Discussion: "We Won't be Missed: Work and Growth in the Era of AGI" by Pascual Restrepo

Neil Thompson* Danial Lashkari*† Omeed Maghzian*

October 2025

In "We Won't be Missed: Work and Growth in the Era of AGI," Pascual Restrepo tackles an essential question: what will happen to the economy and workers as AI automation progresses? His answer: AGI can decouple economic growth from human labor, with growth in computational resources driving progress.

We commend Restrepo for defining AGI in economic terms, enabling a rigorous discussion about AGI's impacts. Paired with a parsimonious and general aggregate production function, he obtains striking predictions about a world with AGI and abundant computing resources. Restrepo also provides two important extensions. The first models the automation of science, which can be particularly important for growth (Romer 1990). The second examines how wages will adjust to AGI automation: smoothly if computing resources constrain AI adoption, abruptly if the development of better AI capabilities is the constraint. The paper contributes valuable insights about a possible — and radical — future.

Our comments are in four parts. We commend Restrepo's approach to modeling AGI, but note that (1) in practice it will likely be important to consider alternatives to the definition of AGI presented in the paper, and (2) the paper's production function has important limitations, particularly the assumption of constant returns to scale in production *and utility*, which rules out many important economic phenomena. For both, we propose directions for future work. We also (3) reevaluate Restrepo's conclusion that AGI will benefit workers, which we consider much less assured, and finally (4) we highlight that the assumption of exogenous increases in computing resources hides some exceptionally important economic dynamics on the transition path to AGI.

*MIT FutureTech. †Federal Reserve Bank of NY. All authors contributed equally. We thank Zachary Brown for excellent research assistance. The views presented here do not reflect those of the Federal Reserve Bank of NY or the Federal Reserve System.

AGI definition

One key contribution of the paper is its formalized notion of "AGI" in the context of a general economic model. For Restrepo, AGI occurs when we understand how to automate *all* work using a finite amount of compute. (This definition differs from more technical definitions of AGI, which focus on cognitive capabilities rather than economic performance (Morris et al. 2024).[1]) But, for some tasks, "compute-equivalent" performance is not perfectly substitutable for human performance. For example, it is hard to imagine that compute could replace the work done by athletes in the sports industry, because *human* struggle, achievement, and connection are essential to sports' cultural relevance and hence economic value. In other words, when "humanness" is a bottleneck characteristic in the utility derived from certain goods or services, it is difficult to conceptualize the AGI-equivalent of the corresponding work.

However, this point may not pose a serious problem, as even a relaxed notion of AGI would be sufficient for sustaining growth. If AGI is even imperfectly substitutable with human tasks for which the marginal product is bounded, then it can effectively achieve full automation of production without being able to perform all human work, as defined in the paper. To be more precise, consider the CRS aggregate production function $Y = F\left(\{X_t(\omega)\}_{\omega\in\Omega}\right)$ where output Y satisfies the following restriction (Kimball aggregator):

$$1 = \sum_{\omega\in\Omega} K\left(\frac{X(\omega)}{Y}\right),$$

and where $K(\cdot)$ is a monotonically increasing and concave function with $K(0) = 0$, $K(1) = 1$, and $K'(0) = b > 0$. In this environment, tasks ω are imperfectly substitutable and are all accessory tasks: we can indefinitely grow the economy by indefinitely growing any one task $\omega \in \Omega$ while all other tasks (and their marginal products) remain bounded.

This example suggests that we can use an alternative, less stringent economic definition of AGI, which is the following: AGI obtains if and only if there exists a subset Ω_B of tasks such that (i) all tasks in Ω_B can be performed with finite computational resources and (ii) F can grow indefinitely when all $X_t(\omega)$ for $\omega \in \Omega_B$ grow indefinitely. In other words, AGI may not require the automation of all productive tasks, but rather the automation of a subset of tasks that are sufficient for output to grow without bound. In light of the general structure of the theory, we believe the paper can relax the premise of what is considered to be AGI while maintaining most of the insights.

1. For an alternative economic definition, see OpenAI (2018).

The aggregate production function

The results from Restrepo's model are general in that they hold for a broad class of aggregate production functions: those exhibiting constant returns to scale (CRS) as well as the standard regularity conditions of monotonicity, concavity, and differentiability. But since the economy does not feature a balanced growth path and in the limit all income accrues to AI, we should ask whether standard growth assumptions are applicable to a world with AGI. For example, constant returns to scale only over labor and compute may be a strong assumption since bottlenecks from scarce factors that cannot be readily made from human work, such as rare earth minerals, may still limit growth regardless of AGI.[2]

It is also important to note that Restrepo's aggregator function F combines both the production technology *and the utility function:* it first maps tasks to the goods and services and then maps those goods and services to consumer welfare. Even if the assumption of constant returns to scale were reasonable for production, it may fail for consumption. Over the long run, there is ample evidence that rising income induces consumers to shift what they demand, e.g., towards healthcare, entertainment, education, etc., violating CRS (Comin, Lashkari, and Mestieri 2021). There are also other important economic phenomena that violate these assumptions: satiation points, positional goods, etc.

Another important feature the production function may not capture is self-fulfillment through work. Suppose, as some of the basic human needs become fulfilled, individuals shift their attention toward self-actualization in certain types of work, which lets them satisfy their desire for creative expression or constant learning (for evidence on this relationship, see Boar and Lashkari 2021). As a result, full automation would run counter to the fulfillment of these demands, unless AI can provide alternative tasks that substitute for this fulfillment.

Benefits to workers

The paper emphasizes that, in the limit, the transition from no automation to complete automation will increase productivity and make workers better off. However, it is unclear whether a transition from today's *partial* automation to full automation would be beneficial for workers. In particular, (Autor and Thompson 2025) point out that historically, when automation occurs disproportionately to workers' most expert tasks, wages fall. Society may also not function well if the wealth inequality between owners of AI

2. To recover the results of the paper, AI must not only be able to replace labor tasks, but also the use of other factors like materials, space, and energy inputs.

systems and workers translates to overwhelming political power for the shareholders of AGI firms. This could produce distributive mechanisms that diverge from workers being paid their marginal products (Restrepo's assumption) and could lead to a need for redistributive transfers to sustain human welfare in a world heading toward AGI.

Computing resources

Restrepo's paper is primarily focused on the end-state of AI automation, *conditional* on the premise of exogenously growing computational resources. But how much compute will it take to develop AGI? Will it be economically sustainable to produce such large amounts of computing? Will the mechanisms for providing computing resources *more efficiently* — hardware and algorithmic improvements — be durable enough? We present four relevant findings.

Some domains have large compute requirements. To apply AI to all problems, each will require a sufficient amount of computing. But domains are heterogenous, and this may be difficult in domains that experience only small improvements in AI performance from large compute investments. For example, the best AI score to date on ARC-AGI-2 (Chollet et al. 2025), a measure of an AI's ability to understand and apply rules it hasn't seen during training, is a mere 16%, compared to an average human performance of 62% (ARC Prize 2025), despite more than three orders-of-magnitude increase in models' cost-per-task on the benchmark. Since the returns to compute diminish strongly with scale (Kaplan et al. 2020; Gundlach, Lynch, and Thompson 2025; Thompson, Ge, and Manso 2022), it could require *many* orders-of-magnitude scale-up in compute to eclipse human capabilities in tasks like this one. This could greatly postpone AGI.

Growth in compute stock might not persist. Progress towards AGI and the economy Restrepo describes may be further impeded by slowing compute accumulation. While the stock of computing resources has historically grown exponentially, much of this has been driven by efficiency improvements that have made computing dramatically more cost-efficient (e.g., Moore's law) (Leiserson et al. 2020). However, as Figure 1 shows, improvements in chip performance-per-dollar have slowed in recent decades. While GPUs and other parallel-computing chips have maintained progress longer than CPUs through architectural and clock-speed trade-offs, they face the same fundamental limits. And, since current semiconductor fabrication is approaching atomic-scale limits, miniaturization will become increasingly costly and complex (Shalf and Leland 2015). Sustaining

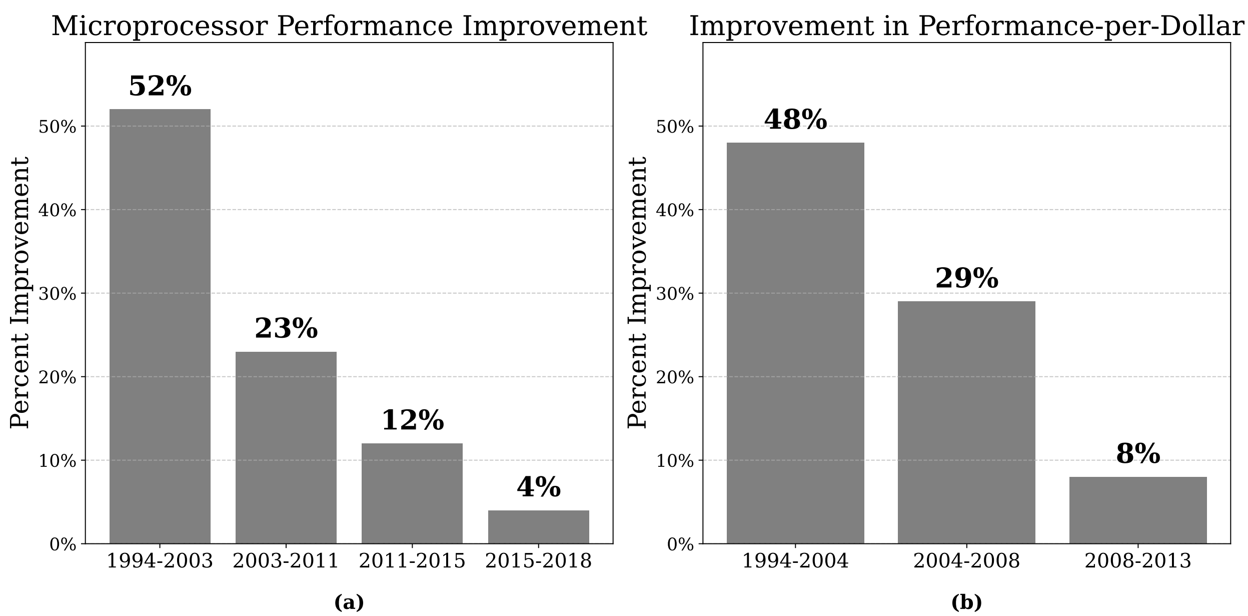

Figure 1: ***Rate of improvement in microprocessors,*** *as measured by (a) Annual performance improvement on the SPECint benchmark, and (b) Annual quality-adjusted price decline. Reproduced from Thompson and Spanuth 2021 based on data from Hennessy and Patterson (2019) and U.S. Bureau of Labor Statistics (Producer Price Index: Semiconductor and Related Device Manufacturing: Microprocessors (Including Microcontrollers)), although see Byrne, Oliner, and Sichel (2015) for a discussion of some of the difficulties in calculating the quality-adjusted price decline.*

exponential growth in compute will therefore require breakthroughs beyond incremental scaling of existing technology, which may not be realistic (Thompson and Spanuth 2021).

Algorithm progress can substitute for compute. The efficiency by which FLOPs (floating point operations) are converted into progress on tasks depends on the algorithm that is used. Innovation in algorithms can therefore compensate for insufficient compute capacity. Historically, algorithmic innovation has sometimes outstripped the pace of Moore's law. (Sherry and Thompson 2021).

Figure 2 summarizes how algorithm efficiency, hardware efficiency, and hardware scale-up have contributed to better language model performance (Ho et al. 2024). Of these sources, efficiency gains from algorithms or hardware are the most sustainable. But, as mentioned above, future hardware progress is likely to be limited. And algorithms themselves ultimately hit theoretical limits (Liu 2021). Absent these two sources of improvement, increased computing power would require proportional increases in cost.[3] This can quickly make technically feasible automation economically unattractive (as discussed below).

The transition to full automation may be slow. Even if AGI makes it technically feasible to automate human labor with compute, it may be economically unattractive to do so for a long period. This is because today many tasks would require orders-of-magnitude higher cost to perform with AI than with humans. AI's cost-per-task in these domains would therefore need to decline substantially before it is economically advantageous to invest in the compute required to automate them. But if there are technical barriers that mean cost-per-task declines slowly (for instance, slow rates of hardware and algorithm improvement, economic frictions, etc.), this could lead to the economy converging to an equilibrium in which compute remains bounded.

These barriers are important in practice: an analysis of tasks which are already technically feasible to automate with computer vision systems (Svanberg et al. 2024) estimates that it will still be cost-effective to employ human labor for decades, even if the costs of automation decrease by 50% annually.

3. While there would be some learning-by-doing or scale economies, such effects are too small to matter in this context.

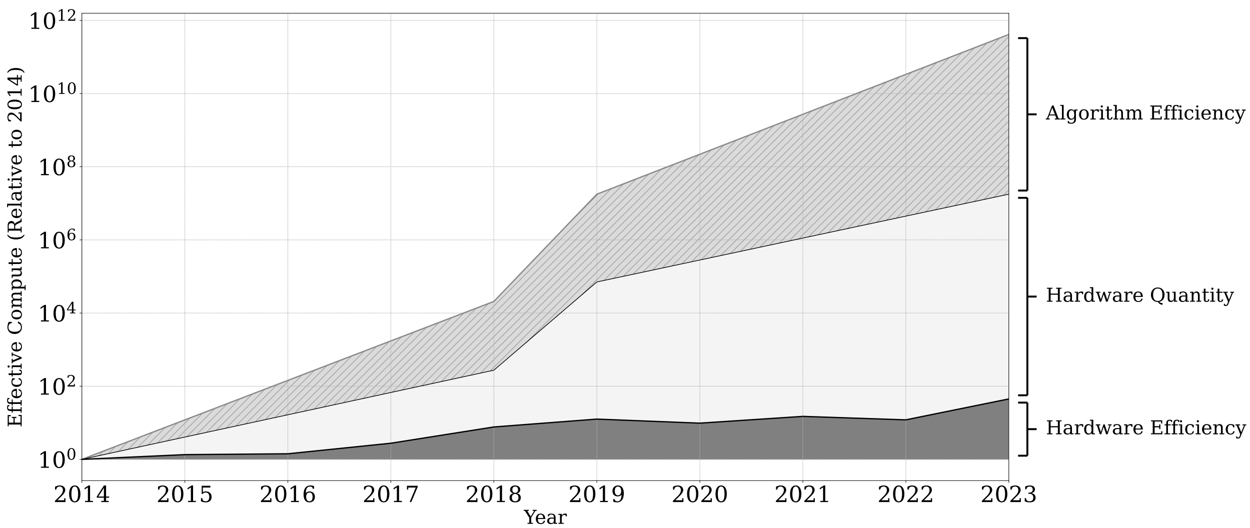

Figure 2: ***Contributions to improved language model performance.*** *Improvements in the efficiency of algorithms from 2014-2023 have been equivalent to a 22,000x scale-up in the compute used to train language models. Improvements in hardware efficiency (for instance, from Moore's law) accounted for a further 45x improvement, while most of the increase in AI performance was due to increased investment in compute. Adapted from Ho et al. 2024 with data on hardware efficiency from Del Sozzo et al. 2025.*

Implications for transition dynamics. If the growth in computing is slower than expected, then dynamics for how the economy converges towards AGI may be much more important than the (potentially long-delayed) final AGI state. Since the pace and character of this transition will depend heavily on the economics of computing resource provision, we strongly encourage future models to endogenize compute production.

In his discussion of the transition to AGI affects labor, Restrepo concludes that it matters whether the binding constraint is the growth in computing resources or new AI capabilities. If computing resources bind, the economic transition will exhibit smoother wage adjustments (proportional to the growth in computing resources) compared to the case where the path to AGI is constrained by progress in model capabilities.

As Restrepo highlights, these two paths are the extreme end of bottlenecks to reaching an AGI economy. However, the vast heterogeneity in computing requirements and algorithmic innovations across domains suggests that the expected transition path will be complicated, with some tasks being limited by capabilities while others are limited by computing resources. We encourage future work on labor dynamics to quantify the relative importance of each type of bottleneck across tasks.

Conclusion

Restrepo's paper provides an intuitive and valuable framework to analyze the changing relationship of labor to economic growth. There is a growing understanding that the future of work — and perhaps the position of humanity — will look substantially different as AI advances. We look forward to future work that applies and extends the line of thought developed in the paper, particularly relaxing the assumptions of constant returns to scale and exogenous computing growth.

References

ARC Prize. 2025. *ARC-AGI Leaderboard.* https://arcprize.org/leaderboard. Accessed: 2025-08-12.

Autor, David, and Neil Thompson. 2025. *Expertise.* Working Paper, Working Paper Series 33941. National Bureau of Economic Research. https://doi.org/10.3386/w33941. http://www.nber.org/papers/w33941.

Boar, Corina, and Danial Lashkari. 2021. *Occupational choice and the intergenerational mobility of welfare.* Technical report. National Bureau of Economic Research.

Byrne, David M., Stephen D. Oliner, and Daniel E. Sichel. 2015. *How Fast Are Semiconductor Prices Falling?* Working Paper, NBER Working Paper Series 21074. Cambridge, MA: National Bureau of Economic Research, April. https://doi.org/10.3386/w21074. http://www.nber.org/papers/w21074.

Chollet, Francois, Mike Knoop, Gregory Kamradt, Bryan Landers, and Henry Pinkard. 2025. *ARC-AGI-2: A New Challenge for Frontier AI Reasoning Systems.* arXiv: 2505.11831 [cs.AI]. https://arxiv.org/abs/2505.11831.

Comin, Diego, Danial Lashkari, and Martí Mestieri. 2021. "Structural change with long-run income and price effects." *Econometrica* 89 (1): 311–374.

Del Sozzo, Emanuele, Martin Fleming, Kenneth Flamm, and Neil Thompson. 2025. "How Much Progress Has There Been For GPUs?" Under review at ACM Economics and Computation.

Gundlach, Hans, Jayson Lynch, and Neil Thompson. 2025. *Meek Models Shall Inherit the Earth.* arXiv: 2507.07931 [cs.AI]. https://arxiv.org/abs/2507.07931.

Hennessy, John L., and David A. Patterson. 2017. "Domain-Specific Architectures." In *Computer Architecture: A Quantitative Approach,* 6th ed., 3. See Fig. 1.1. Cambridge, MA, USA: Morgan Kaufmann. ISBN: 978-0-12-811905-1.

Ho, Anson, Tamay Besiroglu, Ege Erdil, David Owen, Robi Rahman, Zifan Carl Guo, David Atkinson, Neil Thompson, and Jaime Sevilla. 2024. *Algorithmic progress in language models.* arXiv: 2403.05812 [cs.CL]. https://arxiv.org/abs/2403.05812.

Kaplan, Jared, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. *Scaling Laws for Neural Language Models.* arXiv: 2001.08361 [cs.LG]. https://arxiv.org/abs/2001.08361.

Leiserson, Charles E., Neil C. Thompson, Joel S. Emer, Bradley C. Kuszmaul, Butler W. Lampson, Daniel Sanchez, and Tao B. Schardl. 2020. "There's Plenty of Room at the Top: What Will Drive Computer Performance After Moore's Law?" *Science* 368 (6495): eaam9744. https://doi.org/10.1126/science.aam9744.

Liu, Emily. 2021. "A Metastudy of Algorithm Lower Bounds." Advisor: Neil Thompson. Master's thesis, Massachusetts Institute of Technology, September. https://hdl.handle.net/1721.1/140013.

Morris, Meredith Ringel, Jascha Sohl-dickstein, Noah Fiedel, Tris Warkentin, Allan Dafoe, Aleksandra Faust, Clement Farabet, and Shane Legg. 2024. *Levels of AGI for Operationalizing Progress on the Path to AGI.* arXiv: 2311.02462 [cs.AI]. https://arxiv.org/abs/2311.02462.

OpenAI. 2018. *OpenAI Charter.* https://openai.com/charter. Accessed September 30, 2025.

Romer, Paul M. 1990. "Endogenous Technological Change." *Journal of Political Economy* 98 (5): S71–S102. ISSN: 00223808, 1537534X, accessed August 13, 2025. http://www.jstor.org/stable/2937632.

Shalf, John M., and Robert Leland. 2015. "Computing beyond Moore's Law." *Computer* 48 (12): 14–23. https://doi.org/10.1109/MC.2015.374.

Sherry, Yash, and Neil C. Thompson. 2021. "How Fast Do Algorithms Improve? [Point of View]." *Proceedings of the IEEE* 109 (11): 1768–1777. https://doi.org/10.1109/JPROC.2021.3107219.

Svanberg, Maja, Wensu Li, Martin Fleming, Brian Goehring, and Neil Thompson. 2024. *Beyond AI Exposure: Which Tasks are Cost-Effective to Automate with Computer Vision?*

Thompson, Neil C., Shuning Ge, and Gabriel F. Manso. 2022. *The Importance of (Exponentially More) Computing Power.* arXiv: 2206.14007 [cs.AR]. https://arxiv.org/abs/2206.14007.

Thompson, Neil C., and Svenja Spanuth. 2021. "The decline of computers as a general purpose technology." *Communications of the ACM* (New York, NY, USA) 64 (3): 64–72. ISSN: 0001-0782. https://doi.org/10.1145/3430936.

U.S. Bureau of Labor Statistics. *Producer Price Index by Industry: Semiconductor and Related Device Manufacturing: Microprocessors (Including Microcontrollers) [PCU33441333441312].* Retrieved from FRED, Federal Reserve Bank of St. Louis. https://fred.stlouisfed.org/series/PCU33441333441312.