How Much Redistribution Will AI Require?

Alex Tabarrok

Department of Economics
George Mason University
Fairfax, VA 22030

Date: August 2026

Acknowledgments. The author thanks Tyler Cowen, Maxwell Tabarrok, and Kevin Bryan. The author also thanks ChatGPT, Claude, and Refine for error checking and productive comments. The author takes full responsibility for the accuracy and integrity of the work.

How Much Redistribution Will AI Require?

Date: August 2026

Abstract

Artificial intelligence may raise output while reducing labor's share of income. This paper derives an accounting rule for the transfer required to keep aggregate real labor income per capita on its no-AI path. The required share of AI-path GDP depends on relative GDP and labor's share. Faster growth permits a lower labor share without reducing aggregate labor income. Suppose GDP would otherwise grow by 2 percent annually, with labor receiving 60 percent. No transfer is needed if AI raises annual growth to 5 percent and labor's share remains at least 45 percent. At 10 percent growth, labor's share can fall to 28 percent and labor as a whole would remain as well off as under the no-AI path. A three-parameter translator connects beliefs about displacement, productivity, and labor's share of new output to this forecast space. Across three stylized views, including one in which labor's share falls to 23 percent, the required transfer remains below 6 percent of GDP. At that scale, making labor recipients whole while preserving gains for capital recipients appears politically feasible.

JEL codes: O33, E25, H23.

1 Introduction

Artificial intelligence can raise output while reducing labor's share of income. The gains may be large enough to compensate labor recipients, but the required transfer depends on how much real GDP per capita rises and how far labor's share falls.

This paper asks what share of AI-path GDP must be transferred to keep aggregate real labor income per capita on a benchmark no-AI path. Let y denote real GDP per capita and s_L labor's share of GDP. Labor income per capita is $s_L y$. The benchmark assumes that, without AI, labor income remains equal to 60 percent of GDP. Given forecasts for GDP on the AI and no-AI paths and a scenario for labor's share under AI, national-income accounting determines the required transfer. The calculation applies regardless of the model used to produce the forecasts.

Faster growth permits a lower labor share without reducing aggregate labor income. Suppose GDP would otherwise grow by 2 percent annually, with labor receiving 60 percent. No transfer is needed if AI raises annual growth to 2.5 percent and labor's share remains at least 57 percent. At 5 percent growth, labor's share can fall to 45 percent. At 10 percent growth, it can fall to 28 percent and labor as a whole would remain as well off as under the no-AI path. Large transfers arise when labor's share falls sharply while GDP growth remains near baseline. The calculation expresses the required transfer as a share of GDP. A corresponding rate on nonlabor income can be calculated separately.

2 A task-model example

The task-based framework of Acemoglu and Restrepo (2018), in the stripped-down form presented by Hassanein and Steinsson (2026), gives one interpretation of these outcomes. Capital may take over tasks previously performed by labor, reducing labor demand and labor's share. Machines may also become more productive in tasks already assigned to capital. When tasks are complementary, these gains raise output and increase the value of the tasks that remain with labor.

The low-growth, low-labor-share region studied below corresponds to task displacement that runs ahead of machine-productivity growth. Faster machine improvement, firms adding capital, and the creation of new human tasks support higher output and a larger labor share. The Appendix verifies that the forecast grid can be generated within this model while capital and labor input per capita adjust.

3 The accounting rule

Fix a horizon T. Let $\bar{y}_T$ denote real GDP per capita on the no-AI path and y_T real GDP per capita on the AI path.[1] Define

$$G = \frac{y_T}{\bar{y}_T}.$$

Let λ_0 denote labor's benchmark share of GDP on the no-AI path and s_T labor's share on the AI path. Benchmark labor income per capita is $\lambda_0 \bar{y}_T$, while realized labor income per capita is $s_T y_T$. The transfer required to close any shortfall, expressed as a share of AI-path GDP, is

$$\boxed{x_T^*(G, s_T) = \max\left\{0, \frac{\lambda_0}{G} - s_T\right\}.} \tag{1}$$

The transfer equals zero whenever $s_T G \geq \lambda_0$. Results scale linearly with the benchmark λ_0.

For comparison with a tax on nonlabor income, the corresponding share of gross nonlabor income is

$$\tau_T^* = \frac{x_T^*}{1 - s_T}. \tag{2}$$

For $G > 1$ and a positive transfer, labor recipients receive $\lambda_0 \bar{y}_T$ after the transfer. Nonlabor recipients retain

$$y_T - \lambda_0 \bar{y}_T = (G - \lambda_0)\bar{y}_T,$$

compared with $(1 - \lambda_0)\bar{y}_T$ on the no-AI path. Their aggregate gain is $(G - 1)\bar{y}_T$. Interpreting labor and nonlabor recipients as aggregate groups, equation (1) yields a Pareto improvement over the no-AI path: labor recipients retain their benchmark income, while nonlabor recipients gain whenever $G > 1$. Individual household outcomes may differ.

[1] I assume population follows the same path under AI and no AI. Preserving labor income per capita is then equivalent to preserving total aggregate labor income.

4 The benchmark

Equation (1) compares the AI outcome with a simple no-AI path. The central benchmark combines 2 percent annual growth in real GDP per capita with a constant labor share of $\lambda_0 = 0.60$ (International Labour Organization, 2023). Benchmark labor income therefore grows at the same rate as GDP on the no-AI path. Substituting another labor share changes every break-even threshold through λ_0/G.

Over ten years, real GDP per capita and benchmark labor income per capita both rise by $(1.02)^{10} - 1 = 21.9$ percent. The transfer protects this rising labor-income path. Faster AI-path growth increases total income, while a lower labor share reduces the portion paid to labor. Their product determines whether a transfer is needed.

5 AI growth scenarios

Let the horizon be $T = 10$ years. Annual growth in real GDP per capita on the no-AI path is $\bar{g} = 0.02$. Let g denote projected annual growth under AI. The terminal GDP ratio is

$$G = \left(\frac{1+g}{1+\bar{g}}\right)^T. \tag{3}$$

The analysis considers annual AI-path growth rates of $g = 2.5$, 5, and 10 percent. After ten years, these rates place real GDP per capita 5.0, 33.6, and 112.8 percent above the no-AI path. A larger value of G allows labor's share to fall further before labor income drops below the benchmark.

Prominent industry forecasts motivate the upper scenarios. Nadella (2025) uses annual GDP growth of 5–10 percent as his test for an AI industrial revolution. Amodei (2026) discusses growth of 10–20 percent.

5.1 Break-even labor share

Set AI-path labor income relative to no-AI GDP, $s_T G$, equal to the benchmark λ_0. The resulting break-even terminal labor share is

$$\boxed{s_T^{\mathrm{BE}} = \frac{\lambda_0}{G} = \lambda_0 \left(\frac{1+\bar{g}}{1+g}\right)^T.} \tag{4}$$

Table 1: Growth in real GDP per capita and the break-even labor share

Annual growth in real GDP per capita under AI g	GDP per capita relative to no-AI path G	Break-even terminal labor share
2.5%	1.050	57.1%
5.0%	1.336	44.9%
10.0%	2.128	28.2%

Table 1 traces the break-even curve. At 5 percent growth, labor can receive 44.9 percent of

GDP and still reach its benchmark income level.

Table 2 applies the break-even thresholds to selected labor-share outcomes. Each cell reports the required transfer as a share of AI-path GDP. A zero indicates that no transfer is needed.

5.2 Selected transfer rates

Table 2: Required year-10 transfer as a share of AI-path GDP

Terminal labor share	$g = 2.5\%$	$g = 5\%$	$g = 10\%$
45%	12.1%	0.0%	0.0%
35%	22.1%	9.9%	0.0%
25%	32.1%	19.9%	3.2%
15%	42.1%	29.9%	13.2%

Notes: Since $s_T^{\text{BE}} = \lambda_0/G$, the required GDP share is simply the gap below the break-even labor share: $x_T^* = \max\{0, s_T^{\text{BE}} - s_T\}$.

5.3 The growth-share map

Figure 1 maps each combination of AI-path growth and labor's terminal share into the transfer required as a share of AI-path GDP. The dashed line marks combinations that leave aggregate labor income on its no-AI path. The white region above the line requires no transfer.

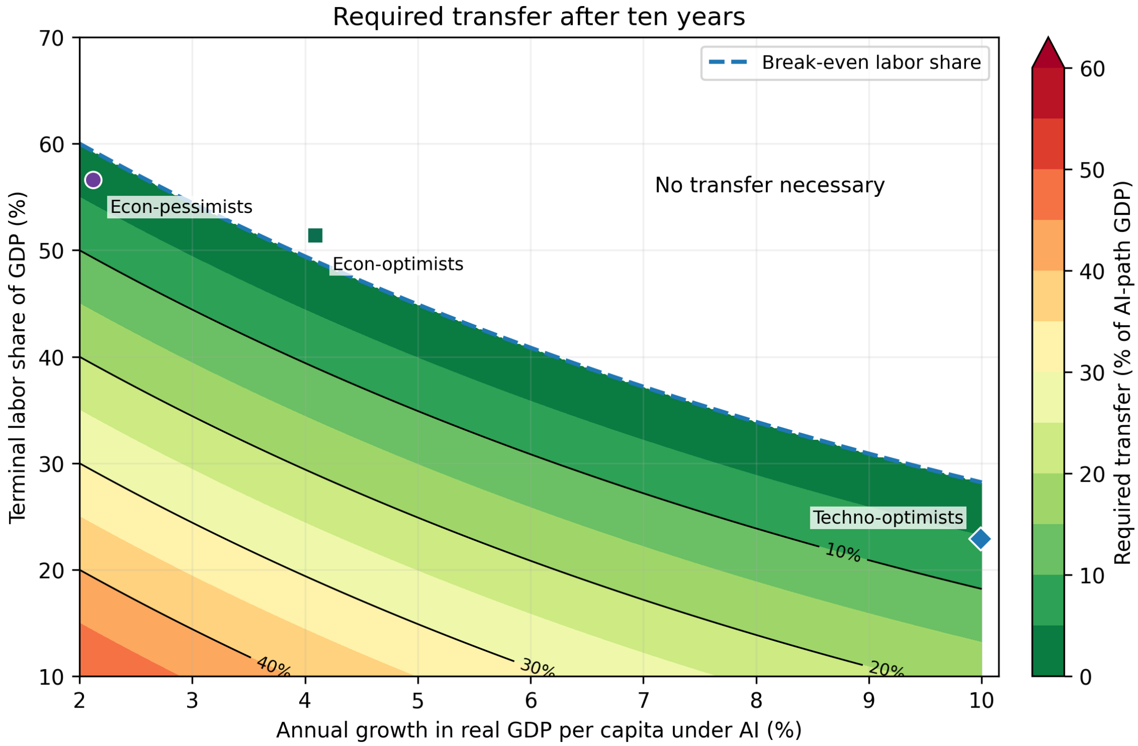

Figure 1: Required transfer as a share of AI-path GDP after ten years

Notes: The figure applies equation (1) using $T = 10$, conventional annual compounding, and $\lambda_0 = 0.60$. The horizontal axis gives annual growth in real GDP per capita under AI; the no-AI benchmark is 2 percent. The dashed line is the break-even labor share. The white region has a zero transfer. Black contours mark transfers equal to 10, 20, 30, 40, and 50 percent of AI-path GDP. The three points are the stylized views in Table 3.

The downward-sloping break-even line shows that higher AI-path growth is compatible with a lower labor share while leaving labor as a whole at least as well off as under the no-AI path.

5.4 A reduced-form scenario translator

Which scenarios are most likely? A reader can place a forecast on Figure 1 with the parameter vector $\theta = (d, q, \beta)$, where d measures displaced benchmark labor income, q measures GDP created per displaced dollar, and β assigns part of that new GDP to labor. The resulting GDP ratio and terminal labor share are

$$G = 1 + q\lambda_0 d, \qquad s_T = \frac{\lambda_0(1-d) + \beta q \lambda_0 d}{G}. \tag{5}$$

Equation (3) converts G into annual growth g, and equation (1) gives the required transfer. The product βq is the new labor income generated per dollar of benchmark labor income displaced. For any $d > 0$, no transfer is required if and only if $\beta q \geq 1$. More displacement moves a scenario southeast when $\beta q < 1$. Higher q moves it rightward, while higher β moves it upward.

Drawing broadly on the literature, we define three stylized views: the econ-pessimist, the econ-optimist, and the techno-optimist. Figure 1 maps them onto the growth–labor-share space, while Table 3 reports their parameter values and year-10 outcomes.

The econ-pessimist view uses Acemoglu's estimate that 4.6 percent of tasks may be profitably affected by AI within ten years and that the resulting GDP gain is likely to be about 0.9–1.6 percent (Acemoglu, 2025). Weak machine productivity limits profitable aggregate displacement, leaving labor's share near its benchmark even when particular industries face substantial disruption. Its low value of β also captures the concern that current AI development emphasizes displacement rather than the creation of good new tasks.

The econ-optimist view treats some direct task displacement as a source of complementarity and new work. This interpretation draws on Cowen (2026)'s rejection of lasting mass unemployment, McCloskey (2017)'s historical argument against technological unemployment, and Bryan (2026)'s emphasis on dynamically acquired human skills. The calibration produces 4.1 percent annual growth, with income gains for labor and nonlabor recipients.

The techno-optimist view combines very high displacement with very high productivity. Amodei (2026) argues that sustained annual GDP growth of 10–20 percent may be possible and warns that AI could displace half of entry-level white-collar jobs within one to five years. Altman (2021) similarly anticipates exceptional abundance alongside a shift in economic power from labor toward capital.

Table 3: Three stylized views mapped into year-10 outcomes

Stylized view	d	q	β	g	s_T	x_T^*
Econ-pessimists	0.05	0.40	0.25	2.1%	56.6%	2.7%
Econ-optimists	0.25	1.50	0.80	4.1%	51.4%	0.0%
Techno-optimists	0.75	2.50	0.30	10.0%	22.9%	5.3%

Notes: Calculations use $T = 10$, $\bar{g} = 0.02$, and $\lambda_0 = 0.60$. The three values of βq are 0.10, 1.20, and 0.75. The labels identify broad mechanisms in the cited work rather than joint views or endorsed numerical forecasts.

Surprisingly, all three views imply transfer shares below 6 percent of GDP. The econ-optimist calibration raises labor income from 60 to 63 percent of no-AI GDP and nonlabor income from 40 to 59.5 percent, so labor and capital both gain without a transfer. Weak productivity in the econ-pessimist calibration limits aggregate displacement; labor's share remains 56.6 percent and the transfer is 2.7 percent, even though disruption may be severe in exposed industries. The techno-optimist calibration takes the opposite route. Labor's share falls to 22.9 percent, while GDP reaches 2.125 times its no-AI level and holds the transfer to 5.3 percent of AI-path GDP.

6 Discussion

Equation (1) is deliberately a year-10 benchmark. It maps alternative year-10 outcomes for GDP per capita and labor's share into the transfer required at that date. Its purpose is to provide a back-of-the-envelope measure of whether the resulting compensation problem is modest and politically and financially feasible, or severe and potentially intractable under normal political conditions.

The required transfer need not imply an equal increase in total taxation. The relevant fiscal question is whether the transfer is smaller than the revenue already raised from taxes on labor. When it is, compensation could in principle be financed by a revenue-neutral shift in the tax base rather than by a net increase in government revenue. The accounting calculation does not determine the appropriate mix of tax reductions and cash transfers; it measures the aggregate resources that must be redirected toward labor recipients.

Compensation could come from broad capital ownership, a tax on capital income or economic rents, profit sharing, a social wealth fund, or general revenue.[2] These instruments differ in incidence, administrative feasibility, and effects on investment.

One possible implementation is a broad consumption tax combined with lower or zero direct taxes on wages and with cash transfers to workers who receive little labor income. A consumption tax still burdens wage income when it is spent, but it also taxes consumption financed by nonlabor income and accumulated wealth. Because the base is broader, it can replace a given amount of labor-tax revenue at a lower effective rate on wage-financed consumption. The wage-tax reduction benefits continuing workers, while the transfers are necessary to reach workers whose employment or earnings have disappeared.[3]

In the more moderate scenarios in Table 2, the required transfer may plausibly be within the scale of existing labor-tax revenue. In scenarios where AI replaces many tasks, the required GDP share is much larger, and additional revenue, lower spending elsewhere, or a less complete compensation target may be required.

Early task takeover can move the economy south in Figure 1 by reducing labor's share before the full productivity gains from automation appear. As automation spreads, machine productivity compounds across a larger base, firms add capital, and new human tasks become more valuable. Meanwhile, diminishing returns make further task displacement harder once the low-hanging opportunities have been exhausted. Labor's share may therefore fall early in the transition and recover later. The timing depends on the relative pace of task-frontier expansion and machine-productivity growth (Hassanein and Steinsson, 2026).

Radiology illustrates the more favorable mechanism. Image-recognition AI was widely expected to displace radiologists, but deployment has so far concentrated on narrower tasks such as triage, measurement, and report drafting, while radiologists remain responsible for integrating clinical information, difficult cases, communication, procedures, and final sign-off. In a prospective clinical study, AI-generated draft radiograph reports improved documentation efficiency by 15.5 percent without reducing peer-reviewed accuracy or textual quality (Huang et al., 2025). At the same time, imaging volumes continue to rise and the specialty faces persistent workforce shortages (Rula, 2026). This is movement toward the northeast of the figure: AI raises the productivity of a scarce complementary professional and expands service output rather than simply eliminating the

[2]The nonlabor GDP residual is not itself an immediately usable statutory tax base. It includes depreciation and other components that do not accrue as disposable owner income. Consumption of fixed capital was about 16.4 percent of GDP in 2025 (BEA, 2026). If depreciation is d as a share of GDP, translating a required GDP share x_T^* into a rate on net nonlabor income gives approximately $\tau_{\text{net}} = x_T^*/(1 - s_T - d)$. At a 25 percent terminal labor share and $d = 0.164$, a transfer equal to 20 percent of GDP corresponds to roughly 34 percent of net nonlabor income.

[3]If B_L is consumption financed by labor resources and B_N consumption financed by nonlabor resources, a wage tax t_L raising $t_L B_L$ can be replaced by a tax-inclusive consumption tax $t_C = t_L B_L/(B_L + B_N) < t_L$ whenever $B_N > 0$.

occupation. The same logic would apply in software engineering if AI removes boilerplate and routine debugging while engineers use the saved time to build more products, integrate systems, and perform higher-level design.

7 Conclusion

AI could bring tremendous economic change. Rapid automation may reshape jobs throughout the economy, raise output at unprecedented rates, and shift a large share of income from labor to capital. Surprisingly, the transfers needed to keep labor as a whole on its no-AI path while preserving gains for capital are modest across the scenarios considered here. They appear politically feasible.

The reason is growth. Scenarios with the largest declines in labor's share also tend to have the highest GDP growth rates. A smaller labor share of a much larger economy can leave aggregate labor income at or above its benchmark path. Redistribution then shares part of AI's surplus from a growing pie.

With a 60 percent benchmark labor share and 2 percent annual no-AI growth, 5 percent AI-path growth lets labor's share fall to about 45 percent before any transfer is required. At 10 percent growth, the threshold is about 28 percent. Across the three stylized views, every required transfer is below 6 percent of AI-path GDP. Such a transfer keeps labor recipients whole as a group while leaving nonlabor recipients better off than under the no-AI path. It produces an aggregate Pareto improvement between labor and capital.

Policy can adjust as realized GDP and labor's share become observable. The reduced-form translator identifies the key uncertainty: whether productivity and labor's share of new output generate at least one dollar of new labor income per dollar displaced, $\beta q \geq 1$. When they do, labor gains without a transfer. When they fall short, equation (1) measures the fiscal adjustment needed to convert AI's aggregate gains into a Pareto improvement between labor and capital.

Appendix

This appendix checks that the forecast grid can be generated by one illustrative parameterization of the stripped-down task model while capital and labor input per capita adjust. With constant returns to scale, write output, capital, and labor hours per capita as y, k, and h. The model reduces to

$$y = \left[(A_K k)^{\frac{\sigma-1}{\sigma}} + (A_L h)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}, \qquad A_K = \alpha^{\frac{1}{\sigma-1}}\psi_K, \quad A_L = (1-\alpha)^{\frac{1}{\sigma-1}}\psi_L. \tag{6}$$

Let barred variables denote the no-AI terminal path and define

$$G = \frac{y_T}{\bar{y}_T}, \qquad \kappa = \frac{k_T}{\bar{k}_T}, \qquad \lambda = \frac{h_T}{\bar{h}_T}. \tag{7}$$

Holding labor's within-task productivity fixed relative to the no-AI path, the forecast pair (G, s_T) implies

$$1 - \alpha_T = (1 - \bar{\alpha}_T) G^{\sigma-1} \left(\frac{s_T}{\lambda_0}\right)^{\sigma} \lambda^{1-\sigma}, \tag{8}$$

and

$$\frac{\psi_{K,T}}{\bar{\psi}_{K,T}} = \frac{G}{\kappa} \left(\frac{1-s_T}{1-\lambda_0}\right)^{\frac{\sigma}{\sigma-1}} \left(\frac{\alpha_T}{\bar{\alpha}_T}\right)^{-\frac{1}{\sigma-1}}. \tag{9}$$

The first equation maps the forecast into task displacement conditional on the change in labor input. The forecast identifies the combined contribution of machines becoming more productive and firms adding capital. I allow labor hours to fall by as much as 30 percent and vary the capital response:

$$\lambda \in \{0.70, 0.85, 1\}, \qquad \kappa = G^{\chi}, \qquad \chi \in \{0.5, 1, 1.25\}. \tag{10}$$

Here $\lambda = 0.70$ allows labor hours per capita to fall 30 percent relative to the no-AI path. The case $\chi = 1$ holds the capital-output ratio constant. With $\chi < 1$, firms add capital more slowly than output rises; with $\chi > 1$, they add capital faster.

Using $\sigma = 0.5$, $\bar{\alpha}_T = 0.5$, $\lambda_0 = 0.60$, and the 2 percent no-AI growth benchmark, the range plotted in Figure 1, with annual AI-path growth from 2 to 10 percent and terminal labor shares from 10 to 70 percent, implies approximately

$$\alpha_T \in [0.460, 0.883]. \tag{11}$$

Machine productivity relative to the no-AI path lies between 0.93 and 3.71 when $\chi = 0.5$, between 0.93 and 2.54 when $\chi = 1$, and between 0.93 and 2.10 when $\chi = 1.25$. The twelve points reported in Table 2 are comfortably interior, with $\alpha_T \in [0.577, 0.857]$. Depending on labor and capital adjustment, their implied machine-productivity ratios range from 0.99 to 2.40 when $\chi = 0.5$, from 0.97 to 1.64 when $\chi = 1$, and from 0.96 to 1.36 when $\chi = 1.25$. In the high-transfer combinations, AI replaces many tasks without commensurate GDP gains. At the same labor shares, high-growth combinations require substantially greater machine productivity or capital deepening.

The calculation places the forecast grid inside one task-model calibration.

References

Acemoglu, Daron, and Pascual Restrepo. 2018. “The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment.” *American Economic Review* 108(6): 1488–1542.

Acemoglu, Daron. 2025. “The Simple Macroeconomics of AI.” *Economic Policy* 40(121): 13–58.

Altman, Sam. 2021. “Moore’s Law for Everything.” March 16. `https://moores.samaltman.com/`.

Amodei, Dario. 2026. “The Adolescence of Technology.” `https://darioamodei.com/essay/the-adolescence-of-technology`.

U.S. Bureau of Economic Analysis. 2026. “GDP (Advance Estimate), Fourth Quarter and Year 2025.” BEA 26–11, February 20.

Bryan, Kevin A. 2026. “The Economic Impacts of Artificial Intelligence: A Multidisciplinary, Multi-Book Review.” *Journal of Economic Literature* 64(1): 281–300.

Cowen, Tyler. 2026. “Some Non-Obvious Reasons Why AI Will Create Some Transitional Problems in Employment.” *Marginal Revolution*, May 13. Online post.

Huang, Jonathan, Matthew T. Wittbrodt, Caitlin N. Teague, et al. 2025. “Efficiency and Quality of Generative AI–Assisted Radiograph Reporting.” *JAMA Network Open* 8(6): e2513921.

Hassanein, Abdelrahman, and Jon Steinsson. 2026. “A Simple Introduction to Task-Based Models of Technical Change.” Working paper.

International Labour Organization. 2023. “10.4.1: Labour Share of GDP (%), SL_EMP_GTOTL.” Dataset, *United Nations Sustainable Development Goals 2023-Q2*. `https://ourworldindata.org/grapher/labor-share-of-gdp`.

McCloskey, Deirdre N. 2017. “The Myth of Technological Unemployment.” *Reason*, August/September. `https://reason.com/2017/07/11/the-myth-of-technological-unem/`.

Nadella, Satya. 2025. Interview by Dwarkesh Patel, “Microsoft’s AGI Plan and Quantum Breakthrough.” February 19. `https://www.dwarkesh.com/p/satya-nadella`.

Rula, Elizabeth Y. 2026. “The Radiologist Shortage: A Workforce Update from HPI.” *ACR Bulletin*, February 5.