The Economics of Superabundant AI: Autonomy, Scarcity, and the Future of Work

Luis Garicano

Submission to NBER Volume on Transformative AI.

September 30, 2025.

---

1. Introduction: When Do Humans Become Horses?

If AI can create a "country of geniuses in a data center," as Dario Amodei has put it, what role remains for human workers? When AI can do the same task for less, why hire a human?

Agrawal, Gans, and Goldfarb (2025) model a planner who routes work to humans or AI "geniuses." With abundant, effective AI, routine workers are displaced, and humans shift to frontier tasks. But the data are mixed. High-frequency payroll data (Brynjolfsson et al. 2025b) do show early-career employment falling in more AI-exposed jobs. At the same time, studies in firm deployment of AI co-pilots show significant augmentation: Brynjolfsson et al. (2025a) found the largest gains accruing to less experienced and lower-skilled customer support agents.

How can AI simultaneously augment entry-level workers in specific tasks yet displace them in aggregate?

This puzzle highlights that the simple logic of factor costs—why pay a human when AI is cheaper? —is partial equilibrium. The general-equilibrium question is: once AI scales, what is scarce?

Following Ide and Talamàs (2025), two dimensions determine the outcome: autonomy (can a problem be solved without a human act?) and the problem-to-compute ratio (are there more valuable problems or "slots" than compute can handle?). When compute is scarce relative to "slots," compute earns a price; when compute is plentiful and slots are scarce, slot rights earn a price. Only with autonomy and a compute glut do lower-skill humans become redundant. I show here that endogenizing the flow of "addressable opportunities" can prevent collapse when the cost of creating a slot that AI can attempt stays below the value the AI can realize.

2. A Minimal Environment

I use a stripped-down version of the standard knowledge-hierarchy model with ex-ante heterogeneous agents (Antràs, Garicano and Rossi-Hansberg, 2006), as extended by Ide and Talamàs (2025) to include AI. This literature studies how organizations structure the use of tacit knowledge. AI, unlike previous Information Technology, can acquire tacit knowledge. The environment treats intelligence, human or artificial, as problem solving.

A solved problem yields value 1. Humans have heterogeneous knowledge $z \in [0,1]$. A worker with knowledge z successfully solves a fraction z of the problems they encounter. If the worker fails (probability $1-z$), they can consult a more knowledgeable agent, a solver. Each escalation consumes solver helping time $h \in (0,1)$.

AI has knowledge z_{AI}. Let μ be per-period compute capacity (AI "worker" units). A finite stock of opportunities φ is the period's assignable problems. We study whether compute is scarce or abundant relative to φ.

As shown by Ide and Talamàs (2025), the competitive equilibrium allocates human labor, compute, and problems across one-layer (independent) or two-layer (worker-solver) structures, where either role may be compute. Matching is positive assortative, and profits are 0. Three prices pin down the equilibrium: the rental price of compute, r^*, the price of a problem (opportunity), p^*, and the wage schedule, *w(·)*. Who gets paid depends on what is scarce. In the baseline, where φ is large relative to compute, the marginal unit of compute must work alone. This implies the rental price of compute equals its standalone productivity, $r^* = z_{AI}$, and the price of opportunities is 0, $p^* = 0$.

3. Results: when are humans redundant?

Compute Scarce Relative to Problems: Co-Pilot Economics. When the stock of valuable problems φ is large relative to the available compute capacity μ, the marginal unit of compute works alone. As established, this pins the $r^* = z_{AI}$ and $p^* = 0$. There is no unemployment. Humans are valuable precisely because compute is scarce. Humans with low knowledge, $z < z_{AI}$ work with assistance; humans with $z > z_{AI}$ act as solvers. Workers are paid for the compute (or human) time they save by solving problems rather than escalating. Solvers are paid for the value generated by leveraging knowledge above the AI's baseline**.** These are the Ide–Talamàs (2025) baseline post-AI predictions. In this regime, compute is a scarce, valuable fuel. Humans are paid for "fuel efficiency." As long as the fuel has a positive price ($r^* > 0$) human time retains value.

Compute Glut with Autonomous AI: When Problems are the Bottleneck. Now, let compute exceed available problems: μ ≥φ, and assume autonomy: AI can attempt a problem with no required human act, though escalation to a human is allowed. Some compute is idle, hence the rental price of compute collapses: $r^* = 0$. The bottleneck becomes φ; the price of opportunities is pinned by AI's stand-alone success $p^* = z_{AI}$ (Ide-Talamas (2025) Section 6). All humans with less knowledge than AI, $z < z_{AI}$, become unemployed. Humans with $z > z_{AI}$ remain as solvers, handling exceptions AI cannot solve. As z_{AI} rises, employment shrinks. Compute earns nothing, most labor earns little, and rents flow to the owners of the scarce opportunities.

Compute Glut with Non-Autonomous AI: Wage Compression, Not Collapse. If μ ≥φ, but AI is non-autonomous, each unit of output requires a human input somewhere—

judgment, accountability, or a physical act. Every problem requires a person, so human time is now the bottleneck. All income flows to labor. Everyone works assisting AI. Wages compress around the ability of the human provider, because deep knowledge adds less when paired with highly capable, free AI. Compression is not equality: better humans still raise success probabilities. Neither compute, not problems carry any rent ($r^* = p^*$ =0). (See Ide- Talamas, Online Appendix 4.3.)

4. The role of autonomy

Autonomy, not raw z_{AI}, determines who gets paid when compute is cheap. The current state of AI in radiology provides a compelling case study (Mousa, 2025). Radiology seemed optimized for replacement. In 2016, Geoffrey Hinton said "people should stop training radiologists now." By 2017, models like CheXNet beat radiologists in some benchmarks. Yet, demand and wages for radiologists are at all-time highs. In practice, these systems remain non-autonomous: real-world variation and liability keep a person in each slot. Income therefore flows to labor.

5. The Stock of Problems

With autonomous AI, the future of work turns on the race between compute growth and the growth of the stock of addressable opportunities, φ. If φ cannot scale with compute, the economy slides into the compute-glut regime.

Endogenize φ. Entrepreneurs can create opportunities at constant unit cost C> 0. A marginal opportunity worked by stand-alone AI must break even, so the slot price plus the rental of the AI must equal the value generated by AI (the probability it solves the problem): $p + r = z_{AI}$. Free entry pins the opportunity price to cost, but not above its value, thus $p^* = \min\{C, z_{AI}\}$.

Proposition (Endogenous opportunities prevent collapse). *In a competitive equilibrium with autonomous AI and free entry into opportunity creation,*

1. *if $C < z_{AI}$, (Entry is profitable), entry expands φ until compute is scarce. The price of compute is positive $r^* = z_{AI} - C$, which sustains positive wages for workers and solvers.*
2. *If $C \geq z_{AI}$* (Entry is unprofitable), opportunity creation stalls, then $p^* = z_{AI}$ and $r^* = 0$, *and the bottleneck logic of Section 3.2. applies, with low-skill worker displacement.*

Proof. Combine the zero-profit condition $p + r = z_{AI}$, with the free entry condition, $p^* = min\ \{C, z_{AI}\}$.

Case 1: $C < z_{AI}$. Then $r^* > 0$. Entry expands φ until it absorbs compute at rental $r^* = z_{AI} - C > 0$. With a positive price for compute, the wage schedule is positive, since workers and solvers are paid for the compute time they save or unlock.

Case 2: If $C(z_{AI}) \geq z_{AI}$, then $r^* \leq 0$. Compute cannot have a negative price, hence $r^* = 0$ and the opportunity price absorbs all value, reproducing the bottleneck case of the second paragraph of Section 3. ∎

Intuitively, think of z_{AI} as the value of handing AI a clean task. Think of C as the "packaging cost" of turning a messy real-world need into such a task—specifying the job, providing the right data, verifying the result, and making it safe to pay for. If this "packaging cost" (C) is lower than the value the AI creates z_{AI}, then firms keep inventing new jobs. Compute stays scarce, people stay involved, and wages stay positive. If packaging cost is as high as value, entry stalls, the flow of addressable work φ stalls, AI fills the limited existing jobs, and low-skill humans are left out.

The condition $C < z_{AI}$ shifts the central question. Collapse is not driven by machine intelligence per se, but by the relative cost of turning latent needs into addressable opportunities. Policy and design shape C : data access, verification, liability, and payment systems. Anything that lowers C moves the economy away from collapse. Radiology illustrates this: the "problem" is patient management, not image classification; much of the value lies in packaging, communication, and responsibility.

6. Conclusion

The "humans as horses" analogy is compelling but relies on partial-equilibrium intuition. General equilibrium asks what remains scarce after AI.

Three regimes follow. When compute is scarce, AI is a co-pilot and human time has value because it saves compute. When compute is abundant and AI is autonomous, AI fills all slots, displacing low-skill workers and shifting rents to slot owners. When compute is abundant and AI is non-autonomous, every slot needs a person, there is no unemployment, and wages compress.

This resolves the opening puzzle: AI can raise entry-level performance on single tasks (augmentation) yet still reduce entry-level jobs in total (displacement). Augmentation dominates when compute is scarce or AI requires supervision. Displacement occurs when autonomous AI is abundant and the flow of new opportunities (ϕ) lags, allowing AI to fill existing slots and push novices out.

Endogenizing problem supply adds a policy lever. When it costs less to turn a messy need into a clear, checkable task than the value AI can produce on it, firms will keep creating such tasks. That keeps compute scarce, makes human time valuable, and sustains broad employment even as AI improves. Hence turning humans into "horses" requires the three conditions to be true at once: AI must be autonomous, compute capacity must exceed available opportunities, and the cost of defining new, well-specified problems must be high enough to bottleneck the creation of new opportunities.

The economic imperative is to turn latent needs into purchasable units of work, growing φ faster than compute. Examples from personalized education to continuous medical care to large scale infrastructure or scientific discovery all point to a vast stock of hard problems still to address.

References

Ajay Agrawal, Joshua S. Gans and Avi Goldfarb. “Genius on Demand: The Value of Transformative Artificial Intelligence”. NBER Forthcoming.

Antràs, Pol, Luis Garicano, and Esteban Rossi-Hansberg. 2006. “Offshoring in a Knowledge Economy.” *The Quarterly Journal of Economics* 121 (1): 31–77.

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond. "Generative AI at work." The Quarterly Journal of Economics 140, no. 2 (2025a): 889-942.

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen. "Canaries in the coal mine? six facts about the recent employment effects of artificial intelligence." Stanford Digital Economy Lab. Published August (2025b).

Ide, Enrique, and Eduard Talamàs. 2024. “Artificial Intelligence in the Knowledge Economy.” JPE Forthcoming.

Mousa, Deena. “The algorithm will see you now.” *Works in Progress Magazine, Sept 25, 2025.*