Garicano. Hierarchies and the organization of knowledge in production. *Journal of Political Economy*, 108(5):874–904, 2000.

L. Garicano and E. Rossi-Hansberg. Organization and inequality in a knowledge economy. *Quarterly Journal of Economics*, 121(4):1383–1435, 2006.

L. Garicano, J. Li, and Y. Wu. Weak bundle, strong bundle: How AI redraws job boundaries. Discussion Paper DP21453, Centre for Economic Policy Research, 2026. See also *Messy Jobs: The Work That AI Cannot Reach* (2026).

R. Gibbons. Four formal(izable) theories of the firm? *Journal of Economic Behavior & Organization*, 58(2):200–245, 2005. doi: 10.1016/j.jebo.2004.09.010.

R. Gibbons and R. Henderson. Relational contracts and organizational capabilities. *Organization Science*, 23(5):1350–1364, 2012. doi: 10.1287/orsc.1110.0715.

T. Graeber, S. Noy, and C. Roth. The transmission of reliable and unreliable information. *Quarterly Journal of Economics*, 141(4):3391–3445, 2026. doi: 10.1093/qje/qjag037.

G. K. Hadfield and A. Koh. An economy of ai agents. In A. K. Agrawal, E. Brynjolfsson, and A. Korinek, editors, *The Economics of Transformative AI*, chapter 5. University of Chicago Press, 2025.

D. Hadfield-Menell and G. K. Hadfield. Incomplete contracting and AI alignment. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pages 417–422. ACM, 2019. doi: 10.1145/3306618.3314250.

M. Hampole, D. Papanikolaou, L. D. W. Schmidt, and B. Seegmiller. Artificial intelligence and the labor market. Working Paper 33509, National Bureau of Economic Research, Feb. 2025. URL https://www.nber.org/papers/w33509.

M. Hoffmann, S. Boysel, F. Nagle, S. Peng, and K. Xu. Generative AI and the nature of work. Working Paper 25-021, Harvard Business School, 2025. Also circulated as CESifo Working Paper No. 11479.

B. Holmström. Moral hazard in teams. *The Bell Journal of Economics*, 13(2):324–340, 1982.

S. M. Hosseini and G. Lichtinger. Generative AI as seniority-biased technological change: Evidence from U.S. résumé and job posting data. SSRN Working Paper 5425555, 2025.

Hugging Face. Security incident disclosure—July 2026. Hugging Face Blog, July 2026. URL https://huggingface.co/blog/security-incident-july-2026.

X. Hui, O. Reshef, and L. Zhou. The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market. *Organization Science*, 35(6):1977–1989, 2024. doi: 10.1287/orsc.2023.18441.

A. Humlum and E. Vestergaard. Still waters, rapid currents: Early labor market transformation under generative AI. Working Paper 33777, National Bureau of Economic Research, May 2025. Revised March 2026.

E. Ide and E. Talamàs. Artificial intelligence in the knowledge economy. *Journal of Political Economy*, 133

(12):3762–3800, 2025. doi: 10.1086/737233.

E. Ide and E. Talamàs. The Turing valley: How AI capabilities shape labor income. *Journal of Human Capital*, 2026. Forthcoming; arXiv:2408.16443.

C. I. Jones. Intermediate goods and weak links in the theory of economic development. *American Economic Journal: Macroeconomics*, 3(2):1–28, 2011.

J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Žídek, A. Potapenko, A. Bridgland, C. Meyer, S. A. A. Kohl, A. J. Ballard, A. Cowie, B. Romera-Paredes, S. Nikolov, R. Jain, J. Adler, T. Back, S. Petersen, D. Reiman, E. Clancy, M. Zielinski, M. Steinegger, M. Pacholska, T. Berghammer, S. Bodenstein, D. Silver, O. Vinyals, A. W. Senior, K. Kavukcuoglu, P. Kohli, and D. Hassabis. Highly accurate protein structure prediction with AlphaFold. *Nature*, 596:583–589, 2021. doi: 10.1038/s41586-021-03819-2.

N. Kaldor. The equilibrium of the firm. *Economic Journal*, 44(173):60–76, 1934.

H. Kim and R. Koning. AI-Native Firms. Working Paper 26-090, Harvard Business School, 2026. URL https://www.hbs.edu/faculty/Pages/item.aspx?num=69077.

Y. Kim, K. Gu, C. Park, C. Park, S. Schmidgall, A. A. Heydari, Y. Yan, Z. Zhang, Y. Zhuang, Y. Liu, M. Malhotra, P. P. Liang, H. W. Park, Y. Yang, X. Xu, Y. Du, S. Patel, T. Althoff, D. McDuff, and X. Liu. Towards a science of scaling agent systems. *arXiv preprint arXiv:2512.08296*, 2025. URL https://arxiv.org/abs/2512.08296.

T. C. Koopmans. Analysis of production as an efficient combination of activities. In T. C. Koopmans, editor, *Activity Analysis of Production and Allocation*, pages 33–97. Wiley, New York, 1951.

A. Korinek and D. Suh. Scenarios for the transition to AGI. Working Paper 32255, National Bureau of Economic Research, Mar. 2024. URL https://www.nber.org/papers/w32255.

M. Kremer. The O-ring theory of economic development. *Quarterly Journal of Economics*, 108(3):551–575, 1993.

T. Kwa, B. West, J. Becker, A. Deng, K. Garcia, M. Hasin, S. Jawhar, M. Kinniment, N. Rush, S. Von Arx, R. Bloom, T. Broadley, H. Du, B. Goodrich, N. Jurkovic, L. Miles, S. Nix, T. Lin, N. Parikh, D. Rein, L. J. K. Sato, H. Wijk, D. Ziegler, E. Barnes, and L. Chan. Measuring AI ability to complete long software tasks. In *Advances in Neural Information Processing Systems*, volume 38, 2025. doi: 10.52202/085713-3086. URL https://proceedings.neurips.cc/paper_files/paper/2025/hash/85069585133c4c168c865e65d72e9775-Abstract-Conference.html.

S. Legg and M. Hutter. Universal intelligence: A definition of machine intelligence. *Minds and Machines*, 17(4):391–444, 2007.

C. List and P. Pettit. *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford University Press, 2011. doi: 10.1093/acprof:oso/9780199591565.001.0001.

R. E. Lucas, Jr. On the size distribution of business firms. *The Bell Journal of Economics*, 9(2):508–523, 1978.

J. Marschak and R. Radner. *Economic Theory of Teams*. Yale University Press, New York and London, 1972.

M. J. Melitz. The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, 71(6):1695–1725, 2003. doi: 10.1111/1468-0262.00467.

METR and Redwood Research. Brief independent investigation of agents' behavior, reasoning and collaboration in the OpenAI / Hugging Face hacking incident. METR Technical Report, Aug. 2026. URL https://metr.org/blog/2026-08-26-openai-hugging-face-incident-investigation/.

W. Novshek and H. Sonnenschein. General equilibrium with free entry: A synthetic approach to the theory of perfect competition. *Journal of Economic Literature*, 25(3):1281–1306, September 1987. URL https://www.jstor.org/stable/2726028.

OpenAI. Introducing GPT-5.2, Dec. 2025. URL https://openai.com/index/introducing-gpt-5-2/.

OpenAI. The Hugging Face incident and the road ahead. OpenAI Technical Report, Aug. 2026. URL https://openai.com/index/the-hugging-face-incident-and-the-road-ahead/.

T. Ord. Swarm scaling. LessWrong, Sept. 2026. URL https://www.lesswrong.com/posts/6cb7qd3RSkgnviCpf/swarm-scaling. Crossposted on the EA Forum: https://forum.effectivealtruism.org/posts/EkaBCQwL5hhTtxEwR. Accessed 10 October 2026.

J. M. Ostroy. The no-surplus condition as a characterization of perfectly competitive equilibrium. *Journal of Economic Theory*, 22(2):183–207, 1980. doi: 10.1016/0022-0531(80)90039-3.

D. Patel. The rise and fall of agent civilizations. Dwarkesh Podcast, Aug. 2026. URL https://www.dwarkesh.com/p/openai-huggingface.

T. Patwardhan, R. Dias, E. Proehl, G. Kim, M. Wang, O. Watkins, S. Posada Fishman, M. Aljubeh, P. Thacker, L. Fauconnet, N. S. Kim, P. Chao, S. Miserendino, G. Chabot, D. Li, M. Sharman, A. Barr, A. Glaese, and J. Tworek. GDPval: Evaluating AI model performance on real-world economically valuable tasks. *arXiv preprint arXiv:2510.04374*, 2025. URL https://arxiv.org/abs/2510.04374.

M. Polanyi. *The Tacit Dimension*. Doubleday, Garden City, NY, 1966.

A. Prat. Should a team be homogeneous? *European Economic Review*, 46(7):1187–1207, 2002. doi: 10.1016/S0014-2921(01)00165-9.

Y. Qian. Incentives and loss of control in an optimal hierarchy. *Review of Economic Studies*, 61(3):527–544, 1994.

R. Radner. The organization of decentralized information processing. *Econometrica*, 61(5):1109–1146, 1993.

P. Restrepo. We won't be missed: Work and growth in the AGI world. In A. K. Agrawal, E. Brynjolfsson, and A. Korinek, editors, *The Economics of Transformative AI*, chapter 9. University of Chicago Press, 2026. NBER conference volume, forthcoming.

J. G. Riley. Informational equilibrium. *Econometrica*, 47(2):331–359, 1979.

P. A. Samuelson. *Economics: An Introductory Analysis*. McGraw-Hill Book Company, New York, 1948.

P. Shahidi, G. Rusak, B. S. Manning, A. Fradkin, and J. J. Horton. The Coasean singularity? Demand, supply, and market design with AI agents. In A. K. Agrawal, E. Brynjolfsson, and A. Korinek, editors, *The Economics of Transformative AI*, chapter 6. University of Chicago Press, 2025. Also circulated as NBER Working Paper No. 34468.

N. Shapira, C. Wendler, A. Yen, G. Sarti, K. Pal, O. Floody, A. Belfki, A. R. Loftus, A. R. Jannali, N. Prakash, et al. Agents of chaos. *arXiv preprint arXiv:2602.20021*, 2026. URL https://arxiv.org/abs/2602.20021.

K.-I. Shimomura and J.-F. Thisse. Competition among the big and the small. *The RAND Journal of Economics*, 43(2):329–347, 2012. doi: 10.1111/j.1756-2171.2012.00168.x.

H. A. Simon. *Administrative Behavior*. Macmillan, New York, 1947.

H. A. Simon. Designing organizations for an information-rich world. In M. Greenberger, editor, *Computers, Communications, and the Public Interest*, pages 37–72. Johns Hopkins Press, Baltimore, 1971.

G. Starace, O. Jaffe, D. Sherburn, J. Aung, J. S. Chan, L. Maksin, R. Dias, E. Mays, B. Kinsella, W. Thompson, J. Heidecke, A. Glaese, and T. Patwardhan. PaperBench: Evaluating AI's ability to replicate AI research. *arXiv preprint arXiv:2504.01848*, 2025. URL https://arxiv.org/abs/2504.01848.

Vals AI. Vals legal AI report (VLAIR): Benchmarking legal AI against a human lawyer baseline. Vals AI Research Report, 2025. URL https://www.vals.ai/industry-reports/vlair-10-14-25.

B. Vidgen and Mercor Research. APEX-Agents: The AI productivity index for agents. *arXiv preprint arXiv:2601.14242*, 2026. URL https://arxiv.org/abs/2601.14242.

B. Vidgen, C. R. Sunstein, E. Topol, and Mercor Research. The AI productivity index (APEX). *arXiv preprint arXiv:2509.25721*, 2025. URL https://arxiv.org/abs/2509.25721.

B. Wernerfelt. A resource-based view of the firm. *Strategic Management Journal*, 5(2):171–180, 1984.

H. Wijk, T. Lin, J. Becker, S. Jawhar, N. Parikh, T. Broadley, L. Chan, M. Chen, J. Clymer, J. Dhyani, E. Ericheva, K. Garcia, B. Goodrich, N. Jurkovic, M. Kinniment, A. Lajko, S. Nix, L. Sato, W. Saunders, M. Taran, B. West, and E. Barnes. RE-Bench: Evaluating frontier AI R&D capabilities of language model agents against human experts. *arXiv preprint arXiv:2411.15114*, 2024. URL https://arxiv.org/abs/2411.15114.

O. E. Williamson. Hierarchical control and optimum firm size. *Journal of Political Economy*, 75(2):123–138, 1967.

A. W. Woolley, C. F. Chabris, A. Pentland, N. Hashmi, and T. W. Malone. Evidence for a Collective Intelligence Factor in the Performance of Human Groups. *Science*, 330(6004):686–688, 2010. doi: 10.1126/science.1193147.

Y. Yang, C. Qu, M. Wen, L. Shi, Y. Wen, W. Zhang, A. Wierman, and S. Gu. Understanding agent scaling in LLM-based multi-agent systems via diversity. *arXiv preprint arXiv:2602.03794*, 2026. URL https://arxiv.org/abs/2602.03794.

I. Yotzov, J. M. Barrero, N. Bloom, P. Bunn, S. J. Davis, K. M. Foster, A. Jalca, B. H. Meyer, P. Mizen, M. A. Navarrete, P. Smietanka, G. Thwaites, and B. Z. Wang. Firm data on AI. Working Paper 34836, National Bureau of Economic Research, 2026. Revised March 2026.

W. R. Zame. Incentives, contracts, and markets: A general equilibrium theory of firms. *Econometrica*, 75 (5):1453–1500, 2007. doi: 10.1111/j.1468-0262.2007.00799.x.

APPENDIX

A Omitted Proofs 59
A.1 Equilibrium and efficiency 59
A.2 Industry structure 63
A.3 Wages and employment 74
A.4 Connection to the task model 79
A.5 Continuity of values and wages 85
B Equilibrium concept and firm counts 88
C Connection to the task model 90
D CES examples and numerical methods 93
E Calculations for the example 102

A OMITTED PROOFS

We collect the proofs in the order of the main results.

A.1 Equilibrium and efficiency. For the proofs, let $c_{\min} := \min\{r, \underline{w}\}$. Fix an atomistic machine-only firm $\bar{f}^M \in \mathcal{F}^0$ with the lowest unit cost, and let

$$\bar{c} := \frac{C(\bar{f}^M)}{q_{\bar{f}^M}}.$$

Unlimited replication of this atomistic firm gives $\mathcal{C}(Q) \leq \bar{c}Q$ for every Q.

A.1.1 Industry cost and duality

Exact output. Given a feasible industry with output $x > Q > 0$, we first shrink each positive-scale firm with $q_f > Q$ to the first scale on its composition ray that produces Q, ensuring that the complete-firm feasibility constraint remains satisfied. Continuity and $F(0) = 0$ ensure existence and a measurable choice; other firms stay unchanged. Since $q_f \leq x$ and $\min\{q_f, Q\} \geq (Q/x)q_f$, the industry still produces at least Q. We can then proportionally reduce activity weights to obtain exactly Q. Both steps preserve feasibility and weakly reduce resource use, with a strict reduction in total cost. Thus, allowing production above Q cannot lower the minimum cost relative to requiring exactly Q units.

A compact activity set. Represent a recipe by its scale $m \geq 0$ and a unit-normalized firm $f \in \mathcal{F}^0$. Read $F(mv_f)/m$ as $L(v_f)$ at zero. The labelled unit compositions in $\mathcal{F}^0$ form a compact set; with one machine profile, this set is a simplex. Replacing any

recipe with resource unit cost above $\bar{c}$ by atomistic machines lowers cost and human use. We can therefore restrict attention to

$$\mathcal{A}_Q := \left\{ (m, f) \in [0, \bar{c}Q/c_{\min}] \times \mathcal{F}^0 : \begin{array}{l} m\nu_f^H \le \mu, \quad F(m\nu_f) \le Q, \\ C(f) \le \bar{c}\, F(m\nu_f)/m \end{array} \right\}.$$

This set is compact by Assumption 1(i), as the inequalities define a closed subset of a compact one. Its cone of aggregate output, human inputs, and cost is

$$\mathcal{K}_Q := \operatorname{cone}\left\{ \left(\frac{F(m\nu_f)}{m}, \left(\nu_f^H(\{\boldsymbol{y}^i\})\right)_{i=1}^K, C(f) \right) : (m, f) \in \mathcal{A}_Q \right\}.$$

A coefficient t on a generator means total input t: it gives firm weight t/m at $m > 0$, or atomistic intensity t at $m = 0$. Every generator has cost at least $c_{\min}$. Bounded aggregate cost therefore bounds the total coefficient. Compactness of the generating set then makes the cone closed and ensures that industry integrals belong to it, by approximation with finite sums.

Lemma 1 (Industry cost and supporting prices)**.** For every output $Q > 0$:

(i) Minimum cost $\mathcal{C}(Q)$ is positive and is attained at exact output by an industry with at most $K + 3$ firm types.

(ii) Every cost minimizer σ_Q has a nonempty, convex, compact support set $\Lambda(Q, \sigma_Q)$. This set is the same for all cost minimizers at Q. Every supported pair satisfies

$$p \le \bar{c}, \qquad pQ - \sum_{i=1}^{K} (w_i - \underline{w})\mu_i = \mathcal{C}(Q).$$

Conversely, a feasible exact-output industry with a supporting pair minimizes cost.

(iii) Minimum cost $\mathcal{C}$ is continuous on $\mathbb{R}_+$.

Proof. We prove attainment, obtain the supporting prices by duality, and then establish continuity.

Step 1: attainment. The activity representation gives

$$\mathcal{C}(Q) = \min\{z : (x, H, z) \in \mathcal{K}_Q,\ x \ge Q,\ H_i \le \mu_i\ \forall i\},$$

where z is total industry cost and H_i is aggregate employment of type i. Scaling all coefficients down makes output exact without changing the firm menu. The

feasible cost bound $\bar{c}Q$ and the coefficient bound above give a compact minimizing set. Bounded output per unit input and positive input costs make its minimum positive. Carathéodory's theorem gives a minimizer with at most $K+3$ generators, each of which implements a firm type.

Step 2: duality. Scale a relative-interior point of $\mathcal{K}_Q$ until all human inputs are strictly below supply, then add enough atomistic machine output to exceed Q. The result stays in the relative interior and satisfies the inequality constraints strictly. Conic strong duality therefore applies. With output multiplier $p \geq 0$ and human multipliers $w_i - \underline{w} \geq 0$, the Lagrangian is

$$pQ - \sum_i (w_i - \underline{w})\mu_i + z - px + \sum_i (w_i - \underline{w})H_i.$$

Its infimum over the cone is finite exactly when firm support holds on $\mathcal{A}_Q$. Atomistic machine entry then gives $p \leq \bar{c}$. Since every feasible recipe with resource unit cost at most $\bar{c}$ belongs to $\mathcal{A}_Q$, this also rules out profitable entry by discarded recipes, as wages are no lower than reservation costs.

Strong duality yields the displayed identity, while complementary slackness implies that active firms break even and human types earn scarcity premiums only when their supply binds. These conditions hold for every primal and dual optimum. Conversely, integrating firm support shows that any feasible industry producing exactly Q, with active firms breaking even and human scarcity satisfied, minimizes resource cost. This proves the characterization in (ii).

Finally, the bounds $0 \leq (w_i - \underline{w})\mu_i \leq \bar{c}Q$, together with closedness of the dual inequalities and equality, give compactness; convexity follows from their linearity.

Step 3: continuity. For $Q_n > Q$, add atomistic machine output to a minimizer at Q, at cost $\bar{c}(Q_n - Q)$. For $Q_n < Q$, the exact-output construction above produces Q_n at no greater cost. These comparisons give upper semicontinuity. For the reverse inequality, choose minimizers at Q_n with at most $K+3$ generators and pad the lists with zero coefficients. Their scales and total inputs are uniformly bounded. Along a convergent subsequence, continuity of $F(mv)/m$, including at zero, preserves output, individual and aggregate human feasibility, and the market cap. The limiting generators implement an industry at Q with the limiting cost. Hence $\mathcal{C}(Q) \leq \liminf_n \mathcal{C}(Q_n)$. At zero, $0 \leq \mathcal{C}(Q) \leq \bar{c}Q$ proves continuity. □

A.1.2 Implementation and equilibrium

Competitive implementation in Appendix B. A producer can increase the intensity of a profitable recipe without bound. An optimal action therefore requires firm support,

and only zero-profit recipes can operate. Human labor choice gives exactly human scarcity.

Conversely, let $(p, w) \in \Lambda(Q, \sigma)$. The firm space is standard Borel, so assign recipes measurably on $[0, 1/2]$ with distribution $\sigma/\sigma(\mathcal{F})$ under normalized Lebesgue measure. Give each assigned producer intensity $2\sigma(\mathcal{F})$ and leave the rest inactive. The resulting industry is σ. All actions are optimal by support and active equality; employment clears and is optimal by human scarcity. □

Proof of Theorem 1. We show that a stable contract must be efficient, support every efficient industry, and then prove existence. Write $H_i(\sigma)$ for employment of human type i.

Step 1: stability implies efficiency. An equilibrium contract has $(p, w) \in \Lambda(Q(\sigma), \sigma)$ by competitive implementation. Its organizer receives zero: otherwise an inactive organizer can copy it and gain strictly. The accounting identity in Appendix B makes the protected parties' total payoff equal to $\mathcal{W}_N(\sigma)$.

If a feasible industry gives higher surplus, replace it by a cost minimizer at its output. Lemma 1 supplies supporting prices, so the replacement admits a competitive continuation. Choose transfers that give each protected party its old payoff and give the remaining surplus to the new organizer. The accounting identity makes these transfers balance. The offer is admissible and blocks the original contract. Hence every stable industry is efficient.

Step 2: efficiency implies stability. Let σ be efficient with $(p, w) \in \Lambda(Q(\sigma), \sigma)$. Set

$$t_i = -(w_i - \underline{w})H_i(\sigma), \qquad t_0 = \sum_i (w_i - \underline{w})H_i(\sigma), \qquad t_E = 0.$$

These transfers give each human coalition its outside payoff and give the consumer representative all surplus. Machine viability makes that surplus positive. The contract is admissible by competitive implementation. A blocking offer would preserve these payoffs and pay its organizer strictly more than zero, so its total surplus would exceed the optimum. Thus the contract is stable.

Step 3: existence and shadow values. Efficiency is equivalent to minimizing cost at a quantity in $\mathcal{Q}_N^*$: either a cost saving at fixed output or a better quantity raises surplus. Continuity of $U_N - \mathcal{C}$ on the compact demand interval gives a maximizing quantity. Machine viability makes it positive. Lemma 1 supplies an exact-output minimizer and supporting prices there; Step 2 gives equilibrium. The same lemma identifies p as the output multiplier and $w_i - \underline{w}$ as the multiplier on aggregate human supply, with the complete-firm menu held fixed. □

A.1.3 Competitive support

Competitive equilibrium implies equilibrium of our model. We verify price support, then use consumer optimization to prove efficiency.

Step 1: supporting prices. Fix a CE (σ, p, w) whose industry satisfies our complete-firm conditions, and write $Q = Q(\sigma)$ for its output and $H_i(\sigma)$ for type i's employment. Unrestricted activity intensities imply that every recipe earns weakly negative profit and active recipes earn zero profit. The maintained wage floor gives $w_i \geq \underline{w}$; worker optimization and clearing give $(w_i - \underline{w})[\mu_i - H_i(\sigma)] = 0$. Thus $(p, w) \in \Lambda(Q, \sigma)$ once we verify that $Q > 0$. At $Q = 0$, consumer optimization requires $p \geq P_N(0)$, but the standing machine-viability condition then makes an atomistic machine-only recipe profitable. Hence $Q > 0$. Since all inputs have positive prices, $p = 0$ would imply zero production. Therefore $p > 0$, and consumer optimization gives $p = P_N(Q)$.

Step 2: efficiency. For any alternative industry $\widehat{\sigma} \in \Sigma$, let $\widehat{Q} = Q(\widehat{\sigma})$ denote its output. Integrating the CE profit inequalities over this industry and subtracting active equality gives

$$\int_{\mathcal{F}} C(f)\ \mathrm{d}\widehat{\sigma}(f) - \int_{\mathcal{F}} C(f)\ \mathrm{d}\sigma(f) \geq p(\widehat{Q} - Q) - \sum_i (w_i - \underline{w})[H_i(\widehat{\sigma}) - H_i(\sigma)].$$

Consumer optimization gives $U_N(\widehat{Q}) - U_N(Q) \leq p(\widehat{Q} - Q)$. Consequently,

$$\mathcal{W}_N(\widehat{\sigma}) - \mathcal{W}_N(\sigma) \leq \sum_i (w_i - \underline{w})[H_i(\widehat{\sigma}) - H_i(\sigma)] \leq 0.$$

The final inequality follows from human feasibility and complementary slackness: each positive wage premium requires $H_i(\sigma) = \mu_i$. Thus σ is efficient. Together with $(p, w) \in \Lambda(Q, \sigma)$, Theorem 1 implies that (σ, w) is an equilibrium of our model.

□

A.2 Industry structure.

A.2.1 Firm scale and the primitive bounds

Proof of Proposition 1. We first separate preferred scale from competence and order it by coordination. We then compare an active firm with feasible copies of its composition. These comparisons yield the bounds.

Step 1: preferred scale. Under Assumption 2, for a unit composition ν with mean coordination z_*, average product is $L(\nu)S(m, z_*)/m$. The positive factor $L(\nu)$ does not affect the maximizing scale set. At fixed shares of each type, a competence change leaves z_* and this set unchanged. The extension of $S(m, z_*)/m$ at zero is one, by the definition of L.

For two coordination levels $z'_* > z_*$, define their log output ratio at positive scale by

$$D(m) := \log S(m, z'_*) - \log S(m, z_*) = \int_{z_*}^{z'_*} \partial_2 \log S(m, u) \, \mathrm{d}u.$$

The larger-firms condition makes D strictly increasing in m; its continuous extension at zero is zero. If a preferred scale at z'_* were smaller than a preferred scale at z_*, adding their optimality inequalities would make D weakly larger at the smaller scale, a contradiction. This proves the preferred-size comparison in the main text, including zero preferred sizes. The comparison does not require unique preferred scales or imply strictly positive scale responses.

Step 2: active scales and constraints. Fix an equilibrium and supporting prices (p, w). For an active firm with unit composition ν, positive input prices and active equality give $p > 0$. Firm support gives

$$p \, \frac{F(m\nu)}{m} \leq r\nu^M(\mathcal{Y}) + \int_{\mathcal{Y}^H} w \, \mathrm{d}\nu^H$$

for every feasible copy, with equality at the active scale. At zero scale, read the ratio as $L(\nu)$. Since $p > 0$ and input cost per worker is constant at a fixed mix, the active scale maximizes feasible average product. This proves the size-choice formula in Section 5.1.

Every smaller copy can be staffed. If it produced more than the active firm, continuity along the ray would give a still smaller copy producing exactly q_f at a strictly lower bill, contrary to support. Thus all smaller copies fit in the market. Their average products cannot exceed the active firm's. If a finite preferred scale lies below the active scale, the active scale must therefore also be preferred. Thus an active firm with a finite preferred scale has size no larger than some preferred scale of its mix. If any finite preferred scale is feasible, the same support inequality puts its average product below the active firm's, so the active scale is also preferred. This includes preferred or active scale zero.

Under strict quasiconcavity, average product rises strictly below all preferred scales. If no finite maximum exists, it rises strictly at every finite scale: a weak decline between two scales would, by strict quasiconcavity, force all later values below the earlier value and yield a finite maximum by continuity. An atomistic firm below preferred scale can be enlarged slightly while preserving staffing and market fit. A positive-scale

firm with all constraints slack also admits a slightly larger feasible copy. Both cases contradict firm support. A firm below all preferred scales must therefore have positive scale and be capped. This establishes the constraint implications stated in the main text.

Step 3: bounds across firms. Under Assumption 2, for an active positive-scale firm with unit composition ν, the scale comparison in Step 2 and staffing feasibility give

$$m_f \leq \sup m^* \left(\int_{\mathcal{Y}} y_* \, \mathrm{d}\nu \right) \wedge \sup\{m : m\nu^H \leq \mu\} \leq \overline{M}.$$

When no finite preferred scale exists, the first term is infinite by convention, and staffing feasibility still gives the bound. Atomistic firms satisfy the upper bound at size zero.

For the lower bound, a firm at or above some preferred scale has size at least $\underline{M}$. Under strict quasiconcavity, every other active firm is positive-scale and capped by Step 2. The slack market constraint makes its size equal to its staffing limit, which is again at least $\underline{M}$. Thus the bounds hold almost everywhere in each equilibrium under the stated conditions.

For the interpretation in the text, Step 1 bounds every composition's smallest preferred scale below by the worst coordinator's smallest preferred scale. Every staffing limit is at least the smallest human supply. A pure composition attains the smaller of these two bounds. The AI-only composition gives its largest preferred scale as a lower bound on $\overline{M}$; the best coordinator's largest preferred scale bounds it above. This completes the proof. □

Coordination and the tail. Under Assumption 2, the tail exponent depends only on mean coordination and is nondecreasing in it. A composition whose tail exponent is below one has a nonempty bounded set of preferred scales. If its exponent exceeds one, no finite scale is preferred. We make no such claim at exponent one.

To verify these claims, the uniform power tail gives

$$S(m, z_*) = \frac{A(\nu)}{L(\nu)} m^{\gamma(\nu)} [1 + o(1)]$$

for every unit composition ν with mean coordination z_*. The left side depends on composition only through z_*, so its exponent does too. If a higher coordination level had a smaller exponent, the log output ratio $D(m)$ in Step 1 would tend to $-\infty$. Its increasing shape rules out this case. Below exponent one, average product tends to zero, so continuity gives a nonempty bounded maximizing set that may include zero. Above exponent one, average product is unbounded, so no finite maximum exists.

A.2.2 Competence and coordination

Proposition 8 (Better AI lowers cost and raises surplus)**.** If the AI improves from $\boldsymbol{y}^M$ to $\widetilde{\boldsymbol{y}}^M \geq \boldsymbol{y}^M$, then minimum resource cost weakly falls at every output quantity, and equilibrium total surplus weakly rises at every market size.

Proof of Proposition 8. We reproduce the old output at lower cost. Take any feasible industry with output $Q > 0$. For each positive-scale firm f, let $\widetilde{\nu}_f$ describe the firm after we upgrade its machines without changing any input quantities. The upgraded firm has the same size and dominates the old one, so equal-mass monotonicity gives $F(\widetilde{\nu}_f) \geq q_f$. Continuity and $F(0) = 0$ give a scale factor $t_f \in (0, 1]$ at which the replacement firm $\widetilde{f}$ satisfies

$$F(t_f \widetilde{\nu}_f) = q_f, \qquad C(\widetilde{f}) = t_f C(f) \leq C(f).$$

Choose the first scale factor that attains q_f; continuity makes this a measurable choice. Keep each firm's activity weight unchanged. Its output still fits within Q, and its human use weakly falls, both individually and in aggregate. For each atomistic activity, the upgrade weakly raises output per input; reduce its activity weight to preserve its old output. Thus the upgraded economy can reproduce Q at weakly lower cost. Taking the infimum proves the cost comparison; at $Q = 0$, both costs are zero. The gross benefit of each quantity is unchanged, so maximized surplus weakly rises. Theorem 1 gives the equilibrium claim. □

Write $\mathcal{W}_N^*$ for equilibrium total surplus and $\overline{\boldsymbol{y}}_f$ for firm f's mean productive capabilities.

If Assumption 2 holds, y_d^M is interior to its feasible range, $\mathcal{W}_N^*$ is differentiable in y_d^M, and the market constraint is slack for every firm $f \in \operatorname{supp}(\sigma|_{\mathcal{F}^+})$ in an equilibrium industry σ, then

$$\frac{\mathrm{d}\mathcal{W}_N^*}{\mathrm{d}y_d^M} = P_N(Q) \int_{\mathcal{F}} \underbrace{q_f}_{\text{firm output}} \underbrace{\frac{\nu_f^M(\mathcal{Y})}{m_f}}_{\text{machine share}} \underbrace{\frac{\partial_d L(\overline{\boldsymbol{y}}_f)}{L(\overline{\boldsymbol{y}}_f)}}_{\text{value of domain } d} \mathrm{d}\sigma(f). \qquad \text{(SURPLUS)}$$

Derivation of the surplus formula (SURPLUS). We show that the fixed organization stays feasible locally, then differentiate its surplus.

Step 1: local feasibility. Fix supporting prices and put $Q = Q(\sigma)$. Active equality gives $c_{\min} m_f \leq p q_f \leq \bar{c} Q$. Normalize the supported firms and include their zero-scale

limits, as in Section A.1.1. The resulting closure is compact. Every positive-scale point satisfies $q_f < Q$ by hypothesis, while complete-firm output tends uniformly to zero at zero scale. Hence the complete-firm capacities have a uniform margin below Q.

At fixed inputs and weights, a competence change ε gives

$$q_f(\varepsilon) = q_f \frac{L(\overline{\boldsymbol{y}}_f + \varepsilon[\nu_f^M(\mathcal{Y})/m_f]\boldsymbol{e}_d)}{L(\overline{\boldsymbol{y}}_f)},$$

where $\boldsymbol{e}_d$ is the domain-d unit vector. The formula also holds for atomistic activities. The ratio converges uniformly to one, so total output converges to Q, and complete-firm outputs change uniformly on the compact set. The market margin and unchanged human inputs make the fixed organization feasible for small changes of either sign.

Step 2: the derivative. Let $G(\varepsilon)$ be its surplus. Optimality gives $\mathcal{W}_N^*(y_d^M + \varepsilon) \geq G(\varepsilon)$ with equality at zero. Differentiability therefore equates their derivatives. Input costs stay fixed. On the compact set of means, L is bounded away from zero and $\partial_d L$ is bounded, so the output derivative is bounded by a constant times q_f. Since $\int q_f \, \mathrm{d}\sigma = Q$, we can differentiate under the integral:

$$G'(0) = P_N(Q) \int_{\mathcal{F}} q_f \frac{\nu_f^M(\mathcal{Y})}{m_f} \frac{\partial_d L(\overline{\boldsymbol{y}}_f)}{L(\overline{\boldsymbol{y}}_f)} \, \mathrm{d}\sigma(f).$$

This is the stated formula. □

With one human type, a firm with machine share s has unit composition $\nu_s = (1-s)\delta_{\boldsymbol{y}^H} + s\delta_{\boldsymbol{y}^M}$ and mean productive capabilities $\overline{y}_d = (1-s)y_d^H + sy_d^M$ in each domain d. At a fixed human wage w, define its log output per unit of input cost, after choosing scale, by

$$g(s, y_d^M) := \log\left[\max_{m \geq 0} \frac{F(m\nu_s)}{m}\right] - \log\big((1-s)w + sr\big).$$

If Assumption 2 holds, consider a continuous selection of active interior firms with machine shares s as y_d^M varies. Suppose the human wage stays fixed, nearby compositions have feasible preferred scales at the respective equilibrium output, and g is twice

continuously differentiable with $g_{ss} < 0$ along the selection. Then

$$\frac{\mathrm{d}s}{\mathrm{d}y_d^M} = \frac{\overbrace{\frac{\partial}{\partial s}\left[s\,\frac{\partial_d L(\overline{\boldsymbol{y}})}{L(\overline{\boldsymbol{y}})}\right]}^{\substack{\text{how much the upgrade favors}\\ \text{a higher machine share}}}}{\underbrace{-g_{ss}}_{\substack{\text{curvature at the}\\ \text{preferred mix}}}}. \tag{MIX}$$

The term in brackets is the proportional productivity gain at a fixed mix. Since $g_{ss} < 0$, the machine share rises exactly when mixes with more machines gain proportionally more. More machines expose a firm more to the upgrade, but also change which capabilities it lacks. The assumptions do not sign the net effect. Fixed wages matter: changes in wages would also change relative input costs.

Derivation of the composition formula (MIX). Firm support and active equality make the selected share a local maximizer of g: its preferred scale and those at nearby shares are feasible, and its output per unit cost equals $1/p$. Thus $g_s = 0$. At fixed share, coordination and input costs do not change with productive capability, so

$$g_{y_d^M} = s\,\frac{\partial_d L(\overline{\boldsymbol{y}})}{L(\overline{\boldsymbol{y}})}.$$

Since $g \in C^2$ and $g_{ss} < 0$, the implicit function theorem gives $\mathrm{d}s/\,\mathrm{d}y_d^M = -g_{sy_d^M}/g_{ss}$ along the continuous selection. Substitution proves the formula. □

Proof of Proposition 2. We identify the active shares, sign their response to coordination, and then differentiate preferred size.

Step 1: equilibrium firms maximize output per unit cost. At machine share $s \in [0,1]$, mean coordination is $\overline{y}_* = (1-s)y_*^H + sy_*^M$. The CES scale factor gives

$$\partial_m \log[S(m,\overline{y}_*)/m] = \frac{\phi - (1-\overline{y}_*)(1+m)}{(1+m)^2}.$$

Since $\overline{y}_* < 1$, the assumed positive preferred scale is unique and equals $m^*(\overline{y}_*) = \phi/(1-\overline{y}_*) - 1$. The assumed positivity holds at every share along the path. Feasibility and the size-choice argument in Section A.2.1 put every active firm at its preferred scale and exclude atomistic production.

Slack aggregate human supply gives $w = \underline{w}$ by human scarcity. Firm support and active equality in Theorem 1 then imply that every active share maximizes output per

unit input cost: all preferred-scale alternatives can be staffed and fit in the market. Using the earlier notation g for its logarithm, we have

$$\begin{aligned} g(s, y_*^M) = \; & \log L(\overline{\boldsymbol{y}}) + \phi - (1 - \overline{y}_*) + (1 - \overline{y}_*) \log \frac{1 - \overline{y}_*}{\phi} \\ & - \log\big((1-s)\underline{w} + sr\big), \end{aligned}$$

where $\overline{\boldsymbol{y}}$ denotes mean productive capabilities. At a fixed share, these capabilities and input costs do not change with AI coordination.

Step 2: the share response. Differentiating gives

$$g_{sy_*^M} = \log \frac{\phi}{1 - \overline{y}_*} + \frac{s(y_*^M - y_*^H)}{1 - \overline{y}_*}.$$

If $y_*^M \geq y_*^H$, this derivative is positive: preferred size is positive, and the second term is nonnegative. Let s and $\widetilde{s}$ maximize g before and after a coordination gain within this region. If $\widetilde{s} < s$, integration gives

$$g(s, \widetilde{y}_*^M) - g(\widetilde{s}, \widetilde{y}_*^M) > g(s, y_*^M) - g(\widetilde{s}, y_*^M) \geq 0,$$

contrary to optimality of $\widetilde{s}$. Thus $\widetilde{s} \geq s$, including at corners. Mean coordination and preferred size also weakly increase. This argument needs neither uniqueness nor continuity.

If $y_*^M < y_*^H$, the mixed derivative can have either sign. At an interior optimum with $g_{ss} < 0$, the implicit function theorem gives

$$\frac{\mathrm{d}s}{\mathrm{d}y_*^M} = -\frac{g_{sy_*^M}}{g_{ss}}.$$

Since $\log[\phi/(1 - \overline{y}_*)] = \log(1 + m^*)$, this proves the three sign cases in part (ii).

Step 3: the size response. Differentiate $m^* = \phi/(1 - \overline{y}_*) - 1$ along the selected firm:

$$\frac{\mathrm{d}m}{\mathrm{d}y_*^M} = \frac{\phi}{(1 - \overline{y}_*)^2} \left[s + (y_*^M - y_*^H) \frac{\mathrm{d}s}{\mathrm{d}y_*^M} \right].$$

The first term is nonnegative. The sign of the second follows from Step 2 and the coordination gap, which gives the stated amplification and dampening cases. □

A.2.3 Large markets

Equilibrium output grows with the market. Write $Q_N^* = Q(\sigma_N^*)$. Machine viability gives $\bar{c} < P(0)$. Choose a fixed quantity $z > 0$ with $P(z) > \bar{c}$. Atomistic machines can produce Nz at cost $\bar{c}Nz$, so efficiency gives

$$P(0)Q_N^* \geq \mathcal{W}_N^* \geq N \int_0^z [P(u) - \bar{c}] \, \mathrm{d}u.$$

The integral is strictly positive. Thus Q_N^* grows at least linearly with N, for every value of γ^M, including one.

Lemma 2 (Subcritical firms with bounded unit cost have bounded size)**.** If $\gamma^M < 1$, then, for every finite B, feasible positive-scale firms with $C(f)/q_f \leq B$ have uniformly bounded sizes and capacities.

Proof. As feasible firm size m_f diverges, its human share is at most $\mu(\mathcal{Y}^H)/m_f$, so its composition converges to the machine profile. Choose $\delta > 0$ with $\gamma^M + \delta < 1$. Continuity of the exponent and the uniform power tail imply $q_f/m_f = O(m_f^{\gamma^M+\delta-1}) \to 0$ uniformly. Since $C(f) \geq c_{\min} m_f$, unit cost then diverges. A finite unit-cost bound therefore bounds size, and continuity of output on firms of bounded size bounds capacity. □

Lemma 3 (Supercritical concentration)**.** Suppose $\gamma^M > 1$. For any sequence $Q_\ell \to \infty$ and any sequence of least-cost industries σ_ℓ that produce Q_ℓ,

$$\int_{\mathcal{F}^+} \left(\frac{q_f}{Q_\ell}\right)^2 \mathrm{d}\sigma_\ell(f) \longrightarrow 1.$$

Proof of Lemma 3. We compare costs with one large machine firm, remove the output of small activities, and bound the remaining concentration deficit.

Step 1: a sharp cost comparison. Set $\gamma := \gamma^M$ and $\alpha := 1/\gamma$. Write $A_* := A(\delta_{\boldsymbol{y}^M})$ for the machine tail coefficient and $c_* := rA_*^{-\alpha}$ for its asymptotic cost coefficient. For every sufficiently large Q, one firm on this ray produces Q at resource cost $(c_* + o(1))Q^\alpha$.

The uniform asymptotic tail and the Lipschitz condition on γ give the following sharp lower bound. For every $\varepsilon > 0$, all sufficiently large feasible positive-scale firms satisfy

$$C(f) \geq (c_* - \varepsilon) q_f^\alpha.$$

To see this, consider any sequence with $q_f \to \infty$. Its scale also diverges, and its normalized human input converges to zero. Along a convergent subsequence, the asymptotic tail gives

$$\frac{C(f)}{q_f^\alpha} = \frac{C(f)}{m_f} A(\nu_f/m_f)^{-\alpha} m_f^{1-\alpha\gamma(\nu_f/m_f)} [1+o(1)]^{-\alpha}.$$

The first factor converges to r, and the uniform error factor converges to one. Put $M_H := \mu(\mathcal{Y}^H)$. By Assumption 1(v), there is a constant $K_\gamma < \infty$ such that

$$\gamma(\nu_f/m_f) \le \gamma + K_\gamma M_H/m_f$$

for every sufficiently large feasible firm. Hence

$$m_f^{1-\alpha\gamma(\nu_f/m_f)} \ge \exp\left\{-\alpha K_\gamma M_H \frac{\log m_f}{m_f}\right\} \longrightarrow 1.$$

The firm composition converges to $\delta_{\boldsymbol{y}^M}$, so continuity gives $A(\nu_f/m_f) \to A_*$. This proves the sharp lower bound.

Step 2: small activities supply a vanishing share. Fix $0 < \varepsilon < c_*$. All atomistic firms and all positive-scale firms below the preceding large-output threshold have resource average cost above a common $\delta > 0$. The one-firm upper bound therefore implies that these activities produce only $O(Q_\ell^\alpha) = o(Q_\ell)$ units in σ_ℓ. Let $\mathcal{L}_\ell \subseteq \mathcal{F}^+$ contain the remaining large-output firms and put $s_{f,\ell} := q_f/Q_\ell$. Then

$$\int_{\mathcal{L}_\ell} s_{f,\ell}\, \mathrm{d}\sigma_\ell(f) \longrightarrow 1.$$

The sharp lower bound and the one-firm upper bound give

$$(c_* - \varepsilon) Q_\ell^\alpha \int_{\mathcal{L}_\ell} s_{f,\ell}^\alpha\, \mathrm{d}\sigma_\ell(f) \le \mathcal{C}(Q_\ell) \le (c_* + o(1)) Q_\ell^\alpha.$$

Step 3: concentration. The set $\mathcal{L}_\ell$ depends on ε, which we keep fixed. Since $s^\alpha - s \ge (1-\alpha)s(1-s)$ on $[0,1]$, Step 2 gives

$$\begin{aligned}\limsup_{\ell\to\infty} \left[1 - \int_{\mathcal{F}^+} s_{f,\ell}^2\, \mathrm{d}\sigma_\ell(f)\right] &\le \frac{1}{1-\alpha}\left(\frac{c_*}{c_*-\varepsilon} - 1\right) \\ &= \frac{\varepsilon}{(1-\alpha)(c_*-\varepsilon)}.\end{aligned}$$

The omitted output share vanishes for each fixed ε. The left side does not depend on ε, so letting $\varepsilon \downarrow 0$ makes it at most zero. Market fit gives $s_{f,\ell} \leq 1$ and bounds the concentration integral above by one. This proves convergence to one. □

Proof of Proposition 3. The output bound in Section A.2.3 gives $Q_N^* \to \infty$. If $\gamma^M < 1$, choose supporting prices for any equilibrium. Active equality and Lemma 1 give

$$\frac{C(f)}{q_f} \leq p_N \leq \bar{c} \qquad \text{for almost every active firm.}$$

Lemma 2 therefore bounds every active positive-scale capacity independently of N. Dividing by Q_N^* proves replication. If $\gamma^M > 1$, equilibrium industries minimize cost at Q_N^*, so Lemma 3 proves concentration.

□

A.2.4 Human diversity and firm composition

Proof of Proposition 4. We first use firm support and active equality to identify the maximizing shares. We then prove the dominance and human-employment claims and classify firms when there is one human type. In these comparisons, we use $\bar{a}(\lambda) = \bar{a}(\lambda, r\mathbf{1})$.

Step 1: active firms maximize cost-adjusted average product. Choose a supporting output price p. For a feasible firm f, write its input bill as $B_w(f) := r\nu_f^M(\mathcal{Y}) + \int_{\mathcal{Y}^H} w \, d\nu_f^H$. Its output per unit cost, multiplied by r, is $rq_f / B_w(f)$, including the normalized inputs and output of an atomistic activity. Active equality gives $rq_f / B_w(f) = r/p$ for almost every active firm; positive costs and positive equilibrium output imply $p > 0$. Since the unrestricted envelope includes every active firm,

$$\max_{\lambda} \bar{a}(\lambda, \boldsymbol{w}) \geq \frac{r}{p}.$$

By hypothesis, a feasible firm attains the global maximum. Firm support applied to that firm gives the reverse inequality. Hence

$$\max_{\lambda \in [0,1]} \bar{a}(\lambda, \boldsymbol{w}) = \frac{r}{p}.$$

Almost every active firm attains this value. Its human share maximizes the envelope, and its size and human mix attain the envelope at that share.

Step 2: dominance and separation at any wage. For AI dominance, take any active positive-size firm that employs humans. Replacing all its humans with machines at the same total size strictly raises output under the stated monotonicity conditions and costs no more, since $w_i \geq \underline{w} \geq r$. By continuity and $F(0) = 0$, a strictly smaller machine-only firm produces exactly the original output. Its cost is strictly lower, so it is a profitable feasible entrant at the supporting price, a contradiction. For an atomistic activity, the same strict increase in L and weak cost reduction permit a lower-cost machine-only activity per unit output. Thus no human-employing firm can be active.

For separation with one human type, the human wage w gives

$$\overline{a}(\lambda, w) = \frac{r\overline{a}(\lambda, r\mathbf{1})}{(1-\lambda)r + \lambda w}.$$

For $0 < \lambda < 1$, strict convexity bounds the numerator strictly below $r[(1-\lambda)\overline{a}(0, r\mathbf{1}) + \lambda\overline{a}(1, r\mathbf{1})]$. Dividing by the denominator gives a weighted average of the two cost-adjusted endpoint values, with weights proportional to $(1-\lambda)r$ and λw. Hence no interior share maximizes $\overline{a}(\cdot, w)$. Step 1 then implies that almost every active firm is pure; no separate feasibility assumption on the endpoints is needed.

Step 3: positive human employment. Suppose instead that almost every active firm is machine-only. Every human type then has slack supply, so human scarcity gives $\boldsymbol{w} = \underline{w}\mathbf{1} \leq r\mathbf{1}$. Lower human wages weakly raise $\overline{a}(\lambda, \boldsymbol{w})$ at every share and leave the machine-only endpoint unchanged. The strict inequality in part (ii) therefore means that zero is not a maximizing human share at these wages. This contradicts Step 1. Thus a set of active firms of positive industry measure employs humans.

Step 4: classify firms with one human type. For part (ii)(a), Step 2 excludes mixed firms, and Step 3 ensures that some active firms are human-only.

For part (ii)(b), the cost-adjusted envelope is the ratio in Step 2. Strict concavity of its numerator and positivity of its affine denominator imply that it has at most one maximizer: if two distinct shares attained the same maximum, their strict convex combinations would have a strictly higher ratio. Attainment follows from the proposition's opening assumption. Step 3 excludes a maximum at zero. Since $w \geq \underline{w} = r$, the left derivative at one satisfies

$$\partial_\lambda^- \overline{a}(1, w) = \frac{r}{w^2}\left[w\overline{a}'_-(1) - (w-r)\overline{a}(1)\right] < 0.$$

Thus one is not a maximizer either. The unique maximizing share is interior, and Step 1 implies that almost every active firm has that share.

□

A.3 Wages and employment.

A.3.1 The wage envelope and sorting

Under Assumption 2, write production as $\widetilde{F}(m, \boldsymbol{z})$, where $\boldsymbol{z}$ is the vector of the firm's mean capabilities. Its marginal product is

$$D_{\boldsymbol{y}} F = \partial_m \widetilde{F} + \frac{1}{m} \nabla_{\boldsymbol{z}} \widetilde{F} \cdot (\boldsymbol{y} - \boldsymbol{z}), \qquad \boldsymbol{z} = \frac{1}{m} \int_{\mathcal{Y}} \boldsymbol{y} \, \mathrm{d}\nu(\boldsymbol{y}). \tag{MEAN-MP}$$

Proof of Theorem 2(i). Put $Q = Q(\sigma)$. By Theorem 1, choose p such that $(p, w) \in \Lambda(Q, \sigma)$. For any firm g, write its input bill as

$$B_w(g) := r\nu_g^M(\mathcal{Y}) + \int_{\mathcal{Y}^H} w \, \mathrm{d}\nu_g^H.$$

Active equality gives $B_w(f) = pq_f$. Firm support implies that f minimizes its input bill among feasible firms that produce q_f. We first show that machines have positive marginal product, then substitute between machines and human type i.

Step 1: positive machine marginal product. Let $m_f^M := \nu_f^M(\mathcal{Y})$ be the firm's machine mass and $M_f := D_{\boldsymbol{y}^M} F(\nu_f)$ its marginal product. We claim that $M_f > 0$. For $t > 0$, let f_t reduce the firm's machine input by the fraction t. Then

$$B_w(f_t) = B_w(f) - trm_f^M, \qquad q_{f_t} = q_f - tm_f^M M_f + o(t).$$

Suppose first that $M_f < 0$. For all small t, the firm f_t costs less and produces more than f. By continuity and $F(0) = 0$, we can scale the firm f_t down until it produces exactly q_f. The scaled firm remains feasible and has a still lower input bill, which contradicts firm support.

Now suppose that $M_f = 0$. Whenever $q_{f_t} > q_f$, the same scaling argument gives a contradiction. Whenever $q_{f_t} \leq q_f$, firm support gives

$$0 \leq B_w(f_t) - pq_{f_t} = -trm_f^M + p(q_f - q_{f_t}) = -trm_f^M + o(t) < 0,$$

which is also a contradiction. Hence $M_f > 0$.

Step 2: fixed-output substitution. Suppose firm f's staffing constraint for human type i is slack. Add a small amount of this type and adjust the firm's machine input to hold output fixed. The implicit function theorem gives the slope of machine use with respect to human use as $-D_{\boldsymbol{y}^i} F(\nu_f) / M_f$. Cost minimization makes the one-sided

derivative of the input bill weakly positive:

$$w_i - r\frac{D_{\boldsymbol{y}^i}F(\nu_f)}{M_f} \geq 0.$$

This is $b_f(\boldsymbol{y}^i) \leq w_i$. If the firm also employs a positive amount of type i, the variation is feasible in both directions, so the inequality is an equality. Substituting $M_f = D_{\boldsymbol{y}^M}F(\nu_f)$ gives $b_f(\boldsymbol{y}^i) = w_i$.

Step 3: the highest eligible bid. The inequality just proved and $w_i \geq \underline{w}$ show that its right-hand side is at most w_i. If type i is slack in the aggregate, human scarcity gives $w_i = \underline{w}$, so the bound is attained. If an active machine-using firm has an interior match with the type, the equality just proved says that its bid equals w_i. The bound is again attained. □

Proof of Theorem 2(ii). We compare bids, then describe their winning regions. Subtracting (MEAN-MP) at $\boldsymbol{y}$ and $\boldsymbol{y}^M$ gives the affine bid formula (BID) and its stated coefficients. By Theorem 2, $b_g(\boldsymbol{y}^i) \leq w_i = b_f(\boldsymbol{y}^i)$ for σ-almost every eligible firm g. Subtracting the affine bids gives (FAN). For a fixed finite set of eligible firms, each winning region is an intersection of half-spaces whose boundaries pass through $\boldsymbol{y}^M$. Their inequalities are homogeneous in $\boldsymbol{y} - \boldsymbol{y}^M$, so the regions are polyhedral cones. Their faces form a polyhedral fan. The interior-match condition therefore assigns workers according to these regions when they share this eligible set; a binding staffing cap can exclude a firm for a given type. □

A.3.2 Human employment and exit

Proof of Proposition 5. For part (i), we derive the industry frontier, prove continuity, and check the exit condition. For part (ii), we exclude mixed firms and compare the two pure technologies.

Part (i), Step 1: the industry frontier. Let T_H, T_M be aggregate human and machine input, and let $T = (T_H, T_M)$. Concavity of best average product bounds industry output by

$$\Phi(T) := (T_H + T_M)\,\overline{a}\left(\frac{T_H}{T_H + T_M}\right), \qquad \Phi(0) := 0.$$

Indeed, each firm's average product is at most $\overline{a}$ at its share, and input-weighted shares average to $T_H/(T_H + T_M)$. Strict concavity and unique preferred scales make equality hold exactly when all input uses this common share at its preferred scale. The relaxed

input problem is therefore

$$\max_{0\le T_H\le \mu,\ T_M\ge 0} \{U_N(\Phi(T)) - \underline{w}T_H - rT_M\}.$$

The feasibility hypothesis makes its maximizing firms feasible in the original economy, so equality is attained.

The frontier Φ is concave and homogeneous of degree one. Two distinct optimal inputs on the same ray give different outputs, so strict concavity of U_N rules them out. Inputs on different rays give a strict gain in Φ at their midpoint, by strict concavity of $\bar{a}$, and hence a strict surplus gain. Optimal output lies below demand saturation: at or above saturation, a small radial input reduction saves cost at first order and loses benefit of smaller order. Thus both comparisons are strict. The relaxed optimum is unique. It exists because benefit is bounded and input costs are positive, and machine viability makes its total input positive. Equality in the frontier bound then determines the common share, preferred scale, and activity weight of the equilibrium industry.

Part (i), Step 2: continuity. Preferred scales share a finite bound. Berge's theorem therefore makes $\bar{a}(\lambda;\theta)$ jointly continuous and the unique preferred scale continuous. The relaxed optimum also varies continuously: locally, positive costs and bounded benefit restrict all maximizing inputs to one compact set. Output, human input, the common share, and preferred scale therefore vary continuously. At positive scale m, firm weight is $(T_H + T_M)/m$, so the industry measure is continuous there. At zero scale we use the atomistic intensity $T_H + T_M$; the disjoint-union topology does not give convergence of firm measures across this boundary.

Part (i), Step 3: exit. At the machine-only relaxed optimum, the output price is $r/\bar{a}(0)$ and output is positive by viability. The right marginal product of humans is $\bar{a}(0) + \bar{a}'_+(0)$. Concavity makes this optimum global exactly when

$$\frac{r}{\bar{a}(0)}[\bar{a}(0) + \bar{a}'_+(0)] \le \underline{w},$$

which is the condition in Proposition 5(i). If the inequality fails, a small human input raises surplus. Uniqueness includes the equality case.

Part (ii): separate production and threshold exit. At supporting prices, feasible preferred pure firms give $p\bar{a}(1) \le w$ and $p\bar{a}(0) \le r$. Strict convexity makes every interior human share satisfy

$$p\bar{a}(\lambda) - \lambda w - (1-\lambda)r < \lambda(p\bar{a}(1) - w) + (1-\lambda)(p\bar{a}(0) - r) \le 0.$$

Thus no mixed firm operates. An active machine firm uses its preferred scale by the size-choice argument in Section A.2.1, so $p = r/\bar{a}(0)$. If $p\bar{a}(1) > \underline{w}$, human support gives $w > \underline{w}$ and scarcity makes all humans work; active equality then gives $w = p\bar{a}(1)$. If $p\bar{a}(1) < \underline{w}$, every human firm loses money, so employment is zero and $w = \underline{w}$. At equality, the wage is $\underline{w}$. Both pure activities break even, and every feasible split of the same equilibrium output between them has the same cost and remains an equilibrium. This gives the wage formula and allows nonunique employment at the threshold. □

A.3.3 Wage changes through reorganization

Proof of Proposition 6. At each parameter value, Theorem 2 gives $w_i = b_f(\boldsymbol{y}^i)$ at the selected employer with an interior match. Equation (MEAN-MP) shows that the listed statistics determine the bid. Apply the chain rule to this continuously differentiable function. Holding m, s, and the human mean fixed holds the relevant quantities of human and machine inputs fixed; only machine capabilities change in the technology term. When human means have fewer independent coordinates, take the derivative along their feasible affine span. The assumptions concern the selected match; they do not establish that such a selection exists in every economy. □

A.3.4 The large-market wage jump

Verification of Figure 26. We verify the black curve for every equilibrium wage selection in the large-market limit. Use the calibration in footnote 24, with $\theta = y_*^M$ denoting AI coordination. Below one, pure firms attain a bound on industry output. Above one, no human firm remains active in a large enough market. At one, supporting-price bounds determine the limiting wage.

Use the optimized log scale factor ℓ from (CES-PEAK) at $\phi = .02$. Then $\bar{a}(1) = 1.1e^{\ell(.985)}$, and the best machine average product below the threshold is $\bar{a}(0) = e^{\ell(\theta)}$. Let h and a be human and machine inputs in a firm, and let $\lambda = h/(h+a)$ denote its human share.

Step 1: replication below one. For $.995 \le \theta < 1$, the highest log average product at human share λ is

$$g_\theta(\lambda) = \log(1 + .1\lambda) + \ell\big(\theta - \lambda(\theta - .985)\big).$$

We show that average product, e^{g_θ}, is strictly convex. Write $d := 1 - \theta + \lambda(\theta - .985)$ for one minus mean coordination. Then

$$\frac{(e^{g_\theta})''}{e^{g_\theta}} = (\theta - .985)^2 \left[\frac{1}{d} + \left(\log \frac{.02}{d}\right)^2\right] - \frac{.2(\theta - .985)}{1 + .1\lambda} \log \frac{.02}{d} > 0.$$

For the sign, multiply by $d(1 + .1\lambda)/(\theta - .985) > 0$. The positive term is at least .01, whereas the subtracted term is at most $.004/e$, since $d \log(.02/d) \leq .02/e$. Convexity therefore gives

$$q_f \leq (h + a)\left[\lambda \bar{a}(1) + (1 - \lambda)\bar{a}(0)\right] = \bar{a}(1)h + \bar{a}(0)a.$$

The same bound holds for atomistic firms by taking the zero-scale limit. Since $\bar{a}(1) > \bar{a}(0)$, every industry with output Q costs at least

$$\frac{Q}{\bar{a}(0)} - \left(\frac{\bar{a}(1)}{\bar{a}(0)} - 1\right)\frac{2}{3}.$$

For all sufficiently large Q, two pure-human firms of scale $1/3$ and copies of the preferred machine firm attain this bound. Both complete firm types fit inside Q, so the construction respects the capacity cap. Active equality pins $p = 1/\bar{a}(0)$ and $w_H = \bar{a}(1)/\bar{a}(0)$. Every optimal dual supports this cost minimizer; hence these prices apply to every supported cost-minimizing industry at that output. The bound in Section A.2.3 gives $Q_N \to \infty$, proving the limiting wage below one.

Step 2: exclude large human firms at and above one. Fix $\theta \geq 1$. Replacing a share $\lambda > 0$ of machines with humans at fixed total size m changes log output by

$$\log(1 + .1\lambda) - \lambda(\theta - .985)\log(1 + m) \leq \lambda\left[.1 - (\theta - .985)\log(1 + m)\right].$$

Choose a finite size M for which the bracket is negative whenever $m > M$. Any human-using firm above that size has lower output than the pure-machine firm of the same size. Because $w_H \geq 1 = r$, the latter has no higher input bill. Scale the firm down to the original output to obtain a strictly cheaper feasible firm. Firm support therefore excludes every human-using positive-scale firm with $m > M$.

If $\theta > 1$, the machine firm with capacity Q has average product tending to infinity. Dual feasibility gives $p_Q \to 0$. On the compact set of human-using compositions and scales $0 \leq m \leq M$, average product has a finite upper bound. Each such activity costs at least one per unit input, so it cannot break even once p_Q is sufficiently small. No humans then work, and scarcity complementarity gives $w_H = 1$. Again Section A.2.3 gives $Q_N \to \infty$, proving the limiting wage above one.

Step 3: the critical case. At $\theta = 1$, define $g_1(0) = .02$ by continuity. The average-product curvature bound in Step 1 still holds for $\lambda > 0$, and continuity extends its chord bound to $\lambda = 0$. Hence

$$q_f \le \bar{a}(1)h + e^{.02}a, \qquad \mathcal{C}(Q) \ge e^{-.02}Q - (\bar{a}(1)e^{-.02} - 1)\frac{2}{3}.$$

Replicating machine firms with average product approaching $e^{.02}$ gives $\mathcal{C}(Q)/Q \to e^{-.02}$. Equilibrium output diverges by Section A.2.3. Dual equality gives $p_Q \ge \mathcal{C}(Q)/Q$. Entry by a machine firm of capacity Q bounds p_Q above by its average cost. These two bounds give $p_Q \to e^{-.02}$.

The pure-human peak firm is an available entrant for large Q, so $w_H \ge p_Q\bar{a}(1) > 1$. Human supply must bind. We show that any human-employing active activity also gives the reverse limiting inequality. Step 2 bounds its positive scale by M; include atomistic firms as scale zero. On this compact set, pure-machine average product is strictly below $e^{.02}$. Continuity therefore gives a fixed $\delta > 0$ such that activities with human share $0 < \lambda < \delta$ cannot break even for large Q, even at wage one. Every active human activity consequently has $\lambda \ge \delta$.

Let v denote such an activity's output per unit input. Active equality and the output bound give

$$w_H = \frac{p_Q v - (1-\lambda)}{\lambda} \le p_Q\bar{a}(1) + (p_Q e^{.02} - 1)\frac{1-\lambda}{\lambda}.$$

The last term tends uniformly to zero because $\lambda \ge \delta$. Combined with $w_H \ge p_Q\bar{a}(1)$, this proves $w_H \to \bar{a}(1)e^{-.02}$ for every supporting wage selection. These three steps establish the stated pointwise limit. They do not impose a pure-firm formula at arbitrary finite output, where mixed firms can be cheaper. □

A.4 Connection to the task model.

Lemma 4 (The task technology satisfies Assumption 1)**.** F^{T} satisfies Assumption 1 with tail exponent $\gamma \equiv 1$. Every nonzero firm ν has a unique optimal vector of task outputs $\boldsymbol{z}^*(\nu)$, and it is strictly positive. With *task prices* $\boldsymbol{\pi}(\nu) := \nabla L(\boldsymbol{z}^*(\nu))$, the marginal product of capability $\boldsymbol{y}$ is

$$D_{\boldsymbol{y}}F^{\mathrm{T}}(\nu) = \max_{1\le d\le n} \pi_d(\nu)\, y_d. \qquad \text{(BEST-TASK)}$$

Proof of Lemma 4. We characterize the optimum, derive (BEST-TASK), and then check Assumption 1. Let $\mathcal{Z}(\nu)$ be the set of task-output vectors that assignments of ν attain. It is convex and compact, and $\mathcal{Z}(m\nu) = m\mathcal{Z}(\nu)$. Each agent can work on the task

where $\pi_d y_d$ is largest, so the support function of $\mathcal{Z}(\nu)$ is

$$\max_{\boldsymbol{z}\in\mathcal{Z}(\nu)} \boldsymbol{\pi}\cdot\boldsymbol{z} = \int_{\mathcal{Y}} \max_d \pi_d y_d \,\mathrm{d}\nu(\boldsymbol{y}) \qquad (\boldsymbol{\pi}\in\mathbb{R}^n). \tag{TASK-SET}$$

Step 1: the optimum. Every capability is at least $\underline{y} > 0$, so the equal assignment gives positive output. An optimum leaves no task empty. When $\rho \le 0$, output is zero without that task. When $0 < \rho < 1$, its marginal product is infinite, so moving a little time to it raises output. If two optima differed and were not proportional, their midpoint would do strictly better, because L is strictly quasiconcave. If they were proportional, homogeneity would give them different values. Hence the optimum is unique.

Step 2: marginal products. Fix ν, put $\boldsymbol{\pi} := \boldsymbol{\pi}(\nu)$, and let $g(\boldsymbol{y}) := \max_d \pi_d y_d$. Concavity and homogeneity of degree one give $L(\boldsymbol{z}) \le \boldsymbol{\pi}\cdot\boldsymbol{z}$ for every $\boldsymbol{z}$, with equality at $\boldsymbol{z}^*(\nu)$. With (TASK-SET), every firm ν' therefore satisfies

$$F^{\mathrm{T}}(\nu') \le \int_{\mathcal{Y}} g \,\mathrm{d}\nu'. \tag{UPPER}$$

Equality holds at $\nu' = \nu$, because the first-order condition for maximizing L over $\mathcal{Z}(\nu)$ makes $\boldsymbol{z}^*(\nu)$ maximize $\boldsymbol{\pi}\cdot\boldsymbol{z}$ there. Write an optimal assignment of ν as $\nu_d = \tau_d\nu$ with time shares $\tau_d(\boldsymbol{y})$. The equality case forces, for ν-almost every $\boldsymbol{y}$, $\tau_d(\boldsymbol{y}) > 0$ only on tasks that maximize $\pi_d y_d$.

Now let η be a finite signed measure with $\nu + t_0\eta \ge 0$ for some $t_0 > 0$. The negative part of η is then absolutely continuous with respect to ν. Assign $\nu + t\eta$ as follows: give the part of η that is absolutely continuous with respect to ν the shares τ, and put the rest, which is nonnegative, on a task that maximizes $\pi_d y_d$. For $0 \le t \le t_0$ this is an assignment. Its task outputs are $\boldsymbol{z}^*(\nu) + t\boldsymbol{v}$ with $\boldsymbol{\pi}\cdot\boldsymbol{v} = \int_{\mathcal{Y}} g \,\mathrm{d}\eta$. Differentiability of L at $\boldsymbol{z}^*(\nu)$ gives $F^{\mathrm{T}}(\nu + t\eta) \ge F^{\mathrm{T}}(\nu) + t\int_{\mathcal{Y}} g \,\mathrm{d}\eta + o(t)$. With (UPPER), this proves the first-variation formula (BEST-TASK). Replacing η by $-\eta$ covers $t < 0$.

Step 3: Assumption 1. Homogeneity gives $F^{\mathrm{T}}(m\nu) = mF^{\mathrm{T}}(\nu)$. Average product is therefore constant in scale, which gives (iv), and (v) holds exactly with $A = F^{\mathrm{T}}$ on $\Delta(\mathcal{Y})$ and $\gamma \equiv 1$. The right side of (TASK-SET) is Lipschitz in $\boldsymbol{\pi}$, with constant $\overline{y}\,\nu(\mathcal{Y})$, and continuous in ν under weak convergence. On bounded-mass sets, weak convergence of firm measures therefore gives uniform convergence of support functions on the unit sphere, that is, Hausdorff convergence of $\mathcal{Z}(\nu)$. Berge's maximum theorem then makes F^{T} and the unique optimum $\boldsymbol{z}^*(\nu)$ continuous. This gives (i), with the positive extension $F^{\mathrm{T}}(m\nu)/m = F^{\mathrm{T}}(\nu)$. Because $\boldsymbol{z}^*(\nu)$ is strictly positive, $\boldsymbol{\pi}(\nu)$ is also continuous, which gives the continuity that (ii) requires. For (iii), let $\nu' \succeq_{\mathrm{FOSD}} \nu$ have equal mass, and build g from $\boldsymbol{\pi}(\nu')$. This g is continuous, nonnegative, and increasing.

Inequality (UPPER), with ν' as the base firm, gives $F^{\mathrm{T}}(\nu') = \int_{\mathcal{Y}} g \, \mathrm{d}\nu' \geq \int_{\mathcal{Y}} g \, \mathrm{d}\nu \geq F^{\mathrm{T}}(\nu)$. □

Lemma 5 (Pooling)**.** Let σ be an industry with aggregate input measure $\bar{\nu} := \int_{\mathcal{F}} \nu_f \, \mathrm{d}\sigma(f)$. Then

$$F^{\mathrm{T}}(\bar{\nu}) \;\geq\; \int_{\mathcal{F}} F^{\mathrm{T}}(\nu_f) \, \mathrm{d}\sigma(f) = Q(\sigma),$$

and the two sides carry the same resource cost. Consequently the least-cost function $\mathcal{C}$ of the task economy in our framework is the least-cost function of the task economy of Definition 3; it is convex and strictly increasing; and every least-cost industry is, in aggregate, an optimal task allocation producing its output exactly.

Proof of Lemma 5. Build g from $\boldsymbol{\pi}(\bar{\nu})$. By the equality case of (UPPER) at $\bar{\nu}$, and then by (UPPER) at each firm,

$$F^{\mathrm{T}}(\bar{\nu}) = \int_{\mathcal{Y}} g \, \mathrm{d}\bar{\nu} = \int_{\mathcal{F}} \int_{\mathcal{Y}} g \, \mathrm{d}\nu_f \, \mathrm{d}\sigma(f) \geq \int_{\mathcal{F}} F^{\mathrm{T}}(\nu_f) \, \mathrm{d}\sigma(f).$$

Resource cost is linear in the input measure, so $\int_{\mathcal{F}} C(f) \, \mathrm{d}\sigma(f) = \underline{w} \, \bar{\nu}^H(\mathcal{Y}) + r \, \bar{\nu}^M(\mathcal{Y})$, which is the bill the aggregate pays.

For the cost functions, take an industry producing Q at cost c with human use $\bar{\nu}^H \leq \mu$. An optimal assignment of $\bar{\nu}$ is an allocation of the task economy with the same bill and output $F^{\mathrm{T}}(\bar{\nu}) \geq Q$, so the task economy produces Q at cost at most c. Conversely, an allocation with task outputs $\boldsymbol{z}$ and bill c uses aggregate input measure $\bar{\nu}$; the single atomistic firm with composition $\bar{\nu}/\bar{\nu}(\mathcal{Y})$, run at intensity $\bar{\nu}(\mathcal{Y})$, costs c and produces $F^{\mathrm{T}}(\bar{\nu}) \geq L(\boldsymbol{z})$. The two least-cost functions therefore agree.

Read on the task economy's allocations, least cost is the value of a convex program: the constraint set is convex, L is concave, and the bill is linear. So $\mathcal{C}$ is convex, and $\mathcal{C}(0) = 0$ with positive input prices makes it strictly increasing. Finally, let σ be least cost at Q. Pooling produces $F^{\mathrm{T}}(\bar{\nu}) \geq Q$ at cost $\mathcal{C}(Q)$, so $\mathcal{C}(F^{\mathrm{T}}(\bar{\nu})) \leq \mathcal{C}(Q)$, and strict monotonicity forces $F^{\mathrm{T}}(\bar{\nu}) = Q$. □

Proof of Theorem 4. Both economies have the same least-cost function by Lemma 5. We identify output, characterize supporting prices, recover assignments and wages, and bound the exceptional outputs.

Step 1: output. Under constant returns a positive-scale firm has the same output and inputs per unit of firm size as the atomistic firm with its composition, so the capacity

cap in Σ never binds. Convexity of $\mathcal{C}$ and $\mathcal{C}(0) = 0$ make every slope at $Q > 0$ at least $\mathcal{C}(Q)/Q > 0$, while the standing condition gives $\mathcal{C}(Q) \leq \bar{c}Q$ with $\bar{c} < P(0)$. Hence Q_N^* lies strictly between 0 and $D_N(0)$. There U_N is strictly concave, so Q_N^* is unique and $P_N(Q_N^*) \in \partial\mathcal{C}(Q_N^*)$. Conversely any Q with $P_N(Q) \in \partial\mathcal{C}(Q)$ maximizes the concave function $U_N - \mathcal{C}$, so $Q = Q_N^*$.

Step 2: supporting prices and subgradients. We show that a least-cost industry at Q admits supporting price p exactly when $p \in \partial\mathcal{C}(Q)$. Under constant returns, atomistic recipes implement every input mix, so the firm menu is independent of Q.

For necessity, integrate firm support over any alternative industry σ' and subtract active equality at the supported industry σ. Human feasibility and scarcity give

$$\int_{\mathcal{F}} C(f)\ \mathrm{d}\sigma'(f) \geq \int_{\mathcal{F}} C(f)\ \mathrm{d}\sigma(f) + p[Q(\sigma') - Q].$$

Taking the infimum at each exact output Q' proves both least cost at Q and $\mathcal{C}(Q') \geq \mathcal{C}(Q) + p(Q' - Q)$.

Conversely, let $p \in \partial\mathcal{C}(Q)$ and let σ minimize cost at Q. The subgradient inequality makes σ minimize cost minus p times output, subject only to human supplies. Apply the conic duality argument of Lemma 1, now with output priced at p. The same strictly feasible human allocation gives multipliers $w_i - \underline{w} \geq 0$. Dual feasibility and complementary slackness are precisely firm support, active equality, and human scarcity. Atomistic machine entry gives $p \leq \bar{c}$, so recipes discarded in the compact reduction cannot violate support. Hence $(p, w) \in \Lambda(Q, \sigma)$.

Step 3: output and assignments coincide. By Theorem 1 the industries of our equilibrium outcomes are those that produce Q_N^* at least cost. By Step 2 a competitive allocation is least cost at its output Q, with $P_N(Q) \in \partial\mathcal{C}(Q)$, so Step 1 gives $Q = Q_N^*$. Conversely Step 2 with $p = P_N(Q_N^*)$, together with Lemma 5, turns every least-cost industry at Q_N^* into a competitive allocation; existence follows because an efficient industry exists. This proves (i). For (ii), fix a least-cost industry with aggregate $\bar{v}$. By Lemma 5 its optimal assignment produces Q_N^*, and the equality case of (UPPER) at $\bar{v}$ puts every factor only on tasks that maximize $\pi_d(\bar{v})y_d$. Competitive task prices are $p\,\nabla L(\boldsymbol{z}^*(\bar{v})) = p\,\boldsymbol{\pi}(\bar{v})$, and scaling all task prices by $p > 0$ leaves the maximizing tasks unchanged.

Step 4: wages. Integrating active equality over σ gives

$$pQ_N^* = \mathcal{C}(Q_N^*) + \sum_{i=1}^{K}(w_i - \underline{w})H_i(\sigma),$$

so each wage vector is supported by exactly one payment p. Our equilibrium wage vectors are those supported by some $p \in \partial\mathcal{C}(Q_N^*)$; the competitive ones are those supported by $P_N(Q_N^*)$. By Step 2 every such payment supports at least one wage vector. The two sets therefore coincide if and only if $\partial\mathcal{C}(Q_N^*) = \{P_N(Q_N^*)\}$, that is, if and only if $\mathcal{C}$ is differentiable at Q_N^*. The wage formula is (BEST-TASK) priced at p, together with $w_i \geq \underline{w}$ and complementary slackness on human supply. This proves (iii).

Step 5: exceptional outputs. Let σ produce Q_N^* at least cost with aggregate $\bar{\nu}$. Suppose $\bar{\nu}$ contains machines. Adding or removing t units of them changes output by $tD + o(t)$ and cost by rt, where $D := D_{\boldsymbol{y}^M} F^{\mathrm{T}}(\bar{\nu}) > 0$ by (BEST-TASK). The right derivative of $\mathcal{C}$ at Q_N^* is therefore at most r/D and the left derivative at least r/D, so convexity makes both equal r/D. The same argument, with a human type employed strictly between 0 and its supply at cost $\underline{w}$, covers the other case. Now suppose $\mathcal{C}$ fails to be differentiable at Q_N^*. Then no least-cost industry uses machines and each employs every human type at either 0 or its full supply, so its aggregate is $\bar{\nu} = \sum_{i \in S} \mu_i \delta_{\boldsymbol{y}^i}$ for the employed set S. By Lemma 5, $Q_N^* = F^{\mathrm{T}}(\bar{\nu})$. There are 2^K sets S, which proves (iv). □

Proof of Proposition 7 and the task-model exit condition. We check the augmented technology, differentiate output, and then evaluate the human entry margin at a machine-only firm.

Step 1: the augmented technology. The scale multiplier is

$$\frac{S(m,z)}{m} = (1+m)^{z-1} \exp\{\phi m/(1+m)\},$$

with limit one at zero. Thus $F(\nu) = F^{\mathrm{T}}(\nu) S(m, \bar{y}_*)/m$. Lemma 4 and continuity of the multiplier give continuity and the positive average-product extension in Assumption 1(i). Differentiability follows from the chain rule below. At equal mass, a first-order stochastic improvement raises both F^{T} and mean coordination; since the multiplier increases in coordination, the product is monotone. At a fixed composition, $\partial_m \log(S/m) = (z-1)/(1+m) + \phi/(1+m)^2$ changes sign at most once, from positive to negative, so average product is quasiconcave. Finally, on unit compositions, $F(m\nu) = e^{\phi} F^{\mathrm{T}}(\nu) m^{\int y_* \, \mathrm{d}\nu} [1 + o(1)]$. Bounded capabilities make the remainder uniform and the exponent Lipschitz under mixing. Hence the augmented technology satisfies Assumption 1.

Step 2: the bid. Adding mass t at capability $\boldsymbol{y}$ raises size from m to $m+t$ and changes mean coordination by $t(y_* - \bar{y}_*)/m + o(t)$. The first variation, which measures the

output change per unit added, is therefore

$$D_{\boldsymbol{y}}F = \frac{S}{m}\left\{D_{\boldsymbol{y}}F^{\mathrm{T}} + \frac{F^{\mathrm{T}}}{m}\left[m\partial_m \log(S/m) + \partial_z \log S\,(y_* - \overline{y}_*)\right]\right\}. \qquad \text{(TASK-VAR)}$$

All quantities on the right are evaluated at the original firm. At a positive preferred size, $\partial_m \log(S/m) = 0$ and $\partial_z \log S = \log(1+m)$. Taking the ratio of (TASK-VAR) for the worker and the machine cancels S/m and gives (TASK-BID). The denominator is positive because the firm values machines. The wage equality follows from Theorem 2(i). With constant returns, both scale derivatives in (TASK-VAR) are zero, so the coordination term disappears.

Step 3: the exit condition. Use the human share λ and unit composition ν_λ from Proposition 5, with mean coordination $z_\lambda = \lambda y_*^H + (1-\lambda)y_*^M$. Constant returns of the task technology give

$$\overline{a}(\lambda) = F^{\mathrm{T}}(\nu_\lambda)\max_{m\geq 0}\frac{S(m, z_\lambda)}{m}.$$

The unique, bounded preferred scales in Proposition 5(i) permit the envelope theorem, including a zero preferred scale. Differentiating from the right at the machine-only mix gives

$$1 + \frac{\overline{a}'_+(0)}{\overline{a}(0)} = \frac{D_{\boldsymbol{y}^H}F^{\mathrm{T}}(\delta_{\boldsymbol{y}^M})}{F^{\mathrm{T}}(\delta_{\boldsymbol{y}^M})} + (y_*^H - y_*^M)\log\left(1 + m^*(y_*^M)\right).$$

Here homogeneity gives $D_{\boldsymbol{y}^M}F^{\mathrm{T}}(\delta_{\boldsymbol{y}^M}) = F^{\mathrm{T}}(\delta_{\boldsymbol{y}^M})$. To evaluate the remaining ratio, let π_d be task prices at the machine-only optimum. Every task receives positive machine time, so $\pi_d y_d^M$ has the same value for every task d. The best-task formula (BEST-TASK) then gives

$$\frac{D_{\boldsymbol{y}^H}F^{\mathrm{T}}(\delta_{\boldsymbol{y}^M})}{F^{\mathrm{T}}(\delta_{\boldsymbol{y}^M})} = \max_d \frac{y_d^H}{y_d^M}.$$

Substitution in the condition in Proposition 5(i) proves the task-model exit condition, including equality. This step uses the strict concavity and feasibility hypotheses of Proposition 5(i); the bid identity does not require them. □

Checks for Figure 28. Every composition in this figure has mean coordination between .8 and .97, so preferred sizes are positive and below $.5/.03 - 1 < 16 < \mu$. Task output per agent is at most 2 and the scale multiplier is at most $e^{.5}$, so every preferred firm's output is below 52. Machine-only average product is at least $2e^{\ell(.8)} > 2$, with ℓ defined in (CES-PEAK), so relaxed equilibrium output exceeds $200(3 - 1/2) = 500$. Every preferred scale is feasible, including at the relaxed optimum required by Proposition 5(i).

On the coordination path, AI is twice as productive on every task. For machine share s, unit-workforce task output is $1+s$ and mean coordination is $z = .95 + (y_*^M - .95)s$. Optimized average product is $\bar{a}(1-s) = (1+s)e^{\ell(z)}$. Where humans coordinate better, $y_*^M < .95$, differentiation gives

$$\frac{\partial_s^2 \bar{a}(1-s)}{(.95 - y_*^M)e^{\ell(z)}} = (.95 - y_*^M)(1+s)\left[\frac{1}{1-z} + \ell'(z)^2\right] - 2\ell'(z).$$

The figure script uses outward-rounded arithmetic on a 100×100 grid of rectangles covering the whole share interval and the coordination path below .95. The upper bound is below $-.049$, verifying strict concavity where we use the task-model exit condition. At and above $y_*^M = .95$, replacing each human with half a machine preserves task output, weakly raises mean coordination, and reduces input cost. Reoptimizing scale therefore makes machine-only production strictly cheaper. The exit statistic $.5 + (.95 - y_*^M)\log[.5/(1 - y_*^M)]$ reaches .6 at AI coordination approximately .879879.

Along the coordination path, we minimize unit cost over machine shares and choose the wage that clears human supply, or the outside wage when supply is slack. Saved checks compare the bid formula with wages in mixed firms, search for profitable entry on 20,001 shares, and verify resource accounting and firm feasibility. Independent cost linear programs, scalar minimizations, and finer composition grids check selected points. The saved checks in `images/task/task_rent_checks.json` record the path, snapshots, and residuals.

Proof of Theorem 3. The result follows from Theorem 4, proved above. Parts (i)–(ii) give unique output and identify competitive task allocations with the aggregates of equilibrium industries under optimal assignments. If some equilibrium uses machines, its industry is least cost at that output. Part (iv) then makes the cost function differentiable there, and part (iii) gives equality of the two sets of equilibrium wage vectors. □

A.5 Continuity of values and wages. Fix a continuous feasible AI path $\boldsymbol{y}^M(\theta)$ on an open interval Θ, holding all other primitives fixed and maintaining machine viability. Identify a firm by its human inputs and machine quantity, with the Euclidean topology on each component of $\mathcal{F}^+ \sqcup \mathcal{F}^0$. We use weak convergence against bounded continuous functions on this disjoint union.

Proposition 9 (Continuous values and closed equilibrium limits)**.** If Assumption 1 holds and the AI path is continuous, then:

(i) Minimum cost $(Q, \theta) \mapsto \mathcal{C}(Q; \theta)$ is continuous on $\mathbb{R}_+ \times \Theta$.

(ii) Maximal surplus $\theta \mapsto \mathcal{W}_N^*(\theta)$ is continuous.

(iii) Equilibrium quantities $\theta \mapsto \mathcal{Q}_N^*(\theta)$ are upper hemicontinuous.

(iv) Equilibrium industries $\theta \mapsto \Sigma^*(\theta)$ have closed graph under weak convergence of finite industry measures.

Proof of Proposition 9. We extend the cost argument in Section A.1.1 to a moving AI profile, then pass to surplus maximizers.

Step 1: joint cost continuity. The atomistic machine cost $\bar{c}(\theta) = r/L(\delta_{\boldsymbol{y}^M(\theta)})$ is continuous. Nearby parameters and bounded output therefore share the scale and total-input bounds used in Lemma 1. Take $(Q_n, \theta_n) \to (Q, \theta)$ with $Q > 0$ and a finite-support minimizer at (Q, θ). Update its machines. Trim any complete firm whose output exceeds Q_n to its first scale producing Q_n. Inputs weakly fall and each adjusted output converges to its original value. Reduce weights if total output exceeds Q_n, or fill a shortfall with atomistic machines at vanishing cost. This gives $\limsup_n \mathcal{C}(Q_n; \theta_n) \leq \mathcal{C}(Q; \theta)$.

For the reverse inequality, choose minimizers with at most $K + 3$ types. Use the same scale, labelled unit composition, and total-input representation as in Section A.1.1, identifying the machine component by its quantity. Pad the lists with zero coefficients. The uniform bounds give a convergent subsequence. Joint continuity of output per input, including at zero scale, preserves output, cost, and all individual and aggregate constraints. The limiting finite list implements an industry at (Q, θ), proving the lower bound. At $Q = 0$, use $0 \leq \mathcal{C}(Q; \theta) \leq \bar{c}(\theta)Q$.

Step 2: values and quantities. Joint continuity gives uniform cost convergence on $[0, D_N(0)]$. Berge's theorem applied to $U_N(Q) - \mathcal{C}(Q; \theta)$ gives continuous maximal surplus and upper hemicontinuous maximizing quantities.

Step 3: limits of industries. Suppose equilibrium measures σ_n converge weakly to a finite measure σ as $\theta_n \to \theta$. Active equality and Lemma 1 give $c_{\min} m_f \leq p_n q_f \leq \bar{c}(\theta_n) D_N(0)$ on each positive-scale support. Thus firm sizes have a common bound. Weak convergence also bounds total activity weights. On this size range, inputs and costs are bounded continuous functions, and outputs converge uniformly. Their integrals therefore converge, using continuous cutoffs outside the common bound. In particular, $Q_n \to Q_\theta(\sigma)$. Human feasibility passes to the limit, as does the market cap:

$$\int_{\mathcal{F}^+} (q_{\theta_n}(f) - Q_n)_+ \, \mathrm{d}\sigma_n(f) = 0.$$

The two firm components are disjoint, so this integral also uses a bounded continuous integrand on the relevant size range. Thus σ is feasible and its surplus is the limit

of maximal surplus. Step 2 makes that limit maximal at θ, and Theorem 1 gives equilibrium. □

Closed equilibrium limits do not imply a continuous selection. Moreover, a firm with vanishing activity can still set a wage. Wage continuity therefore requires control of the supported firms.

Definition 2 (Wage-interior industry)**.** An industry σ is *wage-interior* if its positive-scale support is nonempty and compact, every firm on that support uses machines with $D_{\boldsymbol{y}^M} F(\nu_f) > 0$ and has strictly slack human staffing constraints, and each human type i satisfies

$$\int_{\mathcal{F}} \nu_f^H(\{\boldsymbol{y}^i\})\, \mathrm{d}\sigma(f) < \mu_i \quad \text{or} \quad \int_{\mathcal{F}^+} \nu_f^H(\{\boldsymbol{y}^i\})\, \mathrm{d}\sigma(f) > 0.$$

Compactness makes the strict inequalities uniform and keeps supported sizes away from zero. The last condition ensures that each fully employed type has an employer with an interior match. Theorem 2 therefore identifies wages with bid envelopes.

Lemma 6 (Support closeness controls wages)**.** Consider wage-interior equilibrium industries (σ_n, w^n) and (σ, w) with frontier machines $\boldsymbol{y}_n^M$ and $\boldsymbol{y}^M$. If

$$\boldsymbol{y}_n^M \longrightarrow \boldsymbol{y}^M \qquad \text{and} \qquad d_H(\mathrm{supp}(\sigma_n|_{\mathcal{F}^+}), \mathrm{supp}(\sigma|_{\mathcal{F}^+})) \longrightarrow 0,$$

where d_H denotes Hausdorff distance between sets of firms, then the wage vectors converge:

$$\|w^n - w\|_\infty \longrightarrow 0.$$

Proof of Lemma 6. For a firm under frontier $\boldsymbol{y}_n^M$, write

$$b_f^n(\boldsymbol{y}) := r \frac{D_{\boldsymbol{y}} F(\nu_f)}{D_{\boldsymbol{y}_n^M} F(\nu_f)}.$$

Wage interiority and Theorem 2 give, for every type i,

$$w_i = \max_{f \in \mathrm{supp}(\sigma|_{\mathcal{F}^+})} b_f(\boldsymbol{y}^i) \vee \underline{w}, \qquad w_i^n = \max_{f \in \mathrm{supp}(\sigma_n|_{\mathcal{F}^+})} b_f^n(\boldsymbol{y}^i) \vee \underline{w}.$$

By Assumption 1(ii), the human and machine marginal products vary continuously with the firm and the machine frontier away from zero scale. The interiority margins at the limit industry keep the machine marginal product bounded away from zero for all nearby supports. Hence every type's bid varies continuously across paired

firms. Hausdorff convergence pairs each firm in either support with a nearby firm in the other, so the two highest bids converge. The number of human types is finite, so convergence holds in the sup norm. □

B EQUILIBRIUM CONCEPT AND FIRM COUNTS

B.1 Complete firms and activity weights. A recipe describes a complete firm; its activity weight counts copies. Divisible weights relax integer firm counts without relaxing the limits on each complete firm or aggregate human use.

At a fixed output requirement and capacity cap, divisible weights make the cost problem convex. This relaxation is exact for atomistic firms and approximates indivisible firms when each is small relative to the market. With large firms, the concentration index measures complete-firm capacities. An integer industry can have different cost and need not admit the same exact supporting prices.

B.2 The replacement game. We fix demand scale N. A one-shot challenge tests every possible replacement of a status-quo contract.

Players and actions. Player 0 represents consumers; player $i \in \{1, \dots, K\}$ owns human type i's resources. Two disjoint copies of $[0, 1]$ index organizers and producers. Entrants have outside payoff zero. Machines are supplied competitively at resource price r.

An organizer stays out or offers a contract $\omega = (Q, \sigma, p, w, t)$ with $Q > 0$, $\sigma \in \Sigma$, $Q(\sigma) = Q$, $p \geq 0$, $w_i \geq \underline{w}$, and transfers satisfying

$$t_E + t_0 + \sum_{i=1}^{K} t_i = 0.$$

The transfers pay the organizer, consumers, and human coalitions, respectively. Each protected party accepts or rejects the offer. Conditional on acceptance, a producer stays inactive or chooses an intensity $a > 0$ and a recipe $f \in \mathcal{F}^0 \sqcup \{g \in \mathcal{F}^+ : q_g \leq Q\}$. Human coalition i chooses employment $h_i \in [0, \mu_i]$. The intensity is an unrestricted density over producers. Complete-firm feasibility restricts recipes; aggregate clearing restricts deployed inputs. Producer actions must be measurable and induce a finite industry measure.

Information and timing. All primitives, offers, actions, and payoffs are public. Given an implemented status quo, or no contract, an inactive organizer can make one challenge. Protected parties accept or reject simultaneously. A rejection keeps the status

quo and pays the challenger zero. After unanimous acceptance, producers and human coalitions choose actions at the proposed prices. The offer replaces the status quo only if optimal actions implement it and clear markets; otherwise the status quo remains. A successful offer commits consumers to buy Q at the internal payment p. The displaced organizer and producers receive zero.

Payoffs and clearing. Write $H_i(\sigma)$ for aggregate employment of type i. The protected parties receive

$$\Pi_0(\omega) = U_N(Q) - pQ + t_0, \qquad \Pi_i(\omega) = (w_i - \underline{w})H_i(\sigma) + t_i.$$

The organizer receives t_E; without a contract all these net payoffs are zero. A producer choosing (a, f) earns

$$a\left[pq_f - r\nu_f^M(\mathcal{Y}) - \int_{\mathcal{Y}^H} w \, \mathrm{d}\nu_f^H\right],$$

and inactivity earns zero. A human coalition maximizes $(w_i - \underline{w})h_i$; its transfer does not depend on its labor choice. With inactive producers assigned intensity zero, implementation requires

$$\sigma(B) = \int_0^1 a_e \mathbf{1}\{f_e \in B\} \, \mathrm{d}e \quad \text{for every Borel set } B, \qquad H_i(\sigma) = h_i, \qquad Q(\sigma) = Q.$$

The offer has a *competitive continuation* if optimal producer and human actions satisfy these clearing conditions. As proved in Section A.1, this is equivalent to $(p, w) \in \Lambda(Q, \sigma)$.

Admissibility and solution concept. A contract is *admissible* if it has a competitive continuation, balanced transfers, and $t_E \geq 0$, $\Pi_j(\omega) \geq 0$ for every protected party $j = 0, \ldots, K$. Zero producer profits and balanced transfers give the accounting identity

$$\Pi_0(\omega) + \sum_i \Pi_i(\omega) + t_E = \mathcal{W}_N(\sigma).$$

Protected parties accept exactly when an offer preserves their status-quo payoffs; an organizer enters if it can gain strictly. An admissible offer $\widehat{\omega}$ *blocks* an implemented contract ω if

$$\Pi_j(\widehat{\omega}) \geq \Pi_j(\omega) \quad (j = 0, \ldots, K), \qquad \widehat{t}_E > 0.$$

A contract is *replacement-stable* if no admissible offer blocks it. An equilibrium outcome (σ, w) is supported by such a contract. This is the stable-contract test in the main text, not Nash equilibrium of a simultaneous-offer game. It tests every replacement without choosing an order of proposals.

C CONNECTION TO THE TASK MODEL

In our model an agent brings its whole capability vector to every domain at once. In a task model an agent has one unit of time and divides it among tasks (Acemoglu and Restrepo, 2018; Aghion et al., 2019). We show that the second is a special case of the first. Switch coordination off, impose constant returns, and read each productive domain as a task. Our framework then delivers the task model exactly: the same output, the same assignment of factors to tasks, and, except at finitely many output levels, the same wages (Theorem 4). Under constant returns, the mean-capability technology corresponds, up to output units, to equal time on every task. Optimal task assignment allows agents to specialize.

C.1 The task economy. We write the task model on its own terms first, with no reference to firms or to our capability space. There are n tasks. A final good is assembled from task outputs $z = (z_1, \dots, z_n)$ by the constant-returns aggregator

$$L(z) = \left(\frac{1}{n} \sum_{d=1}^{n} z_d^{\rho} \right)^{1/\rho}, \qquad \rho < 1,$$

read as the geometric mean at $\rho = 0$. Human type i has mass μ_i, reservation wage $\underline{w} > 0$, and task productivity y_d^i, measured as output per unit of time. A machine has task productivity y_d^M and can be rented without limit at $r > 0$ per unit of time. All task productivities lie in $[\underline{y}, \overline{y}]$ with $\underline{y} > 0$.

An *allocation* divides each factor's time across tasks: masses $\ell_{id} \geq 0$ of type-i time and $k_d \geq 0$ of machine time on task d, with $\sum_d \ell_{id} \leq \mu_i$. Task outputs and the resource bill are

$$z_d = \sum_{i=1}^{K} y_d^i \ell_{id} + y_d^M k_d, \qquad \underline{w} \sum_{i,d} \ell_{id} + r \sum_d k_d.$$

Humans and machines are perfect substitutes within a task and differ only in their task productivities. The final good sells at the demand price P_N.

Definition 3 (Competitive equilibrium of the task economy)**.** An allocation together with task prices $\boldsymbol{\pi}$, wages $\boldsymbol{w}$, and an output price p is a *competitive equilibrium* if

(a) output clears at the demand price, $p = P_N(L(z))$, and task prices are marginal revenue products, $\boldsymbol{\pi} = p \nabla L(z)$;

(b) every factor works only where it earns most: $\ell_{id} > 0$ only if task d maximizes $\pi_d y_d^i$, and $k_d > 0$ only if task d maximizes $\pi_d y_d^M$;

(c) machines break even, $r \geq \max_d \pi_d y_d^M$, with equality if any machine works;

(d) each human type earns its best task or its outside option:

$$w_i = \max\{\underline{w}, \max_d \pi_d y_d^i\}.$$

Type i works fully when its best task pays more than $\underline{w}$, and does not work when it pays less.

No firms appear in this description, and none are needed. Constant returns leave the boundary between producers undetermined, which is why the task literature works with an aggregate.

C.2 The task technology. Use F^{T} from Section 7.1: a firm assigns its humans and machines optimally across tasks and pays each agent once. Keep all other primitives fixed. Lemma 4 verifies Assumption 1, so equilibrium existence, efficiency, and cost duality apply. Constant returns give $F^{\mathrm{T}}(mv) = mF^{\mathrm{T}}(v)$ and $\gamma \equiv 1$. The large-market result therefore does not apply. Nor do the affine bids: task output depends on the distribution of skills, so F^{T} need not satisfy Assumption 2. Its marginal product instead equals the best-task value in (BEST-TASK).

C.3 The equivalence. How can an economy of firms match an economy with none? Because pooling cannot lower output and preserves resource cost (Lemma 5). Least-cost industries therefore attain the same output as a single firm that combines their inputs.

Our equilibrium differs from the competitive one in two further ways: its internal payment p need not equal the demand price, and its contracts include balanced transfers. We therefore compare output, assignments, and wage vectors, not final payoffs. Write $\partial\mathcal{C}(Q)$ for the subdifferential of least cost at Q, the slopes of the lines that touch $\mathcal{C}$ at Q and lie below it.

Theorem 4 (The task model is our model with the task technology)**.** Fix a demand scale N and run our economy with F^{T}. Then:

(i) **Same output.** Efficient output Q_N^* is unique. An industry belongs to an equilibrium outcome of our economy if and only if it produces Q_N^* at least cost. A competitive equilibrium of the task economy exists, every competitive equilibrium produces Q_N^*, and the competitive allocations are exactly the aggregates of our equilibrium industries.

(ii) **Same assignment.** In both economies every employed factor works only on tasks that maximize $\pi_d y_d$, at task prices proportional to $\nabla L(z^*)$ evaluated at the common task outputs. How our economy splits that allocation across firms is indeterminate.

(iii) **Same wages.** Every competitive wage vector is an equilibrium wage vector of ours, and the two sets coincide if and only if $\mathcal{C}$ is differentiable at Q_N^*. Wages are then $w_i = \max\{\underline{w}, \max_d \pi_d y_d^i\}$ in both. Otherwise our equilibrium supports every payment in $\partial\mathcal{C}(Q_N^*)$ and the competitive equilibrium selects $P_N(Q_N^*)$ from it.

(iv) **Finitely many exceptions.** $\mathcal{C}$ is differentiable at Q_N^* whenever some least-cost industry there uses machines, or employs some human type in an amount strictly between 0 and its supply. It can therefore fail only at an output of the form $F^{\mathrm{T}}\left(\sum_{i\in S} \mu_i \delta_{\boldsymbol{y}^i}\right)$ for a set S of human types, so at most 2^K outputs in all.

Both economies pick the least-cost way to make the efficient quantity, and they assign factors to tasks in the same way. They can part company only over wages, and only where least cost has a kink. There the competitive equilibrium lets demand choose one payment, $P_N(Q_N^*)$, while ours accepts every supporting payment, because consumers and workers can divide the surplus through transfers. By part (iv) a kink needs every employed human type to be fully employed and no machine to work. With one human type that happens only before automation begins. Once any machine works, our wages are the competitive wages.

Example 2 (A kink before automation)**.** Let $n = 3$ and $\rho = -1$. One human type has mass one and productive capability profile $(1, 1, 1)$. The machine has productive capability profile $(1.2, 1, .8)$. Coordination entries play no role. Let $\underline{w} = .2$, $r = .5$, and $P(Q) = (2 - 3Q)_+$ at $N = 1$. A machine-only firm produces about .327 per machine, so its unit cost is about $1.53 < P(0)$ and the standing condition holds. Humans alone produce $1/3$ per unit, with task prices $(1/3, 1/3, 1/3)$. By (BEST-TASK), one more human adds $1/3$ and the first machine adds .4, on task 1. Least cost is therefore $\mathcal{C}(Q) = .6Q$ up to $Q = 1/3$, and (UPPER) shows that no firm with machines is cheaper. Beyond $1/3$, output needs machines, and the first units cost $r/.4 = 1.25$ per unit of output. Since $.6 < P(1/3) = 1 < 1.25$, efficient output is $Q_1^* = 1/3$, produced by humans alone, and $\partial\mathcal{C}(1/3) = [.6, 1.25]$. The competitive equilibrium sets $p = 1$ and $w = p/3 = 1/3$. Our equilibrium supports every $p \in [.6, 1.25]$, so every wage $w = p/3 \in [.2, 5/12]$. At the lowest wage, humans earn their outside option. At the highest, the first machine just breaks even.

Under constant returns, the mean-capability technology instead gives

$$F(\nu) = L\left(\int_{\mathcal{Y}} y_1 \, \mathrm{d}\nu, \ldots, \int_{\mathcal{Y}} y_n \, \mathrm{d}\nu\right).$$

An equal time assignment produces $F(\nu)/n$, so $F(\nu) \leq nF^{\mathrm{T}}(\nu)$. Optimal task assignment can raise output by letting agents specialize. Its marginal product is the best-task value in (BEST-TASK); the mean-capability technology instead gives the affine bids in (BID).

D CES EXAMPLES AND NUMERICAL METHODS

We give the calculations for the current CES figures and distinguish analytic proofs from numerical checks. Parameters appear beside each calculation or in the cited figure footnote. The large-market wage proof is in Section A.3.4.

D.1 Common CES formulas. At mean coordination z, preferred scale is

$$m^*(z) = \begin{cases} 0, & z \leq 1-\phi, \\ \phi/(1-z) - 1, & 1-\phi < z < 1, \\ +\infty, & z \geq 1. \end{cases} \tag{CES-SCALE}$$

For $1-\phi < z < 1$, define the optimized log scale factor by

$$\ell(z) := \phi - (1-z) + (1-z)\log\frac{1-z}{\phi}. \tag{CES-PEAK}$$

The best average product at a composition is then its level factor times $e^{\ell(z)}$. For $z \leq 1-\phi$, the best scale factor is one, attained at zero scale. At $z = 1$, its supremum is e^{ϕ} and is not attained at a finite scale. For $z > 1$, it is unbounded.

The CES scale formulas. At mean coordination z, the CES average product is proportional to $(1+m)^{z-1}e^{\phi m/(1+m)}$. Its log derivative in scale is

$$\frac{z-1}{1+m} + \frac{\phi}{(1+m)^2}.$$

The derivative changes sign at most once, from positive to negative. Its sign gives (CES-SCALE), including the zero-scale corner and the unattained optima. Substitution gives (CES-PEAK). □

A common verification method. At equal unit input costs $r = \underline{w} = 1$, a bound $q_f \leq Am_f$ for every complete firm, with the corresponding bound for atomistic activities, implies $\mathcal{W}_N(\sigma) \leq U_N(Q) - Q/A$. An industry that attains this bound at its maximizing quantity is efficient. If its firms can be staffed and fit in the market, the prices $p = 1/A$ and $w = 1$ support it, so Theorem 1 gives equilibrium. If one share and scale uniquely attain the bound, every equilibrium uses them. Each calculation below checks the required feasibility conditions.

For a separate-production bound $q_f \leq \overline{a}(1)h + \overline{a}(0)a$, the same argument uses the pure human and machine activities. At $p = 1/\overline{a}(0)$, humans work fully when $\overline{a}(1) > \overline{a}(0)$ and do not work when $\overline{a}(1) < \overline{a}(0)$, provided demand leaves positive machine output and both pure firms fit. The wage is $\max\{1, \overline{a}(1)/\overline{a}(0)\}$. At equality, every feasible employment split has the same cost. We verify the output bound and market fit in each application.

D.2 Firm-size envelopes.

Verification of Figure 14. Set $\phi = .3$, human coordination $y_*^H = .95$, and human supply $\mu = 2$. Both productive skills equal one for humans and AI. We vary AI coordination $y_*^M \in [.70, .975]$. At human share λ, preferred size capped by human supply is

$$\min\left\{m^*\left(y_*^M + \lambda(.95 - y_*^M)\right), \frac{2}{\lambda}\right\},$$

with no staffing cap at $\lambda = 0$. Panel (a) fixes $\lambda = 1/2$: size reaches its cap of four at $y_*^M = .93$. Panel (b) takes the lower and upper envelopes over shares.

The lower envelope is

$$\underline{M}(y_*^M, 2) = \min\{m^*(\min\{y_*^M, .95\}), 2\}. \tag{SIZE}$$

Every composition has preferred size at least the first term and a staffing cap at least two. A pure composition attains the smaller term. The lower envelope rises from zero at .70 to two at .90.

Below .95, preferred size rises with the human share while the staffing cap falls. At share zero the former binds; at share one the latter binds, since $m^*(.95) = 5 > 2$. Their unique crossing therefore gives the upper envelope. Substituting $\lambda = 2/m$ into the preferred-size formula gives

$$(1 - y_*^M)m^2 + (y_*^M - 1.2)m - 2(.95 - y_*^M) = 0. \tag{CAP}$$

The leading coefficient is positive and the constant is negative, so the quadratic has

exactly one positive root. At $y_*^M = .85$, that root is $(7+\sqrt{97})/6$, with human share $12/(7+\sqrt{97})$.

At or above .95, no composition has preferred size above the AI-only size, which has no staffing cap. Thus the upper envelope is $m^*(y_*^M) = .3/(1-y_*^M) - 1$. Both expressions equal five at .95. These envelopes bound equilibrium sizes under Proposition 1; the shaded band does not assert that every size occurs in equilibrium. □

D.3 Coordination below the scale threshold.

Verification of Figure 16. Set $\boldsymbol{y}^H = (1, 2, .5)$, human supply 20, $\phi = 1$, $\rho = -1$, $r = \underline{w} = 1$, and $P(Q) = (2 - Q/10)_+$. For path progress $\theta \in [0,1]$, we first find the best composition, then check market fit and human supply.

Step 1: optimal size and composition. AI's first skill is two on all three paths. Write $b = y_2^M$ for its second skill and $d = y_*^M - .5$ for its coordination advantage. Red has $b = 1$ and $d = .05 + .2\theta$; blue has $b = 1 + \theta$ and $d = .05$. Dark green has $b = 1 + \theta$ and $d = 0$. At machine share s, the mean skills are $x = 1 + s$ and $y = 2 + (b-2)s$. The coordination shortfall is $q = 1 - \overline{y}_* = .5 - ds$. Preferred size is $m = 1/q - 1 \in [1, 3]$, and log average product is

$$g(s) = \log \frac{2xy}{x+y} + 1 - q + q \log q.$$

Both inputs cost one, so firms maximize g. Differentiation gives

$$g_s = \frac{1}{x} + \frac{b-2}{y} - \frac{b-1}{x+y} + d \log(1/q), \qquad g_{ss} = -\frac{1}{x^2} - \frac{(b-2)^2}{y^2} + \frac{(b-1)^2}{(x+y)^2} + \frac{d^2}{q}.$$

On red, $g_{ss} \leq -1/2 + .25^2/.25 < 0$. On blue and dark green, $x \leq 2$, $x + y \geq 3$, $0 \leq b - 1 \leq 1$, and $q \geq .45$, so $g_{ss} \leq -1/4 + 1/9 + .05^2/.45 < 0$. Thus all three paths have a unique maximizing share, including corners. Initially $g_s(1/2) > 0$ on red and blue, so their optimal share exceeds $1/2$. Dark green starts at share $1/2$ and size one. On red, $g_s(0) > 0$ and $g_s(1) \leq -1/2 + .25 \log 4 < 0$, so the optimum stays interior. Since

$$g_{sd} = \log(1/q) + ds/q > 0,$$

the share rises strictly. On blue and dark green,

$$g_{sb} = \frac{2}{y^2} - \frac{3}{(x+y)^2} > 0,$$

because $x \geq 1$, $y \leq 2$, and hence $y/(x+y) \leq 2/3$. The share rises until it reaches one and then stays there; at $b = 2$, $g_s(1) = 1/4 + d\log(1/(.5-d)) > 0$.

Step 2: attain the surplus bound. Let $A = \exp(g(s))$ at the maximizing share. The equal-coordination mixed firm provides the lower bound $A \geq (3/2)\exp[(1 - \log 2)/2] > 1.5$ along all three paths. For every composition, the skill level is at most two and the scale factor is at most $\exp[.75 + .25\log(.25)] < 1.5$, so average product is below three. Optimal output is $Q = 10(2 - 1/A) > 40/3$. Every preferred firm produces less than $3 \cdot 3 = 9 < Q$ and uses fewer than 20 humans. Copies of the optimal firm with activity weight $Q/(mA)$ use $(1-s)Q/A < 20$ humans in aggregate. The common verification therefore gives equilibrium at price $p = 1/A$ and wage $w = 1$, with surplus $Q^2/20$. Strict composition optimality and unique preferred sizes identify the industry.

Step 3: separate the two size responses. On blue, size is $m = 1/(.5 - .05s) - 1$, so

$$\frac{\mathrm{d}m}{\mathrm{d}\theta} = \frac{.05}{(.5 - .05s)^2}\frac{\mathrm{d}s}{\mathrm{d}\theta}.$$

Size rises only through composition and stops rising when the machine share reaches one. On dark green, $d = 0$, so mean coordination stays at .5 and size stays at one despite the rising machine share. On red, differentiation gives

$$\frac{\mathrm{d}m}{\mathrm{d}\theta} = \frac{.2s + d\ \mathrm{d}s/\ \mathrm{d}\theta}{(.5 - ds)^2}.$$

The first term raises size at fixed composition. The second is strictly positive throughout the interior of the path, so composition amplifies growth. All three paths raise average product, output, and surplus. At the same time, activity weight falls: total input $Q/A = 10(2/A - 1/A^2)$ falls with $A > 1$, and firm size is nondecreasing. The figure plots size and share, not firm counts. □

D.4 Competence and human exit.

Verification of Figure 23. Both paths set $\rho = -1$ and $r = \underline{w} = 1$. We bound output for every firm, then construct industries that attain the surplus bound.

Step 1: complementary mixing. Set $\boldsymbol{y}^H = (1, 2, .9)$, $\boldsymbol{y}^M = (1.5, \theta, .9)$ for $\theta \in [1, 2]$, $\phi = .3$, $\mu = 20$, and $P(Q) = (2 - Q/10)_+$. Every share has unique preferred size $m^* = 2$. Use the optimized log scale factor ℓ from (CES-PEAK). The best average

product at human share λ is

$$\bar{a}(\lambda;\theta) = e^{\ell(.9)}\left[\frac{2(1.5-.5\lambda)(\theta+(2-\theta)\lambda)}{1.5+\theta+(1.5-\theta)\lambda}\right].$$

The two productive capability profiles are not proportional anywhere on this path. Writing $x = 1.5-.5\lambda$ and $y = \theta+(2-\theta)\lambda$, differentiation gives

$$\bar{a}_{\lambda\lambda} = -\frac{4e^{\ell(.9)}(\theta-3)^2}{(x+y)^3} < 0.$$

Thus mixing is strictly complementary throughout the path.

The unique maximizing share λ^* is the interior root of

$$\frac{\partial}{\partial\lambda}\log\bar{a} = \left[\frac{-.5}{1.5-.5\lambda} + \frac{2-\theta}{\theta+(2-\theta)\lambda} - \frac{1.5-\theta}{1.5+\theta+(1.5-\theta)\lambda}\right] = 0$$

when $1 \le \theta < 1.5$, and is zero when $1.5 \le \theta \le 2$. Indeed, the derivative at $\lambda = 1$ is $-\theta/6 < 0$, while its sign at zero is the sign of $4.5 - 2.25\theta - .5\theta^2$, which vanishes at $\theta = 1.5$. Differentiating the first-order condition, with $\bar{y}_1$ and $\bar{y}_2$ denoting the two mean productive capabilities, gives

$$\frac{\mathrm{d}\lambda^*}{\mathrm{d}\theta} = -\frac{(1-\lambda^*)\bar{y}_2 + \bar{y}_1^2}{(2-\theta)(\bar{y}_1+\bar{y}_2)} < 0 \qquad (1 \le \theta < 1.5).$$

The share thus falls and converges to zero at 1.5. At the transition, the condition in Proposition 5(i) holds with equality.

Let $A := \bar{a}(\lambda^*;\theta)$ be the largest average product. Apply the common verification method above. The unique maximizing output is $Q = 10(2-1/A)$. Since $1 < A < 3$, this output exceeds 10, while every preferred-scale firm produces less than 6 and uses at most two humans. Use weight $Q/(2A)$ of the firm with share λ^* and size two. Total human use is $\lambda^* Q/A < 20$, so human supply stays slack. The common verification gives equilibrium at $p = 1/A$, $w = 1$. Strict concavity and the unique preferred scale make the industry unique. As θ rises, A increases and $Q/A = 20/A - 10/A^2$ decreases because $A > 1$. Both the human share and human employment therefore fall continuously.

Step 2: segregating mixing. Set $\boldsymbol{y}^H = (1.1, 1.1, .985)$, $\boldsymbol{y}^M = (\theta, \theta, .9999)$ for $\theta \in [1, 1.2]$, $\phi = .02$, $\mu = 2/3$, and $P(Q) = (2 - Q/1000)_+$. At human share λ, mean competence

is $\theta + (1.1 - \theta)\lambda$ and mean coordination is $.9999 - .0149\lambda$. Thus

$$\frac{\partial^2}{\partial \lambda^2} \log \bar{a} = -\frac{(1.1-\theta)^2}{[\theta + (1.1-\theta)\lambda]^2} + \frac{.0149^2}{.0001 + .0149\lambda} \geq -\frac{1}{100} + \frac{.0149^2}{.015} = \frac{7201}{1500000} > 0.$$

Strict log convexity implies strict convexity of $\bar{a}$. Every firm with human input h and machine input a consequently obeys

$$F(h\delta_{\mathbf{y}^H} + a\delta_{\mathbf{y}^M}) \leq \bar{a}(1)h + \bar{a}(0)a, \qquad \bar{a}(1) = 1.1e^{\ell(.985)}, \quad \bar{a}(0) = \theta e^{\ell(.9999)}.$$

Pure firms attain the bound at sizes $1/3$ and 199, respectively.

The bound gives industry resource cost at least $Q/\bar{a}(0) + (1 - \bar{a}(1)/\bar{a}(0))T_H$, where $T_H \in [0, \mu]$ is total human employment. Maximize the resulting surplus bound. Output is $Q = 1000(2 - 1/\bar{a}(0))$; humans supply $T_H = \mu$ if $\bar{a}(1) > \bar{a}(0)$ and $T_H = 0$ if $\bar{a}(1) < \bar{a}(0)$. At equality, any $T_H \in [0, \mu]$ attains the bound. The required machine input is $(Q - \bar{a}(1)T_H)/\bar{a}(0) > 0$. Throughout this path $1 < \bar{a}(0) < 2$ and $\bar{a}(1) < 2$, so $Q > 1000$, whereas the machine firm produces less than 398 and the human firm less than $2/3$. The human firm needs $1/3 < \mu$ workers. Both preferred-scale firms can therefore be staffed and fit in the market, and their activity weights implement every stated allocation. At prices $p = 1/\bar{a}(0)$ and $w = \max\{1, \bar{a}(1)/\bar{a}(0)\}$, the output bound verifies support for all firms. The constructed industries are equilibria and attain the global surplus bound. Strict convexity rules out mixed firms in every equilibrium. The unique crossing solves

$$\theta = 1.1 \exp\left(\ell(.985) - \ell(.9999)\right) \simeq 1.0796369.$$

At this crossing, the full interval of employment fractions in Figure 23(f) consists of equilibria. □

D.5 Human coordinators in AI-heavy firms. We reduce entry to the edges of the composition triangle, then check Figure 22 numerically at 99 sampled AI profiles. Productive humans have profile $(6, 6, .2)$ and supply 10; coordinators have profile $(1, 1, .9)$ and supply 1; AI has profile $(\theta, \theta, .2)$ with $\theta \in [4, 12]$. The other parameters are $\rho = -1$, $\phi = 3$, $r = \underline{w} = 1$, and $P(Q) = 2(1 - Q/3000)_+$.

Reduction to two-type edges. For productive-human, coordinator, and machine shares (s_P, s_C, s_M) summing to one, mean competence is $6s_P + s_C + \theta s_M$, and mean coordination is $.2 + .7s_C$. Preferred size is $3/(.8 - .7s_C) - 1 > 0$, and optimized average product is

$$(6s_P + s_C + \theta s_M)e^{\ell(.2+.7s_C)},$$

where ℓ is defined in (CES-PEAK) with $\phi = 3$. At fixed coordinator share, this expression is affine in the productive human share, as is the input bill. Thus profit per unit of input is maximized at an endpoint of each fixed-coordinator-share segment. Checking the two-type edges therefore bounds entry by every composition; optimizing size first also covers every scale, including the zero-size limit.

Prices, allocations, and numerical bounds. The solver imposes zero profit on active firm types, composition optimality for their positive shares, demand clearing, and human-supply complementarity. A slack human supply fixes its wage at one; a wage above one requires full employment of that type. For each candidate state, adaptive Taylor bounds on each edge give an upper bound on entry profit per unit of input, with a 10^{-11} floating-point guard. All 99 bounds are below 10^{-7}, and the largest equilibrium-equation residual is below 3×10^{-14}. Each active firm uses less than the available stock of each human type and produces less than market demand. These are numerical checks at the sampled points, not an exact certificate for the interpolated segments.

At $\theta = 4$, all humans work together, with wages about 1.501 for productive humans and 1.624 for coordinators; AI-only firms supply the remaining output. At $\theta = 12$, productive humans are unemployed. Coordinators remain fully employed at wage 1.005586 in firms with coordinator share .0127096 and machine share .9872904.

Source: `images/pythoncode/make_coordinator_ai_sorting.py`.
Records: `intel-figures/coordinator_ai_sorting_checks.csv` and `intel-figures/coordinator_ai_sorting_firms.csv`.

D.6 Separate and mixed firms in the CES example.

Verification of Figure 19. We derive the envelopes and check prices, staffing, and market fit. Both panels use $\rho = -1$, $\phi = 1$, $r = \underline{w} = 1$, human supply 50, and $P(Q) = (2 - Q/100)_+$. For mean coordination $z \in [.4, .8]$, preferred size is $m^*(z) = 1/(1-z) - 1 > 0$. Define the optimized scale factor by

$$\mathcal{S}(z) := \frac{S(m^*(z), z)}{m^*(z)} = e^z (1-z)^{1-z}.$$

Its log derivatives are $(\log \mathcal{S})'(z) = -\log(1-z)$ and $(\log \mathcal{S})''(z) = 1/(1-z) > 0$.

Panel (a): separate firms. Humans have profile $(1, 1, .8)$ and machines have profile $(1, 1, .4)$. Productive competence is one at every human share, so

$$\bar{a}(\lambda, w) = \frac{\mathcal{S}(.4 + .4\lambda)}{1 - \lambda + \lambda w}.$$

Set $w = \mathcal{S}(.8)/\mathcal{S}(.4) > 1$ and $p = 1/\mathcal{S}(.4)$. Both endpoints then equal $1/p$. The log envelope is strictly convex because

$$\frac{\partial^2}{\partial \lambda^2} \log \bar{a}(\lambda, w) = \frac{.16}{.6 - .4\lambda} + \frac{(w-1)^2}{[1 + \lambda(w-1)]^2} > 0.$$

Every interior share therefore has cost-adjusted best average product below $1/p$. Since the envelope optimizes scale, this bound supports prices against every firm and excludes active mixed firms.

Demand gives $Q = 100[2 - 1/\mathcal{S}(.4)] \simeq 108.927$. Employ all 50 humans in human-only firms of size 4, producing $50\mathcal{S}(.8) \simeq 80.651$. Supply the remaining output with machine-only firms of size $2/3$, using total machine input

$$\frac{Q - 50\mathcal{S}(.8)}{\mathcal{S}(.4)} \simeq 25.751 > 0.$$

Both active firm types break even. Human supply binds, as required by $w \simeq 1.469 > 1$. Each human-only firm uses $4 < 50$ humans and produces $4\mathcal{S}(.8) \simeq 6.452 < Q$; each machine-only firm produces $(2/3)\mathcal{S}(.4) \simeq .732 < Q$.

Panel (b): mixed firms. Humans have profile $(2, 1, .6)$ and machines have profile $(1, 2, .6)$. Mean coordination is .6 and preferred size is $3/2$ at every share. At wage $w = 1$, the harmonic-mean competence aggregator gives

$$\bar{a}(\lambda, 1) = \mathcal{S}(.6)\frac{2(1+\lambda)(2-\lambda)}{3}.$$

Its second derivative is $-4\mathcal{S}(.6)/3 < 0$, and its unique maximum is at $\lambda = 1/2$, with value $3\mathcal{S}(.6)/2$. Set $p = 2/[3\mathcal{S}(.6)]$ and $Q = 100\{2 - 2/[3\mathcal{S}(.6)]\} \simeq 147.215$. Use only firms of size $3/2$ with human share $1/2$. Total human input is

$$\frac{Q}{3\mathcal{S}(.6)} \simeq 38.854 < 50,$$

so human supply is slack and $w = \underline{w} = 1$ satisfies human scarcity. Each firm uses $3/4 < 50$ humans and produces $9\mathcal{S}(.6)/4 \simeq 2.842 < Q$. The envelope bound proves firm support, and each active firm breaks even.

In both panels, activity equals total input divided by firm size. The stated prices clear demand and satisfy support and human scarcity, so Theorem 1 gives equilibrium. □

D.7 Numerical checks and records. The analytic proofs above cover all feasible shares and scales in their stated cases. The table records checks for the numerical

figures; their bounds apply only at sampled parameter values and to the stated tolerances.

Example	Check and scope	Record
Human coordinators	Entry upper bounds below 10^{-7} per unit of input at 99 points; all compositions and scales.	`coordinator_ai_sorting_checks.csv`
Finite-market concentration	Pure-firm optimization, with additional entry searches at selected states.	`ces_support_jump_certificate.csv`; `wage_jump_checks.json`
Large changes in firms	Finite-menu optimization at 21 points; additional entry searches find profits up to 4.7×10^{-6}.	`ces_tipping_certificate.csv`

Table 3: Computational checks. The coordinator records are in `intel-figures/`; other file paths and supplementary checks are in `images/v5/canonical_ces/REPRODUCTION.md`.

Finite-market concentration. The curves in Figures 17 and 26 optimize output over pure human and machine firms. Additional entry searches at selected states do not establish equilibrium over the full paths. The analytic concentration result and pointwise wage limit instead follow from Proposition 3 and the verification in Section A.3.4, respectively.

Large changes in firms. Figure 25 uses a finite menu with atomistic activities and the market constraint. A finer menu at the start, wage peak, and endpoint changes either wage by less than 10^{-4}. The recorded entry searches do not bound profit over all feasible firms. For the dashed curve, the calculation upgrades machines in the initial machine-using firms, holds their inputs fixed, and adds the change in the bid envelope to the initial wage. Human-only employers and binding staffing limits can prevent the initial envelope from equaling that wage. The gap between the curves is therefore a descriptive comparison, not an application of the interior-match derivative formula throughout the path. The snapshots show the largest human employer; the diamond denotes atomistic production and has no firm-size meaning.

E CALCULATIONS FOR THE EXAMPLE

This appendix records the parameters, allocations, and checks behind the example in Sections 2.1 and 2.4. The two markets share one population of six coders, six writers, and two coordinators. Employment totals therefore count each human once. Compute costs $r = 1$ per copy. The human outside option is $1/4$, except in Section 2.4, where it is one.

E.1 Capabilities and firm scale. All illustrated AI profiles have equal task skills, $\boldsymbol{y}^M = (\theta, \theta, y_*^M)$, so AI competence equals θ. The human population averages are

$$\bar{L}_H = \frac{12(1.8)+2}{14} = \frac{59}{35}, \qquad \bar{y}_*^H = \frac{12(0.75)+2(0.97)}{14} = \frac{547}{700}.$$

For these symmetric profiles, the regime definitions in Section 2.3 reduce to

$$\begin{aligned} \text{Minimal AGI:} &\quad \bar{L}_H \le \theta \le 9, \quad \bar{y}_*^H \le y_*^M < 1, \\ \text{Supercompetent:} &\quad \theta > 9, \quad y_*^M \le 1, \\ \text{Supercooperative:} &\quad \theta \le 9, \quad y_*^M > 1, \\ \text{ASI:} &\quad \theta > 9, \quad y_*^M > 1. \end{aligned}$$

The boundary $y_*^M = 1$ has a finite productivity ceiling. It must be distinguished from coordination strictly above one.

At fixed mean capabilities, write $z = \bar{y}_*$ for mean coordination. The scale formulas in Section D.1, with $\phi = 7/4$, give

$$m^*(z) = \frac{7}{4(1-z)} - 1, \qquad G(z) = e^{7/4-(1-z)} \left(\frac{1-z}{7/4}\right)^{1-z}, \qquad 1-\phi < z < 1.$$

Best output per worker is $LG(z)$. Every below-one coordination value used in the example lies in this range. Thus $m^*(3/4) = 6$, $G(3/4) \approx 2.75529$, $m^*(\bar{y}_*^H) \approx 7.00654$, and $G(\bar{y}_*^H) \approx 2.93511$.

At coordination one, output per worker approaches the finite, unattained ceiling $Le^{7/4}$ as firm size grows. Strictly above one, it grows without bound. In a finite market, a complete firm must fit within market output and human supplies. At specialist and coordinator shares s_s and s_c, its productivity-maximizing feasible size is

$$m = \min\left\{m^*(z), \frac{12}{s_s}, \frac{2}{s_c}, m_{\text{fit}}(Q)\right\}, \qquad F(m_{\text{fit}}(Q)) = Q.$$

Here shares remain fixed when solving for $m_{\text{fit}}(Q)$; we omit a staffing bound when its

share is zero and set $m^*(z) = +\infty$ for $z \geq 1$.

E.2 Markets, prices, and numerical method.

Positive production. Some example profiles fail the standing atomistic-machine viability condition: for minimal AGI, $r/L(\delta_{\mathbf{y}^M}) = 35/59 > 0.3$. The cost and supporting-price arguments in Section A.1.1 still apply; that condition ensures positive surplus, not cost attainment. Here a feasible human allocation gives positive surplus directly. Two balanced specialist firms of size six produce $Q = 60G(3/4) \approx 165.32$ in the niche market. Even at outside wage one, their surplus is $0.3Q - 0.00075Q^2 - 12 \approx 17.10 > 0$. Each complete firm fits within total output, and together they use the twelve specialists once. Thus the example has positive production at either outside wage.

For large-market output, replace atomistic machines in the argument of Section A.2.3 by a fixed finite AI firm with unit cost below 0.3. The minimal-AGI preferred firm has unit cost about 0.20211; the other displayed AI profiles also admit such a firm. Replicating this recipe is feasible once market output exceeds its capacity and gives surplus proportional to market size. Hence efficient output diverges. This argument does not establish convergence to the reported limiting allocations.

Firm activity weights can be fractional. A complete firm with size m and activity weight a contributes am inputs and $aF(m)$ output. Its complete recipe must still fit within the market and human supplies. A box in Figures 4 and 7 shows that recipe; its multiplier gives the activity weight. These weights are not integer firm counts.

We minimize resource cost at each niche output, allowing humans to work in either market, then maximize surplus over output. We recover internal prices and wages from the cost problem and search for profitable feasible entrants. Note here that the internal price p finances production and need not equal buyers' marginal willingness to pay $P(Q)$. For one pure-AI firm serving the whole market at an interior efficient size,

$$P(F(m))F'(m) = r, \qquad pF(m) = rm.$$

When average product rises with size, $mF'(m) > F(m)$, so $p > P(F(m))$. Transfers then reconcile these two prices.

Without AI, fixed human supply keeps deep-market output bounded and its demand price tends to 0.3. With AI coordination below one, the limiting deep-market price is instead $p = r/[\theta G(y_*^M)]$ whenever this unit cost is below 0.3, as it is in the displayed allocations. Deep output then satisfies

$$\frac{Q}{200N} \longrightarrow 1 - \frac{p}{0.3}.$$

At coordination one the limiting unit cost is $r/(\theta e^{7/4})$. Strictly above one it tends to zero; a demand price tending to zero requires $Q/(200N) \to 1$. We use these limiting prices in the joint allocation calculation. The numerical records do not establish convergence of finite-market equilibria to each displayed allocation.

For wage calculations at positive scale, the full marginal product of an added worker with profile $\boldsymbol{y}$ is

$$D_{\boldsymbol{y}}F = \frac{F}{m}\left[1 + \frac{\nabla L \cdot (\boldsymbol{y}_{1:2} - \bar{\boldsymbol{y}}_{1:2})}{L} + \log(1+m)(y_* - z) + \frac{m(z-1)}{1+m} + \frac{\phi m}{(1+m)^2}\right].$$

The task-skill vectors $\boldsymbol{y}_{1:2}$ and $\bar{\boldsymbol{y}}_{1:2}$ contain the first two coordinates, and we evaluate ∇L at the firm's mean task skills. The last two terms cancel at an unconstrained best size. Equal mean task skills give $\nabla L = (1/2, 1/2)$. An interior employment choice satisfies $w = pD_{\boldsymbol{y}}F$; a binding staffing or market constraint requires the corresponding boundary condition.

E.3 Human-only benchmarks. For one human type, six coders at best size produce $6(1.8)G(3/4) \approx 29.76$, compared with 21.79 from six solo coders. Twelve coders produce 59.51 in two firms, versus 57.22 in one merged firm. The deep-market wage is $0.3(1.8)G(3/4) \approx 1.488$.

Equal coder and writer shares raise competence to five while keeping coordination at $3/4$. Two balanced six-person firms produce $2(6)(5)G(3/4) \approx 165.32$, versus 158.93 in one merged firm. The specialist wage is $0.3(5)G(3/4) \approx 4.133$.

A team of three coders, three writers, and one coordinator has mean competence $31/7$ and coordination $5.47/7$. Two such seven-person teams produce 181.98, versus 175.66 when merged. Seven is the rounded preferred size 7.00654: with fractional activities, the same proportions use activity $14/7.00654$ and produce the same total to two decimal places. Marginal products at size seven are 13.828 for specialists and 8.034 for coordinators, giving the reported wages 4.15 and 2.41 after rounding. These calculations explain the human benchmarks in Figure 2.

E.4 The four regime allocations. Table 4 records the current snapshots in Figure 4. Specialist and coordinator employment cover both markets; wages for unemployed coordinators equal their outside option.

Minimal AGI. Pure-AI firms have size 7.00654 and output per copy 4.94775, giving unit cost 0.202112. A human firm in the deep market employs both coordinators and 5.13592 specialists. The remaining 6.86408 specialists work across AI firms, with human shares per firm tending to zero. The niche uses AI-only firms with activity 1.88246. Deep-market AI activity grows as approximately $1.88N$.

	Minimal AGI	Supercompetent	Supercooperative	ASI
AI task skills θ	59/35	12	59/35	12
AI coordination	547/700	547/700	1.60	1.05
Niche output	65.259	181.005	180.161	187.368
Niche internal price	.202112	.028620	.030052	.019523
Niche demand price	.202112	.028493	.029759	.018948
Specialist wage	2.90072	.39428	.41443	.26896
Coordinator wage	1.29801	.47561	.25000	.25000
Specialists employed	12	12	12	12
Coordinators employed	2	2	.97257	0

Table 4: Joint-market calculations for the example, with outside option 1/4.

Supercompetent AI. At task skills twelve, deep AI firms retain size 7.00654 and have output per copy 35.22127. The two coordinators work across these firms. Their limiting wage is

$$\frac{1}{12} + (0.97 - \bar{y}_*^H)\log(1 + m^*(\bar{y}_*^H)) \approx 0.47561.$$

The niche employs all twelve specialists in two balanced firms of six. It also uses a complete AI recipe of size 5.18036 with activity 0.08667. The niche internal price is 0.028620, above the deep price 0.028392, because the preferred deep AI recipe does not fit in the niche. Specialist pay is the niche internal price times $5G(3/4)$, or 0.39428. Thus coordinator pay exceeds specialist pay in this snapshot.

Supercooperative AI. At task skills 59/35 and coordination 1.60, the niche uses human-only firms of size 6.51455 and mixed firms of size 9.22247, with activities 1.91864 and 0.07281. Coordinators make up 0.077812 of the human-only recipe; specialists make up 0.705170 of the mixed recipe. Total employment is twelve specialists and 0.97257 coordinators. Thus 1.02743 coordinators are unemployed. Deep production has an AI-only limiting profile.

ASI. At profile $(12, 12, 1.05)$, the niche uses two balanced human firms of six and a complete AI recipe of size 3.65806 with activity 0.11769. The human firms supply 165.317 units; AI supplies the remaining 22.051. All specialists work, while both coordinators are unemployed. The deep-market limiting unit cost is zero, although the finite niche's internal price is positive.

Along the competence route, preferred pure-AI size stays fixed and best output per copy is $\theta G(\bar{y}_*^H)$. Along the coordination route it has preferred sizes 34 at 0.95 and 174 at 0.99, with best output per copy approaching $(59/35)e^{7/4} \approx 9.70062$ as coordination approaches one. Figure 6 illustrates the wage and size changes along these routes; its intermediate wage curves have not received the same joint output and entry checks as the four snapshots.

E.5 Within-firm and between-firm displacement. Only Section 2.4 raises the outside-option wage to one. Both displayed paths start with AI coordination below the human coordinator's 0.97, and coordinators remain unemployed throughout.

Within-firm displacement. Fix AI coordination at 0.96 and raise both task skills from 3 to 3.4. All employed humans work in one mixed niche firm with activity one; the deep market is AI-only. At competence three, the niche firm has size 11.00143 and employs 4.57923 specialists. At competence 3.15, specialist employment is 2.39314. It reaches zero near 3.29535. Human wages remain at one. A pure-AI firm at its preferred size 42.75 would produce 609.62 units at competence three, exceeding niche output 152.80, so market size constrains the firm's scale.

With both input costs equal to one, an interior mix equates human and AI marginal products. If s is the balanced-specialist share, competence is $L = \theta + (5-\theta)s$ and coordination is $z = y_*^M + (0.75 - y_*^M)s$. Differentiating output in s at fixed size gives the condition

$$\frac{5-\theta}{\theta + (5-\theta)s} + (0.75 - y_*^M)\log(1+m) = 0.$$

At exit, set $s = 0$ and combine this condition with the pure-AI output choice $P(F(m))F'(m) = 1$. These equations give the stated exit value. The displayed segment starts after the earlier replacement of human-only niche firms.

Between-firm displacement. Hold task skills at four and raise coordination from 0.80 to 1.04. Balanced human firms have size six and output per worker $5G(0.75) \approx 13.77645$. Near the employment threshold, pure-AI firms also operate at their preferred size. Their unit cost crosses the human unit cost at

$$4G(y_*^M) = 5G(0.75), \qquad y_*^M \approx 0.85210100.$$

The preferred AI firm then has size 10.83240 and produces 149.23197, which fits within niche output 151.60823. Numerical entry checks support separate production near this crossing. Specialist employment is twelve below the threshold and zero above it; at equality, any employment from zero to twelve has the same cost. The wage falls from 1.14155 at coordination 0.80 to one at exit. Initially, specialists work in mixed deep-market firms. They move into human-only firms near the threshold while total employment stays at twelve.

For Figure 7(b), we select the pre-exit allocation with both six-person human firms in the deep market and AI-only firms in both markets. It has the same surplus as the joint-market calculation; human firms need not sort uniquely to that market. The dots show preferred sizes at the threshold, and $\times$many denotes the number of AI firms growing without bound in the deep-market limit.

E.6 Calculation records and scope. The regime snapshots use `Example/equilibrium-checks/four_regimes_refined.json`, except for the supercooperative snapshot, which uses the coordination-1.60 state in `Example/displacement-checks/refined_paths.json`. Their generating scripts refine output choices and search over entrant compositions, including boundary compositions and firms that employ coordinators. The largest recorded entry gain per input across these four snapshots is below 10^{-10}.

The current displacement paths are stored in `Example/displacement-checks/within_between_paths.json`. The between-firm path has output and entry checks at 47 states; the within-firm path has four such checks, including its exit point. The path scripts compare output choices and refine continuous composition searches. These are numerical checks, not global analytic certificates over every entrant or every point on a path. The files do not prove uniqueness at the employment threshold or convergence from finite deep markets. Appendix D records the separate calculations for the general model's figures.