Artificial Intelligence, Mass Unemployment, and the Optimal Financing of Universal Basic Income

Frank Schmid

frank.schmid@frankschmid.com

This version: 8 November 2026

First version: 17 April 2026

Abstract. This paper develops a dynamic general-equilibrium model of an economy in which artificial intelligence (AI) capital substitutes for labor at the task level, and characterizes the tax structure that finances a universal basic income (UBI) along the transition. Production combines AI capital, non-AI capital, and heterogeneous labor in a nested constant-elasticity-of-substitution (CES) technology. AI capital is supplied against scarce complementary inputs, so its supply slopes upward and part of its return is a quasi-rent. The paper states conditions under which rising AI intensity lowers the wage and raises non-participation, and it formulates the planner's problem as a multi-instrument Ramsey problem with an endogenous skill distribution and an insurance role for UBI against skill devaluation. In the computed optimum the AI capital tax falls on the quasi-rent, the normal return to AI capital is untaxed, and the labor income tax is progressive with a zero rate at the participation margin. These results are stated for an instrument set that excludes subsidies: where an earnings subsidy at the participation margin is admitted, the computed optima use it and non-participation is lower in every year. Along the computed transition the effective AI tax rises with AI intensity and lies below its static optimum in the early years, because the planner values faster adoption. With mobile AI capital, uncoordinated taxation sets the rate below the cooperative one, so the financing structure depends on coordination across jurisdictions.

JEL classification: E24, H21, H23, J24, O33, O41.

This paper was developed with assistance from generative AI tools across the design, analysis, coding, and drafting stages. The author supervised all stages, verified the results, and takes full responsibility for the content.

Table of Contents

Abbreviations and Acronyms . . . . . **iv**

1 Introduction **1**

2 Related Literature **2**

3 A Task-Based Foundation for AI Production **4**
3.1 Tasks, Assignment, and Aggregation 5
3.2 From Task-Level to Aggregate Substitution 5
3.3 Implications for Reinstatement 6

4 The Baseline Economy **6**
4.1 A Three-Factor Production Structure 6
4.2 Heterogeneous Labor and the Employment Margin 7
4.3 AI Capital Supply with Scarce Complementary Inputs 8
4.4 Ownership, Wealth, and Distribution 9
4.5 Equilibrium 9
4.6 General-Equilibrium Comparative Statics of AI Adoption 10
4.7 Conditions for Mass Unemployment 11

5 Risk, Insurance, and the Welfare Role of UBI **12**
5.1 Idiosyncratic Skill-Devaluation Risk 13
5.2 The Insurance Value of UBI 13
5.3 Implications for the Welfare Objective 14

6 The Optimal Tax Structure in Steady State **14**
6.1 The Planner's Problem 14
6.2 The Multi-Instrument Ramsey System 15
6.3 Atkinson–Stiglitz and the Role of the VAT 16
6.4 Pigouvian Correction with Endogenous Externality 16
6.5 Mirrlees Progressivity and the Participation Margin 17
6.6 The Optimal Structure in Steady State 18

7 The Dynamic Transition Path **19**
7.1 The Dynamic Environment 19
7.2 Endogenous Skill Formation 20
7.3 The Planner's Dynamic Problem 20
7.4 First-Order and Co-State Conditions 21
7.5 The Deceleration Component as a Congestion Externality 22
7.6 The Four Phases of Optimal Policy 22
7.7 A Calibrated Transition Path 23
7.8 Sensitivity to Four Parameters 26
7.9 An Earnings Subsidy at the Participation Margin 28
7.10 Decomposing the Welfare Gain from Optimal Policy 28
7.11 A Comparative-Statics Summary 29
7.12 Hysteresis and Irreversibility 29

8 The International Dimension **30**

8.1 A Two-Country Model . . . 30
8.2 Nash Equilibrium in AI Capital Taxes . . . 30
8.3 Destination-Based Taxation and Coordination . . . 31

9 Implementation Challenges . . . 31
9.1 Measurement and Definition . . . 32
9.2 Political Economy . . . 32
9.3 Welfare Beyond Income . . . 32
9.4 Endogenous Technical Change . . . 32
9.5 Demographic Heterogeneity . . . 33
9.6 Scope and Limitations of the Framework . . . 33

10 Conclusion . . . 34

Bibliography . . . 36

Appendix A Derivation of the Aggregate CES from Task-Level Primitives . . . 38

Appendix B Notation Summary . . . 38
B.1 Factors and Production . . . 38
B.2 Tasks . . . 39
B.3 Labor, Skills, and Unemployment . . . 39
B.4 Policy Instruments and Welfare . . . 39
B.5 Ramsey System . . . 40
B.6 Dynamic Environment . . . 40
B.7 Risk and Insurance . . . 40
B.8 International Dimension . . . 40

Appendix C Numerical Solution and Calibration . . . 41
C.1 Method . . . 41
C.2 Welfare and Instruments . . . 41
C.3 Accuracy of the Solution . . . 41
C.4 Rent Base . . . 42
C.5 Adoption Law . . . 42
C.6 Quantities Fixed Outside the Stated Calibration . . . 42
C.7 Earnings-Subsidy Variant . . . 43
C.8 Two-Country Solution . . . 43

List of Tables and Figures

Table 1. Computed Optimal Tax Structure by AI Intensity 19
Table 2. Computed Optimal Tax Path 25
Table 3. Comparative-Statics Summary: Sign of Endogenous Responses 29
Table C.1. Quantities the Model Fixes Outside Its Stated Calibration 43
Figure 1. Computed Optimal Transition Path, Years 0 through 30 26
Figure 2. Sensitivity of the Computed Path at Year 30 to Four Parameters 27

Abbreviations and Acronyms

AI	Artificial intelligence
CE	Certainty equivalent
CES	Constant elasticity of substitution
CRRA	Constant relative risk aversion
CV	Coefficient of variation
GDP	Gross domestic product
JEL	Journal of Economic Literature
MSC	Marginal social cost
OECD	Organisation for Economic Co-operation and Development
OLG	Overlapping generations
R&D	Research and development
TFP	Total factor productivity
UBI	Universal basic income
VAT	Value-added tax

Greek Letters

Each entry lists the letter, its name, and its pronunciation.

α	alpha (AL-fah)
γ, Γ	gamma (GAM-ah)
δ, Δ	delta (DEL-tah)
ε	epsilon (EP-sih-lon)
ζ	zeta (ZAY-tah)
η	eta (AY-tah)
θ, Θ	theta (THAY-tah)
ι	iota (eye-OH-tah)
κ	kappa (KAP-ah)
λ	lambda (LAM-dah)
μ	mu (MYOO)
ν	nu (NOO)
ξ	xi (ZYE)
π, Π	pi (PYE)
ρ, ϱ	rho (ROH); ϱ is the variant form
σ, ς	sigma (SIG-mah); ς is the variant form
τ	tau (TAW)
Υ	upsilon (UP-sih-lon)
φ	phi (FYE)
χ	chi (KYE)
ψ, Ψ	psi (SYE)
ω	omega (oh-MAY-gah)

1 Introduction

The accelerating deployment of artificial intelligence (AI) is a candidate structural change in the economy with potentially large distributional consequences. Unlike previous waves of automation, which primarily substituted for physical labor, contemporary AI systems substitute for cognitive labor across a range of occupations. This paper develops a formal theoretical framework for analyzing the conditions under which AI adoption generates mass unemployment and for designing optimal policy responses centered on universal basic income (UBI) financed through taxation.

Our contribution is fivefold. First, we provide a task-based microfoundation for the aggregate substitution elasticity between AI and labor, so that reinstatement becomes a modeled object rather than an assumption, and the aggregate elasticity itself rises endogenously as AI's reach across the task space expands. Second, we work with a three-factor production technology (AI capital, non-AI capital, effective labor) in which the elasticity between AI capital and labor differs from the elasticity between non-AI capital and labor; this gives the paper's proposed tax instruments content by making AI capital a structurally distinct factor. Third, we replace the conventional assumption of perfectly elastic capital supply with an upward-sloping supply curve arising from scarce complementary inputs, which allows a share of AI capital returns to take the form of quasi-rents and alters the incidence of capital taxation. Fourth, we formulate the planner's problem as a proper multi-instrument Ramsey problem with cross-elasticities, derive the first-order and co-state conditions of the dynamic Hamiltonian explicitly, and embed UBI as an instrument whose welfare gain includes an insurance component against uninsurable idiosyncratic skill-devaluation risk. Fifth, we formalize the international dimension as a two-country Nash game in AI capital taxes, showing how capital mobility shifts the feasible policy frontier and how destination-based taxation restores a share of the welfare loss.

The central finding is that the optimal tax structure is time-varying and state-contingent. The tax on AI capital falls on the quasi-rent to scarce complementary inputs, and its effective rate rises as AI intensity rises, because the rent share of the AI return grows as demand for AI capital rises against an inelastic supply. Along the computed transition the rate lies below its static optimum in the early years: adoption responds to the tax, and the planner values faster adoption at the margin. The labor income tax is progressive across terciles, with a zero rate at the bottom because participation responds at the extensive margin there. The system is self-stabilizing in the sense that rising AI displacement generates both the fiscal requirement for UBI and the tax base to finance it.

The paper is organized as follows. Section 2 reviews the relevant literature on optimal taxation, technological unemployment, and UBI. Section 3 develops the task-based foundation. Section 4 develops the baseline three-factor economy, derives the general-equilibrium comparative statics of AI adoption, and states the conditions for mass unemployment. Section 5 introduces idiosyncratic skill-devaluation risk and characterizes the insurance value of UBI. Section 6 derives the optimal static tax structure as a multi-instrument Ramsey problem with Pigouvian, Mirrlees, and Atkinson–Stiglitz components. Section 7, the central contribution, develops the dynamic transition path, writes out the Hamiltonian, derives the first-order and co-state conditions, motivates the deceleration term as a congestion externality, and presents a calibrated transition path. Section 8 formalizes the international dimension. Section 9 discusses measurement, political economy, and related implementation challenges. Section 10 concludes.

2 Related Literature

This paper synthesizes four strands of literature. The first is the optimal taxation literature following Ramsey, Mirrlees, Diamond and Mirrlees, Atkinson and Stiglitz, Saez, Diamond and Saez, and Piketty, Saez, and Stantcheva.[1] This tradition provides the analytical apparatus for characterizing optimal tax rates as functions of base elasticities and distributional weights, and for reconciling differentiated consumption taxation with a progressive nonlinear income tax. We extend this framework by incorporating a large exogenous revenue requirement arising from UBI and the endogenous evolution of the tax base composition as the economy transitions. A closely related literature studies the optimal taxation of labor-displacing capital directly: Guerreiro, Rebelo, and Teles[2] show that a robot tax is part of the optimum during the transition when displaced workers cannot be fully compensated through the income tax, with the optimal rate declining toward zero in the long run; Thuemmel[3] finds quantitatively small optimal robot taxes once the nonlinear income tax adjusts; and Costinot and Werning[4] derive optimal technology regulation from sufficient statistics, obtaining modest optimal rates. Our setting differs in the scale of the revenue requirement, a UBI of 23 through 54 percent of per-capita gross domestic product (GDP) in the computed optimum, and in the presence of an upward-sloping AI capital supply with taxable quasi-rents; Section 6.6 returns to the comparison with Guerreiro, Rebelo, and Teles and with Thuemmel. Beraja and Zorzi[5] show that where reallocation is slow and borrowing limits are tight enough that displaced workers become borrowing constrained, the laissez-faire equilibrium is generically constrained inefficient, and automation is excessive. Their government can tax automation but lacks the redistributive transfers that would fully relieve those borrowing constraints. Under welfare weights that give it no motive to redistribute, it finds it optimal to slow automation while labor reallocates, and it has no reason to tax automation once labor reallocation has ended. In their quantitative model the optimal tax on investment in automation starts at about 5 percent and rises to 7 percent over a decade before it declines, and it raises the half-life of automation from 16 years at the laissez-faire to 22 years at the optimum. The planner of this paper chooses a universal transfer jointly with its taxes, and along the computed transition it taxes AI capital less than the static optimum at the same AI intensity until the rate on the quasi-rent, the rent rate, reaches the upper limit of 100 percent (Section 7.7).

The second strand is the literature on technological unemployment and the task-based approach to production, developed by Autor, Levy, and Murnane; Acemoglu and Autor; Acemoglu

[1] Frank P. Ramsey, "A Contribution to the Theory of Taxation," *Economic Journal* 37, no. 145 (1927): 47–61; James A. Mirrlees, "An Exploration in the Theory of Optimum Income Taxation," *Review of Economic Studies* 38, no. 2 (1971): 175–208; Peter A. Diamond and James A. Mirrlees, "Optimal Taxation and Public Production I: Production Efficiency," *American Economic Review* 61, no. 1 (1971): 8–27; Anthony B. Atkinson and Joseph E. Stiglitz, "The Design of Tax Structure: Direct versus Indirect Taxation," *Journal of Public Economics* 6, no. 1–2 (1976): 55–75; Emmanuel Saez, "Using Elasticities to Derive Optimal Income Tax Rates," *Review of Economic Studies* 68, no. 1 (2001): 205–229; Peter A. Diamond and Emmanuel Saez, "The Case for a Progressive Tax: From Basic Research to Policy Recommendations," *Journal of Economic Perspectives* 25, no. 4 (2011): 165–190; Thomas Piketty, Emmanuel Saez, and Stefanie Stantcheva, "Optimal Taxation of Top Labor Incomes: A Tale of Three Elasticities," *American Economic Journal: Economic Policy* 6, no. 1 (2014): 230–271.

[2] Joao Guerreiro, Sergio Rebelo, and Pedro Teles, "Should Robots Be Taxed?" *Review of Economic Studies* 89, no. 1 (2022): 279–311.

[3] Uwe Thuemmel, "Optimal Taxation of Robots," *Journal of the European Economic Association* 21, no. 3 (2023): 1154–1190.

[4] Arnaud Costinot and Iván Werning, "Robots, Trade, and Luddism: A Sufficient Statistic Approach to Optimal Technology Regulation," *Review of Economic Studies* 90, no. 5 (2023): 2261–2291.

[5] Martin Beraja and Nathan Zorzi, "Inefficient Automation," *Review of Economic Studies* 92, no. 1 (2025): 69–96.

and Restrepo; Autor; and Korinek and Stiglitz.[6] This literature provides the theoretical foundation for our unemployment mechanism and motivates our focus on the elasticity of substitution between AI capital and labor as the parameter that governs the mechanism. The Acemoglu–Restrepo distinction between displacement effects and reinstatement effects is central to our analysis: in Section 3 we provide a task-based microfoundation that allows reinstatement to be modeled explicitly, and we show how the aggregate elasticity of substitution rises endogenously with the range of tasks AI can perform. Acemoglu and Restrepo[7] study automation where some jobs pay workers rents—premia over what the same workers could earn in other tasks—and show that automation targets high-rent tasks, which dissipates the rents, amplifies the wage losses of exposed groups, and is inefficient. For the United States between 1980 and 2016 they estimate that automation accounts for 52 percent of the rise in between-group wage inequality, about one-fifth of that share through rent dissipation, and that inefficient rent dissipation has offset between 60 and 90 percent of the productivity gains from automation. Their rents arise from a wedge in the labor market between wages and opportunity costs; the quasi-rents on which the computed optimum of this paper places the AI capital tax are returns to the scarce inputs complementary to AI capital (Sections 4.3 and 6.6). Ide and Talamàs[8] model AI as computing resources converted into agents of a given knowledge, in an economy where the less knowledgeable do routine work and the more knowledgeable handle the exceptions. The least knowledgeable gain more from AI that only advises humans than from AI that also works on its own, the most knowledgeable gain more from AI that works on its own, and output is higher with it. In their baseline every human is employed, so AI changes the roles humans take and their earnings; in an extension in which computing resources outstrip production opportunities, every human less knowledgeable than AI is unemployed. Jones[9] sets out a task-based view in which tasks are complements, so that output is limited by the tasks not yet automated. With an elasticity of substitution of 0.2 across tasks, making the tasks that account for half of GDP infinitely productive raises GDP by 19 percent, and doubling income per person requires the same for the tasks that account for 94 percent of GDP. He notes that the labor share can fall toward zero while labor income rises, and that larger output makes redistribution through taxes and benefits more affordable.[10] In this paper, the mapping of Section 3, which takes AI intensity to the share of tasks AI performs, assumes equal task productivities and a task-level elasticity of substitution equal to

[6] David H. Autor, Frank Levy, and Richard J. Murnane, "The Skill Content of Recent Technological Change: An Empirical Exploration," *Quarterly Journal of Economics* 118, no. 4 (2003): 1279–1333; Daron Acemoglu and David Autor, "Skills, Tasks and Technologies: Implications for Employment and Earnings," in *Handbook of Labor Economics*, vol. 4B, ed. David Card and Orley Ashenfelter (Amsterdam: Elsevier, 2011), 1043–1171; Daron Acemoglu and Pascual Restrepo, "The Race Between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment," *American Economic Review* 108, no. 6 (2018): 1488–1542; Acemoglu and Restrepo, "Robots and Jobs: Evidence from US Labor Markets," *Journal of Political Economy* 128, no. 6 (2020): 2188–2244; Acemoglu and Restrepo, "Tasks, Automation, and the Rise in U.S. Wage Inequality," *Econometrica* 90, no. 5 (2022): 1973–2016; David H. Autor, "Work of the Past, Work of the Future," *AEA Papers and Proceedings* 109 (2019): 1–32; Anton Korinek and Joseph E. Stiglitz, "Artificial Intelligence and Its Implications for Income Distribution and Unemployment," in *The Economics of Artificial Intelligence: An Agenda*, ed. Ajay Agrawal, Joshua Gans, and Avi Goldfarb (Chicago: University of Chicago Press, 2019), 349–390.

[7] Daron Acemoglu and Pascual Restrepo, "Automation and Rent Dissipation: Implications for Wages, Inequality, and Productivity," *Quarterly Journal of Economics* 141, no. 2 (2026): 1521–1579.

[8] Enrique Ide and Eduard Talamàs, "Artificial Intelligence in the Knowledge Economy," *Journal of Political Economy* 133, no. 12 (2025): 3762–3800.

[9] Charles I. Jones, "AI and Our Economic Future," *Journal of Economic Perspectives* 40, no. 3 (2026): 9–11. Both numbers are stated on page 11, in his discussion of his Figure 2, whose source is Charles I. Jones and Christopher Tonetti, "Past Automation and Future A.I.: How Weak Links Tame the Growth Explosion" (unpublished manuscript, 2026), as reported in Jones, "AI and Our Economic Future," 11.

[10] Jones, "AI and Our Economic Future," 14.

the aggregate elasticity between AI and labor, 1.8 in the calibrated path (Section 7.7). Within that mapping, tasks are substitutes rather than complements.

The third strand is the literature on UBI as a response to technological change, including Hoynes and Rothstein, Ghatak and Maniquet, and Banerjee, Niehaus, and Suri.[11] Much of this literature is empirical, studying existing or experimental UBI programs at modest scales. The experiment of Vivalt et al.[12] assigned 1,000 low-income adults aged 21 through 40 in Illinois and Texas at random to receive $1,000 a month for three years, against 2,000 who received $50 a month. The transfer lowered employment by 4.2 percentage points and total individual income excluding the transfers by about $1,900 a year, and recipients used the time released mainly for leisure; the authors estimate an income elasticity of participation of −0.20. The response measured is the effect of an unconditional transfer on whether recipients work, the margin through which the UBI of the computed optimum raises non-participation at year 0 above the laissez-faire rate (Section 7.7). The authors leave the general-equilibrium effects of such a program at scale to further work. Our contribution is to embed UBI in a general-equilibrium model where the magnitude of the program is driven by the pace of AI adoption, where the financing structure is derived as part of the optimum, and where the welfare gain from UBI includes a formal insurance component against idiosyncratic skill-devaluation risk.

The fourth strand is the directed technical change literature,[13] which suggests that the pace and direction of AI development are themselves endogenous to factor prices and policy. We take the evolution of AI intensity as predetermined in our baseline model and examine the implications of endogenous technical change in Section 9. The paper also connects to the emerging literature on data as a factor of production,[14] which informs our classification of data rents as pure economic rents outside AI (Section 6.1), separate from the quasi-rent to the scarce inputs complementary to AI capital.

3 A Task-Based Foundation for AI Production

Before presenting the aggregate production structure used in the remainder of the paper, we sketch a task-based microfoundation that generates a constant-elasticity-of-substitution (CES) aggregator with two properties that matter for the analysis: first, the aggregate elasticity of substitution between AI capital and labor is an equilibrium object that rises as AI's task reach expands; second, the strength of the reinstatement channel is a modeled feature rather than an assumption. The microfoundation is standard in structure; we state it compactly here and use it in Sections 4 and 7 to give content to otherwise free parameters.

[11] Hilary Hoynes and Jesse Rothstein, "Universal Basic Income in the United States and Advanced Countries," *Annual Review of Economics* 11 (2019): 929–958; Maitreesh Ghatak and François Maniquet, "Universal Basic Income: Some Theoretical Aspects," *Annual Review of Economics* 11 (2019): 895–928; Abhijit Banerjee, Paul Niehaus, and Tavneet Suri, "Universal Basic Income in the Developing World," *Annual Review of Economics* 11 (2019): 959–983.

[12] Eva Vivalt et al., "The Employment Effects of a Guaranteed Income: Experimental Evidence from Two U.S. States," *Quarterly Journal of Economics* 141, no. 4 (2026): 2675–2734.

[13] Daron Acemoglu, "Directed Technical Change," *Review of Economic Studies* 69, no. 4 (2002): 781–809; Acemoglu, "Harms of AI," in *The Oxford Handbook of AI Governance*, ed. Justin B. Bullock et al. (Oxford: Oxford University Press, 2023).

[14] Charles I. Jones and Christopher Tonetti, "Nonrivalry and the Economics of Data," *American Economic Review* 110, no. 9 (2020): 2819–2858.

3.1 Tasks, Assignment, and Aggregation

Final output is produced by combining a unit continuum of tasks indexed by $j \in [0,1]$:

$$Y = A\left[\int_0^1 y(j)^{\frac{\eta_T - 1}{\eta_T}}\, dj\right]^{\frac{\eta_T}{\eta_T - 1}}, \tag{1}$$

where $\eta_T > 0$ is the task-level elasticity of substitution and A is Hicks-neutral total factor productivity (TFP). Each task $y(j)$ can be produced either by labor or by AI capital:

$$y(j) = \begin{cases} \phi_L(j)\,\ell(j) & \text{if labor performs task } j, \\ \phi_K(j)\,k(j) & \text{if AI performs task } j, \end{cases} \tag{2}$$

where $\phi_L(j)$ and $\phi_K(j)$ are the productivities of labor and AI on task j respectively, and $\ell(j), k(j)$ are the factor inputs. Cost minimization assigns task j to AI whenever $r_A/\phi_K(j) < w/\phi_L(j)$, and to labor otherwise. Let $\mathcal{J}_K \subseteq [0,1]$ denote the set of tasks performed by AI at equilibrium prices, and let $n_K \equiv |\mathcal{J}_K|$ be its measure, the share of tasks AI performs.

3.2 From Task-Level to Aggregate Substitution

Given an assignment $\mathcal{J}_K$ and its complement $\mathcal{J}_L$, aggregating across tasks yields an aggregate production function of the form

$$Y = A\left[\Gamma_K(\mathcal{J}_K)\, K_A^{\frac{\sigma-1}{\sigma}} + \Gamma_L(\mathcal{J}_L)\, L^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}, \tag{3}$$

where $\sigma = \sigma(\eta_T, \mathcal{J}_K)$ is the aggregate elasticity of substitution and Γ_K, Γ_L are distributional weights that depend on the task-specific productivities aggregated over $\mathcal{J}_K$ and $\mathcal{J}_L$. The point is that σ is not a deep parameter: holding η_T fixed, expanding $\mathcal{J}_K$ (AI performs more tasks) raises the aggregate σ because the share of output produced in regions of the task space where AI and labor are close substitutes increases.

Defining AI intensity $\alpha \in (0,1)$ as the productivity-weighted share of AI-performed tasks,

$$\alpha \equiv \frac{\Gamma_K(\mathcal{J}_K)}{\Gamma_K(\mathcal{J}_K) + \Gamma_L(\mathcal{J}_L)}, \tag{4}$$

equation (3) takes the form of a standard two-factor CES aggregate. This is the aggregator we use as a component of the production structure in Section 4, with the understanding that both α and σ are endogenous to the task assignment. AI intensity and the share of tasks AI performs are distinct objects. With equal task productivities, and with the aggregate elasticity at its intensive-margin value $\sigma = \eta_T$, the weights are $\Gamma_K = n_K^{1/\sigma}$ and $\Gamma_L = (1 - n_K)^{1/\sigma}$ (Appendix A), so that

$$\frac{\alpha}{1-\alpha} = \left(\frac{n_K}{1-n_K}\right)^{1/\sigma}. \tag{5}$$

At the calibrated $\sigma = 1.8$, the initial intensity $\alpha_0 = 0.15$ of Section 7.7 corresponds to 4.2 percent of tasks and the adoption ceiling $\alpha_{\max} = 0.75$ of Section 7.1 to 87.8 percent.

3.3 Implications for Reinstatement

In the Acemoglu–Restrepo framework, reinstatement corresponds to the creation of new tasks that, at prevailing factor prices, are assigned to labor. Reinstatement offsets displacement when new tasks arrive disproportionately in regions of the task space where $\phi_L(j)/\phi_K(j)$ is large, so that new tasks are labor-assigned at prevailing prices. Reinstatement is weak when new tasks arrive uniformly over the task space, or when AI's productivity advantage $\phi_K(j)/\phi_L(j)$ is broadly rising, so that new tasks are increasingly AI-assigned.

Condition 6 of Section 4.7 is a condition on the relative rates of task creation and task displacement: the fraction of new tasks assigned to labor must be large enough that $|\mathcal{J}_L|$ does not contract. Whether this condition holds for AI is an empirical question the paper does not resolve; the microfoundation states what is being assumed.

4 The Baseline Economy

This section sets out the production structure, the labor market with its employment margin, the AI capital supply with its scarce complementary inputs, the ownership structure, and the equilibrium, in that order, before the comparative statics of Section 4.6 and the unemployment conditions that close the section.

4.1 A Three-Factor Production Structure

The economy produces a single final good using three factors: AI capital K_A, non-AI (traditional) capital K_N, and effective labor L. The production technology is a nested CES:

$$Y = A\left[\xi_X X^{\frac{\nu-1}{\nu}} + (1-\xi_X) K_N^{\frac{\nu-1}{\nu}}\right]^{\frac{\nu}{\nu-1}}, \tag{6}$$

$$X = \left[\alpha K_A^{\frac{\sigma-1}{\sigma}} + (1-\alpha) L^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}, \tag{7}$$

where X is the AI–labor composite, $\sigma > 0$ is the elasticity of substitution between K_A and L within the inner nest, $\nu > 0$ is the elasticity of substitution between X and K_N in the outer nest, $\xi_X \in (0,1)$ is the outer-nest weight, and $\alpha \in (0,1)$ is AI intensity within the inner nest. We assume $\sigma > \nu$: AI substitutes more readily for labor than non-AI capital substitutes for labor, consistent with the task-based microfoundation of Section 3 (AI and labor compete over a common task space; non-AI capital and labor less so). Setting $\xi_X = 1$ gives the two-factor CES case; we retain the three-factor structure because it permits differentiated tax treatment of AI and non-AI capital and is closer to the tax-policy question under study.

Firms are competitive and maximize profits taking factor prices as given. The first-order conditions yield factor demands:

$$w = \frac{\partial Y}{\partial L} = A\,\xi_X\,(1-\alpha)\left(\frac{Y}{AX}\right)^{1/\nu}\left(\frac{X}{L}\right)^{1/\sigma}, \tag{8}$$

$$r_A = \frac{\partial Y}{\partial K_A} = A\,\xi_X\,\alpha\left(\frac{Y}{AX}\right)^{1/\nu}\left(\frac{X}{K_A}\right)^{1/\sigma}, \tag{9}$$

$$r_N = \frac{\partial Y}{\partial K_N} = A\,(1-\xi_X)\left(\frac{Y}{AK_N}\right)^{1/\nu}. \tag{10}$$

Factor shares satisfy $s_L + s_{K_A} + s_{K_N} = 1$ by constant returns to scale. The partial two-factor subresults from equation (3) of Section 3 apply within the inner nest, so the algebra of the comparative statics below mirrors the two-factor case for responses to α conditional on the inner composite X.

4.2 Heterogeneous Labor and the Employment Margin

The economy contains N workers indexed by skill type $\theta \in [0,1]$ with cumulative distribution $F_t(\theta)$ and density $f_t(\theta)$. Skill reflects accumulated human capital, cognitive complementarity with AI, and occupational fit. The subscript t anticipates Section 7.2, in which F_t evolves through education and retraining; in the static analysis we write F for brevity. Each worker of type θ supplies effective labor

$$\ell(\theta) = \theta^{\gamma}, \qquad \gamma > 0, \tag{11}$$

with $\gamma > 1$ capturing super-linear productivity of high-skill workers in AI-complementary tasks.

Each worker has an outside option yielding utility-equivalent income $w_r(b)$, where the dependence on the UBI level b is explicit because UBI raises the reservation wage. In the baseline without UBI, w_r is low. A worker of type θ participates in market work if and only if

$$w \cdot \theta^{\gamma} \geq w_r(b), \tag{12}$$

which defines the employment threshold

$$\theta^{*} = \left(\frac{w_r(b)}{w}\right)^{1/\gamma}. \tag{13}$$

Three specifications the numerical solution uses are stated here (Appendix C). Hours are chosen on an intensive margin, $h(\theta) = [(1-\tau_L(\theta))\,w\,\theta^{\gamma}]^{e}$ with $e = 0.25$ and disutility $h^{1+1/e}/(1+1/e)$, so that the labor tax carries the elasticity the Saez formula of Section 6.5 presupposes. The marginal worker compares after-tax earnings with $w_r(b)$, so that the bottom labor rate enters the participation margin as Section 6.5 requires. And the labor income tax schedule $\tau_L(\theta)$ is a rate for each tercile of the employed, bottom, middle, and top, joined smoothly at the tercile boundaries. Without the first specification the static optimum is a corner with every labor rate at its bound; without the second the labor tax is non-distortionary. Aggregate employed effective labor is

$$L = N\int_{\theta^{*}}^{1} \theta^{\gamma}\, f(\theta)\, d\theta, \tag{14}$$

and the *non-participation rate*, the share of the population that does not work, is $u = F(\theta^{*})$; *mass unemployment* is a large value of u.

4.3 AI Capital Supply with Scarce Complementary Inputs

Where AI capital is supplied at a fixed required return $r + \delta$, supply is perfectly elastic and the incidence of AI capital taxation falls entirely on other factors. We replace that benchmark with an upward-sloping supply curve arising from scarce complementary inputs: frontier compute capacity, electricity and cooling infrastructure, specialized chips, and specialized human capital in model training and deployment. Let Z denote the stock of these complementary inputs, which is in fixed supply in the short run and expands slowly over time. The perfectly elastic benchmark underlies the classical zero-long-run-capital-tax results of Judd and Chamley;[15] Straub and Werning show that those results are less general than long understood.[16] The upward-sloping supply developed here departs from that benchmark on economic rather than technical grounds: the inputs complementary to AI capital are scarce.

The AI capital production technology is

$$K_A = \Phi(I_K, Z), \qquad \Phi_I > 0,\ \Phi_{II} < 0,\ \Phi_Z > 0, \tag{15}$$

where I_K is gross investment in AI capital and Φ is concave in I_K for given Z. Funds are supplied elastically at the normal return $r+\delta$ per unit of AI capital, and the scarce input earns the remainder of the supply price, so the supply price of AI capital, $s(K_A)$, decomposes as

$$s(K_A) = r + \delta + \Pi_Z / K_A, \tag{16}$$

where Π_Z is the quasi-rent accruing to the scarce complementary input: the supply price less the normal return, on every unit of the stock. The supply price rises with the stock as the complementary input binds,

$$s(K_A) = (r + \delta)\,(K_A/\bar{K})^{1/\varepsilon_K}, \tag{17}$$

with $\bar{K}$ the stock at which the supply price equals the normal return. A tax at rate $\tau_{K_A}^{norm}$ on the normal return (Section 6.1) is a wedge of $\tau_{K_A}^{norm}\,(r + \delta)$ per unit between the rental rate and the supply price,

$$r_A - \tau_{K_A}^{norm}\,(r + \delta) = s(K_A), \tag{18}$$

so the AI factor payment divides, without remainder, into the normal-return tax, the normal return, and the quasi-rent: $r_A K_A = \tau_{K_A}^{norm}\,(r + \delta)\,K_A + (r + \delta)\,K_A + \Pi_Z$. A tax on the quasi-rent does not enter (18), because Π_Z accrues to the input in fixed supply. The term Π_Z/K_A is strictly positive whenever Z is scarce ($K_A > \bar{K}$), and it rises with K_A as long as the complementary input binds. The AI capital supply elasticity is then

$$\varepsilon_K \equiv \frac{\partial \ln K_A}{\partial \ln r_A} = \frac{\Phi_I \cdot I_K/K_A}{-\Phi_{II} \cdot I_K + \Phi_I \cdot I_K/K_A}, \tag{19}$$

which is finite when $\Phi_{II} < 0$ and which falls as Z binds more tightly.

The stock carries a resource cost, which its owners bear: the replacement of depreciated capital, δK_A, and the annuity at the interest rate on net investment from the stock $K_{A,0}$ inherited at the start of the transition,

$$c_A \equiv \delta\, K_A + r\,(K_A - K_{A,0}), \tag{20}$$

[15] Kenneth L. Judd, "Redistributive Taxation in a Simple Perfect Foresight Model," *Journal of Public Economics* 28, no. 1 (1985): 59–83; Christophe Chamley, "Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives," *Econometrica* 54, no. 3 (1986): 607–622.

[16] Ludwig Straub and Iván Werning, "Positive Long-Run Capital Taxation: Chamley–Judd Revisited," *American Economic Review* 110, no. 1 (2020): 86–119.

deducted from the owners' capital income in (21), so that consumption is net of the cost of maintaining and expanding the stock. Where the interest rate equals the discount rate, as in the computed solution, the present value of c_A along a path equals that of gross investment in AI capital from the inherited stock; the annuity form charges at each date the replacement of the stock and the interest on the net investment made since the start, rather than the investment when it is made.

This structure has three consequences. First, a share of AI capital returns consists of the quasi-rent Π_Z to the fixed input, which can be taxed without distorting the marginal investment decision. Second, the AI capital tax incidence is shared between AI capital owners and the owners of scarce complementary inputs, rather than falling entirely on other factors. Third, the supply elasticity ε_K is an empirical object that varies with the tightness of complementary-input constraints, providing content to Condition 3 of the unemployment conditions (Section 4.7).

4.4 Ownership, Wealth, and Distribution

We formalize capital ownership through a measure $\omega(\theta)$ giving the share of aggregate capital, AI and non-AI alike, held by households with skill type θ, normalized so that $\int_0^1 \omega(\theta)\, f(\theta)\, d\theta = 1$. Empirically ω is increasing in θ and, in the upper tail, concentrated. Aggregate consumption for type θ is

$$c(\theta) = (1-\tau_C)\left[(1-\tau_L(\theta))\, w\, \theta^{\gamma}\, \mathbf{1}\{\theta \geq \theta^*\} + \left((1-\tau_K)\, r_A\, K_A - c_A + (1-\tau_{K_N})\, r_N\, K_N\right)\omega(\theta) + b\right], \tag{21}$$

where τ_K is the effective tax rate on the AI return, so that $(1-\tau_K)\, r_A K_A$ is the AI factor payment net of the normal-return and rent taxes of Section 6.1, and c_A is the resource cost (20); workers and capital owners are one population, distinguished by θ and $\omega(\theta)$. The distributional welfare weight μ_i used in the Ramsey system of Section 6.2 is then the covariance of tax base i with the social marginal utility of consumption, taken over the full joint distribution of θ and ω. As inequality in ω rises, $|\mu_K|$ rises and the redistributive value of AI capital taxation rises with it.

4.5 Equilibrium

Given exogenous K_A, K_N, Z, and b, equilibrium is characterized by a wage w^* and threshold θ^* jointly satisfying equations (8), (13), and (14). Substituting (13) into (14) yields labor supply $L(w, b)$, which is increasing in w and decreasing in b (through $w_r(b)$). Substituting this into (8) gives the equilibrium condition

$$w = A\, \xi_X\, (1-\alpha) \left(\frac{Y}{AX}\right)^{1/\nu} \left(\frac{X(K_A, L(w,b))}{L(w,b)}\right)^{1/\sigma}. \tag{22}$$

Lemma 1 (Existence and uniqueness)**.** *Suppose f is continuous and strictly positive on $[0,1]$, $\gamma > 0$, $\sigma \in (1,\infty)$, and $\nu \in (0,\infty)$. Then for any $K_A > 0$, $K_N > 0$, and $b \geq 0$ with $w_r(b) > 0$, equation (22) admits a unique solution $w^* \in (w_r(b), \infty)$.*

Proof. The right-hand side of (22) is continuous and strictly decreasing in w: higher w lowers θ^*, raises L, and reduces the marginal product of labor. The left-hand side is the identity. As $w \downarrow w_r(b)$ we have $\theta^* \to 1$ and $L \to 0$; with $\sigma > 1$ the inner-nest exponent $(\sigma-1)/\sigma$ is positive, so $X \to \alpha^{\sigma/(\sigma-1)} K_A > 0$ and $X/L \to \infty$, and the right-hand side is unbounded above. As $w \to \infty$, $\theta^* \to 0$ and $L \to L_{\max} = N \int_0^1 \theta^{\gamma} f(\theta)\, d\theta < \infty$, so the right-hand side is bounded. Continuity and

monotonicity give existence and uniqueness by the intermediate value theorem. For $\sigma < 1$ the marginal product of labor is bounded as $L \to 0$; existence then additionally requires $w_r(b)$ below that bound, and uniqueness follows from the same monotonicity. The analysis below concerns the gross-substitutes case $\sigma > 1$ (Condition 1 of Section 4.7), so the lemma is stated for that case. ■

4.6 General-Equilibrium Comparative Statics of AI Adoption

We now derive the response of the equilibrium wage to rising AI intensity α. Because K_A and L both respond to α in general equilibrium, the partial-equilibrium derivative (differentiating the marginal product of labor at fixed K_A, L) is not sufficient. We decompose the general-equilibrium response into three channels.

Let a superscript $*$ denote equilibrium values. Totally differentiating the equilibrium condition with respect to α gives

$$\frac{dw^*}{d\alpha} = \underbrace{\left.\frac{\partial w}{\partial \alpha}\right|_{K_A,L}}_{\text{direct substitution}} + \underbrace{\left.\frac{\partial w}{\partial K_A}\right|_{\alpha,L} \cdot \frac{dK_A^*}{d\alpha}}_{\text{capital-stock response}} + \underbrace{\left.\frac{\partial w}{\partial L}\right|_{\alpha,K_A} \cdot \frac{dL^*}{d\alpha}}_{\text{labor-supply response}}. \tag{23}$$

The direct-substitution term is negative under the conditions of Proposition 1 below; the capital-stock response is positive (higher α raises r_A, which raises K_A^* through the supply schedule (17)), and this positive response dampens the wage decline. The labor-supply response sign depends on whether the wage decline reduces L (positive feedback on the wage) or whether UBI-induced rise in w_r dominates (further labor-supply contraction, reinforcing the wage decline).

Proposition 1 (Wage response to AI intensity)**.** *Let $\varrho \equiv (\sigma-1)/\sigma$, let $x \equiv (K_A/L)^{\varrho}$, and let $s_X \in (0,1]$ denote the outer-nest income share of the composite X. Holding K_A, L, and K_N fixed, log-differentiation of (8) gives*

$$\frac{1}{w}\left.\frac{\partial w}{\partial \alpha}\right|_{K_A,L,K_N} = -\frac{1}{1-\alpha} + \left(\frac{1}{\sigma} - \frac{1-s_X}{\nu}\right)\frac{\sigma}{\sigma-1}\cdot\frac{x-1}{\alpha x + (1-\alpha)}. \tag{24}$$

Suppose $\sigma > 1$ and $1/\tilde{\sigma} \equiv 1/\sigma - (1-s_X)/\nu > 0$, and define $c_\sigma \equiv \tilde{\sigma}(\sigma-1)/\sigma > 0$. Then: (i) *if $\alpha \geq 1/(1+c_\sigma)$, then $\partial w/\partial\alpha < 0$ for all $K_A/L > 0$;* (ii) *if $\alpha < 1/(1+c_\sigma)$, then $\partial w/\partial\alpha < 0$ if and only if $x < (1+c_\sigma)(1-\alpha)/\big(1-\alpha(1+c_\sigma)\big)$. The two-factor case $s_X = 1$ gives $c_\sigma = \sigma - 1$, and (i)–(ii) reduce to the conditions $\alpha\sigma \geq 1$ and $x < \sigma(1-\alpha)/(1-\alpha\sigma)$ respectively.* (iii) *In general equilibrium, $dw^*/d\alpha < 0$ whenever the right-hand side of (24) is strictly negative and the capital-stock response does not overturn it; for $K_A/L \leq 1$ a sufficient condition is $\varepsilon_K < \tilde{\sigma}$, and for $K_A/L > 1$ the analogous bound scales approximately as $\varepsilon_K \lesssim \tilde{\sigma}/x$.*

Proof sketch. Write $\ln w = \ln(A\xi_X) + \ln(1-\alpha) + (1/\nu)(\ln Y - \ln A - \ln X) + (1/\sigma)(\ln X - \ln L)$ from (8). At fixed factor quantities, $\partial \ln Y/\partial\alpha = s_X\,\partial \ln X/\partial\alpha$ and $\partial \ln X/\partial\alpha = \big(\sigma/(\sigma-1)\big)(x-1)/(\alpha x + 1 - \alpha)$, which together give (24). Multiplying through by the positive quantity $c_\sigma\,(1-\alpha)(\alpha x + 1 - \alpha)$ and collecting terms shows that the right-hand side of (24) is negative if and only if $x\left[(1-\alpha) - c_\sigma\alpha\right] < (1+c_\sigma)(1-\alpha)$, from which (i) and (ii) follow. For (iii): the capital-stock channel in (23) equals $(\partial \ln w/\partial \ln K_A)\,(d\ln K_A^*/d\alpha)$, with $\partial \ln w/\partial \ln K_A = (1/\tilde{\sigma})\,s_K^X$, where $s_K^X \equiv \alpha x/(\alpha x + 1 - \alpha)$ is the AI share within the composite, and $d\ln K_A^*/d\alpha \leq \varepsilon_K\,\partial \ln r_A/\partial\alpha \leq \varepsilon_K/\alpha$ for $x \leq 1$. The channel is therefore bounded by $\varepsilon_K\,x/\big(\tilde{\sigma}\,(\alpha x + 1 - \alpha)\big) \leq \varepsilon_K/\big(\tilde{\sigma}\,(1-\alpha)\big)$, which is smaller than the magnitude $1/(1-\alpha)$ of the leading direct term whenever $\varepsilon_K < \tilde{\sigma}$. The labor-supply channel is a damping

feedback: a lower wage reduces participation, which raises the marginal product of labor, and the fixed point is unique by Lemma 1; it therefore cannot reverse the sign. ■

The economic content is that, under the conditions of Proposition 1, among them that AI and labor are gross substitutes and that the AI capital supply is not too elastic, shifting production weight toward AI reduces the equilibrium wage. Relative to the two-factor case, the outer nest makes a wage decline more likely: with $s_X < 1$, part of the composite's productivity gain accrues to non-AI capital rather than to labor, which shrinks the positive productivity term in (24). When $(1 - s_X)/\nu \geq 1/\sigma$ the productivity term reverses sign and the case analysis differs; the calibrated range ($s_X \in [0.75, 0.90]$, $\nu = 0.6$, $\sigma = 1.8$, so $1/\tilde{\sigma} \approx 0.14$ through 0.39) does not enter this region. For the parameter combinations we treat as empirically relevant ($\sigma \in [1.3, 2.5]$, $\alpha \in [0.15, 0.70]$, K_A/L in the single-digit range, $\varepsilon_K \leq 0.5$ along the calibrated path), condition (i) or (ii) holds and the general-equilibrium sign coincides with the partial-equilibrium sign.

The response of non-participation then follows from (13):

$$\frac{du^*}{d\alpha} = f(\theta^*)\,\frac{d\theta^*}{d\alpha} = -\frac{f(\theta^*)}{\gamma}\left(\frac{w_r(b)}{w^*}\right)^{1/\gamma}\frac{1}{w^*}\,\frac{dw^*}{d\alpha} > 0 \tag{25}$$

under the conditions of Proposition 1 that yield $dw^*/d\alpha < 0$. This is the first main result: under those conditions, among them that AI and labor are gross substitutes and that the AI capital supply is not too elastic, rising AI intensity raises non-participation in general equilibrium.

4.7 Conditions for Mass Unemployment

For non-participation to become quantitatively large, six conditions must be simultaneously satisfied. We state them in the three-factor, upward-sloping-supply environment and in the general-equilibrium form of Proposition 1. Jointly the six conditions are sufficient for mass unemployment, and each is necessary in the sense that its violation blunts the mechanism. Quantifying the non-participation rate that emerges under plausible values of these conditions requires functional specifications beyond those adopted in this section; the calibrated transition path in Section 7.7 illustrates magnitudes.

4.7.1 Condition 1: Substitutability

Condition 1 requires $\sigma > 1$ at the relevant operating margin, combined with parameter values placing the economy in the region of Proposition 1 where $dw^*/d\alpha < 0$. Aggregate historical estimates of the capital–labor elasticity cluster around 0.5 through 0.8, but task-level evidence suggests $\sigma > 1$ for routine cognitive tasks.[17] The task-based microfoundation of Section 3 implies that the effective aggregate σ rises as AI's task reach expands.

[17] Acemoglu and Restrepo, "Robots and Jobs," estimate displacement effects of industrial robots consistent with task-level substitutability in manufacturing; Acemoglu and Restrepo, "Tasks, Automation," estimate that task displacement accounts for 50 through 70 percent of the rise in wage inequality between demographic groups defined by education and experience in the United States between 1980 and 2016, a magnitude consistent with task-level elasticities above unity in the affected task categories. The aggregate σ relevant to our analysis is, per Section 3, a weighted average across tasks; the calibration in Section 7.7 uses $\sigma = 1.8$, which is in the upper range of task-level estimates but below the frictionless substitution limit.

4.7.2 Condition 2: Rising AI Intensity

Condition 2 requires $d\alpha/dt > 0$, driven by capability expansion, cost declines, and complementary investment. This dynamic is modeled explicitly in Section 7.

4.7.3 Condition 3: Finite but Not Extreme AI Capital Supply Elasticity

Condition 3 requires that ε_K satisfy the sufficient condition in Proposition 1 at the relevant operating point. The upward-sloping supply structure of Section 4.3 gives empirical content to this condition: ε_K depends on how tightly complementary inputs Z bind.

4.7.4 Condition 4: Skill Distribution Density

Condition 4 requires that $f(\theta)$ place mass in the range that falls below the employment threshold as θ^* rises. The response of non-participation $du/d\alpha = f(\theta^*)\, d\theta^*/d\alpha$ is large when the density is concentrated near the employment margin. Empirically, skill distributions in advanced economies are approximately log-normal, with mass concentrated in the middle of the distribution, which is the range AI can increasingly substitute for.

4.7.5 Condition 5: Slow Skill Adjustment

Condition 5 requires that F_t adjust slowly relative to $d\alpha/dt$. Section 7.2 models skill adjustment endogenously through an education investment technology; the implication is that policy can speed this adjustment but cannot eliminate the adjustment lag inherent in human-capital formation.

4.7.6 Condition 6: Weak Reinstatement

Condition 6 requires that new labor-using tasks do not emerge at the same rate as old tasks are automated. In the microfoundation of Section 3, this corresponds to a low arrival rate of new tasks in the region of the task space where ϕ_L/ϕ_K is large.

5 Risk, Insurance, and the Welfare Role of UBI

The deterministic framework of Section 4 captures the displacement mechanism but misses a feature of the policy problem that matters for the welfare evaluation: the risk that one's skill θ becomes devalued as AI's task reach expands is itself uninsurable by private markets. Workers cannot write contracts today that compensate them tomorrow for adverse realizations of AI capability developments, because AI capability developments are aggregate rather than idiosyncratic events, subject to adverse selection when the information structure is asymmetric, and occur over horizons longer than standard insurance contracts. UBI acts in part as public insurance against this risk. This section develops the argument formally.

5.1 Idiosyncratic Skill-Devaluation Risk

Consider a worker with initial skill type $\theta_0 \in [0,1]$. As AI's task reach expands, the worker's effective skill evolves stochastically. In keeping with the continuous-time environment of the remainder of the paper, we write the skill process as a diffusion with drift:

$$d\theta_t = -D(\alpha_t, \theta_t)\, dt + s(\theta_t)\, d\mathcal{W}_t, \tag{26}$$

where $D(\alpha_t, \theta_t) \geq 0$ is the expected rate of skill devaluation with $\partial D/\partial \alpha > 0$ and $\partial D/\partial \theta < 0$ (higher-skill workers' skills depreciate less rapidly under rising α), $\mathcal{W}_t$ is a standard Wiener process capturing the idiosyncratic component, and $s(\theta_t) > 0$ is the instantaneous volatility of the idiosyncratic shock. The stochastic term captures the observation that even within a given skill band, some workers find their specific skills more or less devalued by the particular direction AI capability developments take. The cross-sectional distribution of θ at time t evolves according to the Fokker–Planck equation implied by (26) combined with the policy-driven education flow of Section 7.2.

Absent insurance, the consumption risk facing a worker is the residual variation in θ_t combined with the employment threshold: $c(\theta_t) = w\theta_t^{\gamma}$ if $\theta_t \geq \theta^*$, and zero otherwise. The worker bears this risk privately.

5.2 The Insurance Value of UBI

With a UBI transfer b, consumption becomes $c(\theta_t) = w\theta_t^{\gamma}\mathbf{1}\{\theta_t \geq \theta^*\} + b$, which is bounded below by b regardless of the realization of θ_t. For a worker with concave utility $U(c)$, the welfare gain from UBI at level b admits the exact decomposition

$$\mathbb{E}[U(c+b)] - \mathbb{E}[U(c)] = \underbrace{U(\bar{c}+b) - U(\bar{c})}_{\text{mean-income gain}} + \underbrace{\big[\mathbb{E}[U(c+b)] - U(\bar{c}+b)\big] - \big[\mathbb{E}[U(c)] - U(\bar{c})\big]}_{\text{insurance value}}, \tag{27}$$

where $\bar{c} \equiv \mathbb{E}[c]$ is mean pre-UBI consumption; the identity holds by adding and subtracting $U(\bar{c}+b) - U(\bar{c})$. The first brace is the gain from the transfer evaluated at the mean of the distribution, the component a mean-income-only welfare criterion captures. The second brace is the reduction in the utility cost of consumption risk: each of its two bracketed terms is the (negative) gap between expected utility and utility at the mean, and their difference is positive under decreasing absolute risk aversion, since adding the constant b lowers absolute risk aversion at every consumption level. The insurance value is quantitatively large when the pre-UBI distribution places mass at low consumption, because b bounds consumption away from levels at which marginal utility is extreme.

5.2.1 A Constant Relative Risk Aversion Illustration

To quantify the insurance term, we compute the reduction in the proportional risk premium under a constant relative risk aversion (CRRA) utility specification $U(c) = c^{1-\gamma_R}/(1-\gamma_R)$ with $\gamma_R > 0$. Suppose pre-UBI consumption is log-normally distributed with mean $\bar{c}$ and coefficient of variation CV, so that the log-variance is $\varsigma^2 = \ln(1 + \mathrm{CV}^2)$. The certainty equivalent of such a consumption plan is $\mathrm{CE} = \bar{c}\, e^{-\gamma_R \varsigma^2/2}$, so the proportional risk premium is $\pi/\bar{c} = 1 - e^{-\gamma_R \varsigma^2/2}$. A UBI of b shifts the distribution to $c + b$, which lowers the coefficient of variation to $\mathrm{CV}/(1 + b/\bar{c})$; approximating $c + b$ as log-normal with the shifted mean and reduced log-variance $\varsigma'^2 = \ln\!\big(1 + \mathrm{CV}^2/(1 + b/\bar{c})^2\big)$, the insurance value of UBI is the reduction in the risk premium,

$$\frac{\Delta_{\text{ins}}}{\bar{c}} \approx \left(1 - e^{-\gamma_R \varsigma^2/2}\right) - \left(1 + \frac{b}{\bar{c}}\right)\left(1 - e^{-\gamma_R \varsigma'^2/2}\right). \tag{28}$$

For $\gamma_R = 2$, CV $= 0.6$, and $b = 0.23\,\bar{c}$ (stipulated values), equation (28) evaluates to approximately 0.028 of mean consumption per period. Across $\gamma_R \in [1.5, 3]$ the figure ranges from roughly 2.4 through 3.2 percent. The log-normal approximation understates the value when the pre-UBI distribution places mass near zero consumption (displacement without home production or family insurance), since marginal utility under CRRA is unbounded as consumption approaches zero. Treating the insurance value as an additional marginal welfare gain of the consumption floor implies an increment to the optimal b on the order of a few percentage points of per-capita GDP relative to a mean-income-only welfare criterion. A Ramsey planner that ignores the insurance component therefore understates the welfare gain from UBI and, in consequence, understates the optimal b.

5.3 Implications for the Welfare Objective

Incorporating insurance formally modifies the planner's objective in Section 6 from

$$W = \int_0^1 U(c(\theta))\, f(\theta)\, d\theta \tag{29}$$

to the expected-utility form

$$W = \mathbb{E}_0\left[\int_0^1 U(c(\theta))\; f(\theta)\, d\theta\right], \tag{30}$$

where the expectation $\mathbb{E}_0$ is taken over the idiosyncratic skill histories induced by (26). For the remainder of this paper we work with the expected-utility form (30) when formalizing the planner's problem, noting that the deterministic form is the special case $s(\theta) \equiv 0$; the numerical solution of Section 7.7 is computed in that special case (Appendix C).

6 The Optimal Tax Structure in Steady State

Having established the conditions under which the laissez-faire equilibrium features high nonparticipation, we derive the optimal tax structure for financing a UBI transfer b to each citizen. The treatment below incorporates the expected-utility welfare criterion of Section 5 and the multifactor ownership structure of Section 4.4.

6.1 The Planner's Problem

The social planner chooses tax instruments $\{\tau_L(\theta), \tau_{K_A}, \tau_{K_A}^{rent}, \tau_{K_N}, \tau_\Pi, \tau_C\}$ and UBI level b to maximize

$$W = \mathbb{E}_0\left[\int_0^1 U(c(\theta))\, f(\theta)\, d\theta\right], \tag{31}$$

where $c(\theta)$ is household consumption as given by (21) and $\mathbb{E}_0$ is the expectation over idiosyncratic skill histories. The government budget constraint is

$$N \cdot b + G = \tau_L\, wL + \tau_{K_A}^{norm}\,(r+\delta)\,K_A + \tau_{K_A}^{rent}\,\Pi_Z + \tau_{K_N}\, r_N\, K_N + \tau_\Pi\,\Pi + \tau_C\, C, \tag{32}$$

where we have split the AI capital tax into a normal-return component $\tau_{K_A}^{norm}$ applying to $(r+\delta)K_A$ and a rent component $\tau_{K_A}^{rent}$ applying to the quasi-rent Π_Z (cf. Section 4.3). The notation Π covers

pure economic rents outside AI (land, monopoly, and data rents), and C is consumption expenditure inclusive of the tax, so that the consumption tax collects what households pay. We write labor tax revenue as $\tau_L \cdot wL$ using an effective average rate, understanding that the underlying schedule $\tau_L(\theta)$ is nonlinear.

6.2 The Multi-Instrument Ramsey System

Let $\mathcal{I} \equiv \{L, K_A^{norm}, K_A^{rent}, K_N, \Pi, C\}$ denote the set of tax instruments and consider a perturbation $d\tau_i > 0$. The revenue effect, valued at the marginal cost of public funds λ, comprises the mechanical effect B_i plus the behavioral responses of all bases, own and cross; the welfare cost is the marginal-utility-weighted consumption loss (with $\partial c/\partial \tau_i \leq 0$); and the AI capital instruments carry an additional money-metric marginal external benefit $\mathcal{E}_i \geq 0$ through reduced adoption ($\mathcal{E}_i = 0$ for the other instruments). At the optimum, for each $i \in \mathcal{I}$,

$$\lambda \left[B_i + \sum_{j \in \mathcal{I}} \tau_j \frac{\partial B_j}{\partial \tau_i} \right] + \lambda \, \mathcal{E}_i = - \int_0^1 U'(c(\theta)) \frac{\partial c(\theta)}{\partial \tau_i} f(\theta) \, d\theta, \tag{33}$$

where B_j is the equilibrium level of tax base j. The left-hand side is the social value of the revenue and externality gains from raising τ_i; the right-hand side is the welfare cost.

Substituting the standard Slutsky decompositions and solving the system gives, for the AI capital normal-return instrument,

$$\frac{\tau_{K_A}^{norm}}{1 - \tau_{K_A}^{norm}} = \frac{\lambda - \mu_{K_A} + \mathrm{MSC}_K(\alpha)/(r + \delta)}{\varepsilon_{K_A}^{own}} \cdot \mathrm{Y}_{K_A}, \tag{34}$$

where $\varepsilon_{K_A}^{own}$ is the compensated own-elasticity of AI capital supply, μ_{K_A} is the covariance of AI capital ownership with the social marginal utility of consumption (proportional to the joint distribution of ω and θ from Section 4.4), $\mathrm{MSC}_K(\alpha)$ denotes the money-metric marginal social cost (MSC) of AI deployment per unit of AI capital (formalized in Section 6.4), and

$$\mathrm{Y}_{K_A} \equiv 1 + \sum_{j \neq K_A} \frac{\tau_j}{1 - \tau_j} \cdot \frac{\varepsilon_{j,K_A}}{\varepsilon_{K_A}^{own}} \tag{35}$$

captures the cross-elasticity corrections, with ε_{j,K_A} the response of base j to the net-of-tax return on AI capital. Analogous expressions hold for the other instruments. The inverse-elasticity rule $\tau_i/(1 - \tau_i) = (\lambda - \mu_i)/\varepsilon_i$ is the special case $\mathrm{Y}_i = 1$ that obtains when cross-elasticities vanish. For AI capital, labor, and consumption the cross-elasticities are not zero: raising τ_{K_A} lowers the wage (Proposition 1) and hence erodes the labor tax base, so the corresponding term in (35) is negative, $\mathrm{Y}_{K_A} < 1$, and the revenue-leakage channel by itself places the optimal AI capital tax below the inverse-elasticity level. The Pigouvian and quasi-rent components work in the opposite direction, so the position of the total rate relative to that benchmark is a quantitative question. We adopt (34) as the working formula below.

6.3 Atkinson–Stiglitz and the Role of the VAT

The Atkinson–Stiglitz theorem[18] states that, under separable preferences between consumption and leisure and with an optimal nonlinear labor income tax in place, differentiated commodity taxation is dominated by the income tax alone. The theorem raises a question for the framework developed here: why should a uniform value-added tax (VAT) be part of the structure, rather than a richer labor income schedule that does the same distributional work?

Three considerations reconcile the VAT component of the structure with the Atkinson–Stiglitz benchmark. First, the theorem applies to differentiated consumption taxation; a broad-based VAT at a uniform rate falls outside its scope because a uniform commodity tax is equivalent to a uniform labor income tax (up to the treatment of savings). Second, the theorem assumes that the nonlinear income tax can capture all sources of income. In the economy we study, AI capital income accrues disproportionately to a small number of households and crosses international borders, so the labor-income-tax-only instrument cannot fully reach the capital base; a VAT captures consumption out of capital income that would otherwise escape. Third, the VAT's invoice-credit enforcement mechanism reaches consumption financed from income that escapes the income tax (informal earnings, unreported capital income), a base the income-tax instrument cannot reach; this enforcement complementarity is a standard second-best rationale for retaining a uniform VAT alongside a nonlinear income tax.

The VAT's role in the structure is therefore as a broad-based tax on aggregate consumption that captures a share of the AI capital return through capital owners' consumption expenditure, combined with a uniform rate that is consistent with the Atkinson–Stiglitz logic. When combined with UBI, a uniform VAT becomes progressive in net terms because UBI is a constant transfer while VAT liability scales with consumption.

In the model as computed, the consumption tax is not identified as a separate instrument. A uniform consumption tax lowers the real wage, so it carries the same wedge on hours and participation as an equal proportional increase in every tercile's labor rate; on the transfer it is a relabeling; and what remains is a levy on capital income that carries no supply wedge in the model. The computed problem therefore holds the consumption tax at zero (Appendix C), and the case for a VAT rests on the three considerations above, which lie outside the model.

6.4 Pigouvian Correction with Endogenous Externality

When AI capital deployment generates externalities (displaced labor, foregone human capital, disruption of employment-linked institutions, political instability), the optimal tax includes a Pigouvian component. Let $\mathrm{MSC}(\alpha)$ denote the marginal social cost of an increase in AI intensity, expressed in consumption (money-metric) units so that it is commensurable with revenue in the Ramsey system, and let $\mathrm{MSC}_K(\alpha) \equiv \mathrm{MSC}(\alpha) \cdot (d\alpha/dK_A)$ denote the corresponding cost per unit of AI capital, where $d\alpha/dK_A$ is the equilibrium mapping from the capital stock to AI intensity along the adoption path. The Pigouvian component is then

$$\tau_{K_A}^{\text{Pigou}} = \frac{\mathrm{MSC}_K(\alpha)}{r_A}, \tag{36}$$

[18] Atkinson and Stiglitz, "Design of Tax Structure."

which is dimensionless, as an ad valorem rate must be. A formal expression for the money-metric MSC in our setup is

$$\mathrm{MSC}(\alpha) = f(\theta^*) \left| \frac{d\theta^*}{d\alpha} \right| \cdot \left(\tau_L \, w \, \theta^{*\gamma} + b \right) + \chi_H(\alpha) + \chi_I(\alpha), \tag{37}$$

where the first term is the fiscal externality of displacement: the mass of workers at the margin times the per-worker fiscal cost, namely the lost labor-tax revenue plus the UBI outlay. The utility loss of the marginal displaced worker is itself second-order, because the participation condition $w\,\theta^{*\gamma} = w_r(b)$ makes the marginal worker indifferent between employment and the outside option; the first-order social cost of a marginal displacement is fiscal. The terms χ_H and χ_I, also expressed in consumption units, capture foregone human-capital investment (skills that would have been developed are not, reducing future productivity) and institutional externalities (disruption of employment-linked institutions such as health insurance, retirement saving, and civic engagement). Each component can in principle be calibrated; we treat their sum as an increasing function of α in the relevant range.

The monotonicity of $\mathrm{MSC}(\alpha)$ requires a qualification. The first term of (37) is proportional to $f(\theta^*)$, and for a single-peaked skill distribution f is not monotonic in θ^*. When θ^* is below the mode, $f(\theta^*)$ is increasing in α and MSC rises. As θ^* moves through the mode, $f(\theta^*)$ attains its maximum; beyond the mode, $f(\theta^*)$ declines and the first term of MSC declines with it. The second and third terms (χ_H, χ_I) are plausibly monotonic in α. The net monotonicity of MSC therefore depends on parameters; it is a conditional, not a universal, property.

6.5 Mirrlees Progressivity and the Participation Margin

For labor income, the optimal schedule is nonlinear. Following Saez,[19] the optimal top marginal rate is

$$\tau_L^{\mathrm{top}} = \frac{1 - \bar{g}}{1 - \bar{g} + a \cdot e}, \tag{38}$$

where $\bar{g}$ is the social marginal welfare weight on top earners, a is the Pareto parameter of the top income distribution, and e is the elasticity of taxable income. The formula describes the marginal rate on an unbounded Pareto tail. The computed model carries a different instrument, one rate for each tercile of the employed (Section 4.2), on a skill distribution that is bounded above and whose tail is thin: evaluated on the model's own tail, with a Pareto parameter of 27.2, $e = 0.25$, and $\bar{g} = 0$, the formula returns 13 percent, against a computed top-tercile rate of 87 percent at $\alpha = 0.65$ (Table 1). The formula is therefore reported as a benchmark and is not imposed. The computed top-tercile rate also exceeds $1/(1 + e)$, or 80 percent, the rate that maximizes the revenue of a flat rate at a given wage; the fixed-wage benchmark omits that the terciles supply one labor aggregate, so that fewer hours at the top raise the wage per efficiency unit of every worker.

The bottom of the schedule is shaped by the participation margin. Under our framework the extensive-margin response is

$$\frac{du}{d\tau_L^{bottom}} = f(\theta^*) \cdot \frac{d\theta^*}{d\tau_L^{bottom}} > 0, \tag{39}$$

so raising the bottom-tercile labor tax rate has a direct effect on non-participation through the participation margin. In the computed optimum the rate set for the bottom tercile, which is the

[19] Saez, "Using Elasticities."

rate the marginal worker faces, is at zero, the lower limit of the instrument set, at every AI intensity of Table 1 and in every year of the transition path, and the limit binds: the planner would choose a participation subsidy if the instrument set admitted one, as Section 7.9 shows.

6.6 The Optimal Structure in Steady State

Table 1 reports the computed static optimum (Appendix C) at five levels of AI intensity: the initial value, the three phase thresholds of Section 7.6, and the adoption ceiling, where the static optimum is the steady state of the dynamic problem of Section 7. The table separates AI capital into its normal-return and rent components, reflecting the supply structure of Section 4.3, and separates non-AI capital from AI capital, reflecting the three-factor production structure of Section 4.1. Three features of the computed structure are stated. First, the rent rate is at its upper limit of 100 percent and the normal-return rate at zero at every α: the Ramsey system places the AI tax on the rent base alone, and the effective rate on the AI return, 36 through 86 percent, is then the rent share of that return. Second, the labor income tax is progressive across terciles: the rate at the participation margin is zero, and the middle- and top-tercile rates lie between 54 and 76 percent and between 74 and 88 percent. Third, the consumption tax is absent, because the computed problem holds it at zero (Section 6.3).

Table 1: Computed Optimal Tax Structure by AI Intensity

Percent	$\alpha = 0.15$	$\alpha = 0.20$	$\alpha = 0.40$	$\alpha = 0.65$	$\alpha = 0.75$
AI capital, normal return, $\tau_{K_A}^{norm}$	0	0	0	0	0
AI capital, quasi-rent on Z, $\tau_{K_A}^{rent}$	100	100	100	100	100
AI capital, effective rate, τ_{K_A}	36	50	73	84	86
Labor income, bottom tercile	1	3	5	12	24
Labor income, middle tercile	54	57	71	76	69
Labor income, top tercile	74	76	83	87	88
Non-AI capital (fixed)	22.5	22.5	22.5	22.5	22.5
UBI, percent of per-capita GDP	23	26	39	51	54
Revenue, percent of output	40	43	54	63	64
Non-participation, percent	20	25	53	82	91

Note: Each column is the static optimum at the stated AI intensity α (Appendix C); the final column, at the adoption ceiling, is the steady state. Entries are percent, rounded to the whole percent except the non-AI capital rate, shown at one decimal. The effective AI rate is the AI tax as a share of the AI return $r_A K_A$. Labor income rates are the income-weighted average rate within each tercile of the employed; the rate set for the bottom tercile is zero at every α, and the entries of that row reflect the smooth joining of the tercile rates at the tercile boundary. The non-AI capital rate is fixed (Table C.1). Revenue is the sum of the labor, AI capital, and non-AI capital taxes. UBI = universal basic income; GDP = gross domestic product.

Revenue of 40 through 64 percent of output funds UBI of 23 through 54 percent of per-capita GDP and the other government spending of Table C.1. The structure follows the supply logic of the framework: the scarce complementary inputs, whose supply is inelastic, bear the AI tax in full; the normal return, whose taxation would reduce the AI capital stock, bears none; and labor at the participation margin is untaxed. The effective AI rates in Table 1 exceed the optimal robot taxes reported by Guerreiro, Rebelo, and Teles and by Thuemmel;[20] the differences trace to the quasi-rent component of AI capital returns, which is taxable without distorting the marginal investment, and to the scale of the revenue requirement.

7 The Dynamic Transition Path

The static analysis above characterizes optimal policy in the long-run steady state. This section develops the dynamic extension of the model, writes the planner's Hamiltonian in full, derives the first-order and co-state conditions, and characterizes the qualitative features of the optimal transition path.

7.1 The Dynamic Environment

We treat α as time-varying and characterize its evolution. AI intensity evolves according to

$$\frac{d\alpha_t}{dt} = \varphi\,(\alpha_{\max} - \alpha_t)\left[\left(\frac{w_t}{r_{A,t}}\right)^{\eta}(1-\tau_{K_A,t})^{\psi}\,\zeta_t\right], \tag{40}$$

where $\varphi > 0$ is the speed of adjustment, $\alpha_{\max} \leq 1$ the adoption ceiling (the productivity-weighted share of tasks AI can perform, in the sense of Section 3), $\eta > 0$ the wage–rental elasticity of AI

[20] Guerreiro, Rebelo, and Teles, "Should Robots Be Taxed?"; Thuemmel, "Optimal Taxation of Robots."

adoption, $\psi > 0$ the tax sensitivity, $\tau_{K_A,t}$ the effective tax rate on the AI return, and ζ_t an exogenous technological progress term. The wage–rental ratio carries units, so the law requires a normalization: the computed path holds the progress term constant at $\zeta = (r_{A,0}/w_0)^{\eta}$, the value at which the bracket equals one at the year-0 laissez-faire wage–rental ratio, so that φ is the year-0 adoption rate per unit of remaining scope. AI capital accumulates through investment:

$$\frac{dK_{A,t}}{dt} = I_{K,t} - \delta K_{A,t}, \tag{41}$$

with depreciation rate δ. Given the upward-sloping supply structure of Section 4.3, investment is chosen to satisfy the modified no-arbitrage condition

$$r_{A,t} - \tau_{K_A,t}^{norm}(r+\delta) = (r+\delta)(K_{A,t}/\bar{K})^{1/\varepsilon_K(\alpha_t)}, \tag{42}$$

the wedge condition (18) on the schedule (17) with the supply elasticity evaluated at the current AI intensity, so that the stock follows the schedule as the return moves and the quasi-rent $\Pi_{Z,t}$ accrues to the complementary input. In a model with adjustment costs on investment—which the specialized nature of model training, infrastructure, and organizational adaptation makes relevant for AI capital—anticipated future τ_{K_A} would also depress current investment through a Tobin's-q channel. We treat this additional channel as implicit in the subsequent discussion of announcement timing and commitment.

7.2 Endogenous Skill Formation

The skill distribution F_t is an object of policy choice over the relevant horizon. Let $H_t(\theta)$ denote aggregate human capital accumulated toward skill type θ, with $F_t(\theta)$ the implied cumulative distribution. Education and retraining investment $I_{F,t}$ shifts the distribution:

$$\frac{dF_t(\theta)}{dt} = -\kappa_F \left[F_t(\theta) - F^*(\theta; I_{F,t})\right], \tag{43}$$

where $F^*(\theta; I_{F,t})$ is the target distribution under current investment (first-order stochastic dominance rising with $I_{F,t}$) and $\kappa_F > 0$ is the adjustment speed. The adjustment speed is bounded by the biological and institutional constraints on human-capital formation (learning takes time; degree-granting institutions have capacity limits). The planner chooses $I_{F,t}$ alongside tax instruments, financing it from the same budget (32).

This addition matters for the optimum in two ways. First, the planner now faces a trade-off between current transfers (higher b) and future productive capacity (higher I_F), which is absent when F is exogenous. Second, Condition 5 of Section 4.7 becomes a statement about the upper bound on κ_F, rather than an assumption of exogenous slow adjustment; the fact that even policy-maximal κ_F is slow relative to $d\alpha/dt$ is the substantive content of Condition 5.

7.3 The Planner's Dynamic Problem

The planner maximizes

$$\max \int_0^{\infty} e^{-\rho t} W\big(\alpha_t, K_{A,t}, K_{N,t}, F_t, \tau_{L,t}, \tau_{K_A,t}, \tau_{C,t}, b_t\big)\, dt, \tag{44}$$

subject to (40), (41), (43), the government budget constraint (32) extended with I_F on the expenditure side, and the static labor-market equilibrium conditions of Section 4.5. The Hamiltonian is

$$\mathcal{H}_t = W(\cdot) + \lambda_{\alpha,t}\,\frac{d\alpha_t}{dt} + \lambda_{K_A,t}\,\frac{dK_{A,t}}{dt} + \lambda_{F,t}\,\frac{dF_t}{dt}, \tag{45}$$

where $\lambda_\alpha, \lambda_{K_A}, \lambda_F$ are the co-state variables corresponding to the three state variables. The sign of the co-state λ_α is not determined a priori: rising α lowers the wage and displaces labor, which reduces welfare, and raises output, which the transfer distributes. The co-state is negative where the displacement costs dominate and positive where the redistributed output gain does; Section 7.7 reports its sign in the computed path. The co-state λ_{K_A} is positive (AI capital raises output); the co-state λ_F is positive (higher human capital raises labor productivity and reduces displacement). Since F_t is a distribution-valued state, $\lambda_{F,t}$ is formally a co-state functional, and the product $\lambda_{F,t}\,dF_t/dt$ abbreviates the inner product $\int_0^1 \lambda_{F,t}(\theta)\,(\partial F_t(\theta)/\partial t)\,d\theta$; the finite-dimensional notation is retained for readability.

7.4 First-Order and Co-State Conditions

The first-order condition with respect to $\tau_{K_A,t}$ is

$$\frac{\partial W}{\partial \tau_{K_A,t}} + \lambda_{\alpha,t}\,\frac{\partial(d\alpha/dt)}{\partial \tau_{K_A,t}} + \lambda_{K_A,t}\,\frac{\partial(dK_A/dt)}{\partial \tau_{K_A,t}} = 0. \tag{46}$$

Substituting the Ramsey result (34) for $\partial W/\partial \tau_{K_A}$, the adoption response from (40), and the capital response from (41) and (42), this becomes

$$\tau^*_{K_A,t} = \tau^{\text{Ramsey}}_{K_A}(t) + \tau^{\text{Pigou}}_{K_A}(t) + \left(-\frac{\lambda_{\alpha,t}}{\lambda}\right)\cdot\left(-\frac{\partial(d\alpha/dt)}{\partial \tau_{K_A,t}}\right)\cdot\frac{1}{1+\varepsilon^{own}_{K_A}}, \tag{47}$$

where the final term is the dynamic correction. Its second factor is strictly positive (adoption falls when τ_{K_A} rises), so the correction takes the sign of $-\lambda_{\alpha,t}$: it raises the optimal rate above the static Ramsey–Pigouvian level where $\lambda_\alpha < 0$ and lowers it where $\lambda_\alpha > 0$. Section 7.7 reports the correction in the computed path as the difference between the path and the static optimum at the same AI intensity.

The co-state dynamics are given by the Hamiltonian conditions

$$-\frac{d\lambda_{\alpha,t}}{dt} + \rho\,\lambda_{\alpha,t} = \frac{\partial W}{\partial \alpha_t} + \lambda_{\alpha,t}\,\frac{\partial(d\alpha/dt)}{\partial \alpha_t}, \tag{48}$$

$$-\frac{d\lambda_{K_A,t}}{dt} + \rho\,\lambda_{K_A,t} = \frac{\partial W}{\partial K_{A,t}} + \lambda_{K_A,t}\,\frac{\partial(dK_A/dt)}{\partial K_{A,t}}, \tag{49}$$

$$-\frac{d\lambda_{F,t}}{dt} + \rho\,\lambda_{F,t} = \frac{\partial W}{\partial F_t} + \lambda_{F,t}\,\frac{\partial(dF/dt)}{\partial F_t}. \tag{50}$$

The transversality conditions are $\lim_{t\to\infty} e^{-\rho t}\lambda_{i,t}\cdot \text{state}_{i,t} = 0$ for $i \in \{\alpha, K_A, F\}$. Together with the adoption law (40), capital accumulation (41), and skill formation (43), these constitute the dynamic optimality system. The system is a set of coupled partial differential equations in the state and co-state variables that admits no closed-form solution for general functional specifications of W, g, and F^*; the numerical solution of Section 7.7 adopts the functional specifications listed in Appendix C and solves the planner's problem directly.

7.5 The Deceleration Component as a Congestion Externality

A deceleration term captures the welfare consequences of AI adoption outpacing institutional adjustment; this subsection derives it from optimality conditions. Let ι_t denote the rate at which distribution institutions (retraining programs, sectoral reallocation infrastructure, UBI administrative capacity) can absorb displaced workers. Suppose that when $d\alpha_t/dt > \iota_t$, a congestion externality is generated, captured by an additional term in social welfare:

$$\chi_{decel}(d\alpha/dt, \iota) = \frac{\kappa}{2} \cdot \max\Big(\frac{d\alpha}{dt} - \iota,\, 0\Big)^2, \tag{51}$$

with $\kappa > 0$ a congestion parameter. This term enters the instantaneous welfare function with a negative sign. Differentiating the Hamiltonian with respect to τ_{K_A}, the congestion term contributes the marginal benefit

$$-\kappa\, \max\Big(\frac{d\alpha}{dt} - \iota,\, 0\Big) \cdot \frac{\partial(d\alpha/dt)}{\partial \tau_{K_A}} = \kappa\, \max\Big(\frac{d\alpha}{dt} - \iota,\, 0\Big) \cdot \varphi\,(\alpha_{\max} - \alpha_t) \left(\frac{w_t}{r_{A,t}}\right)^{\eta} \psi\,(1-\tau_{K_A,t})^{\psi-1}\, \zeta_t \;>\; 0, \tag{52}$$

using the adoption law (40). Setting the full first-order condition to zero and dividing this marginal benefit by the marginal net-revenue coefficient of the instrument (the bracketed revenue term of (33)) converts it into an additive component of the optimal rate, $\tau_{K_A}^{\text{decel}}$, proportional to $\kappa \cdot \max(d\alpha/dt - \iota,\, 0)$ to first order in the congestion gap. The deceleration component is thus derived from optimality conditions rather than imposed. The total optimal AI capital tax rate is

$$\tau_{K_A}^{\text{total}} = \tau_{K_A}^{\text{Ramsey}} + \tau_{K_A}^{\text{Pigou}} + \tau_{K_A}^{\text{decel}}. \tag{53}$$

7.6 The Four Phases of Optimal Policy

The optimal transition path is read in four phases, distinguished by the level of AI intensity and the state of the labor market. In the calibration of Section 7.7, the phase thresholds are $\underline{\alpha} = 0.20$, $\bar{\alpha} = 0.40$, and $\bar{\bar{\alpha}} = 0.65$; the computed path reaches the first two in years 2 and 17 and the third in year 89, beyond the 30 years reported. These thresholds are properties of the calibration rather than optimal-control features: the computed solution is smooth, and the phases describe its course without implying switches in policy.

7.6.1 Phase I: Pre-Displacement

In the first phase, $\alpha < \underline{\alpha}$, and AI intensity is low. In the computed path the transfer is in place from the start, at 19 percent of per-capita GDP at year 0, financed by the labor income tax; non-participation is 16 percent, against a laissez-faire rate of 5 percent; and the AI tax is zero at year 0 and is phased in within the phase.

7.6.2 Phase II: Early Displacement

In the second phase, $\underline{\alpha} \le \alpha < \bar{\alpha}$. AI intensity has risen, and non-participation runs from 21 to 50 percent over the phase. The rent rate rises from 44 to 95 percent over the phase, and UBI rises from 22 to 38 percent of per-capita GDP. The labor rate at the participation margin stays at zero, since an increase would compress participation through the extensive margin. Retraining investment I_F is zero in this phase, as it is in every year of the path: its lower limit binds, so in this calibration a unit of output raises welfare by more as a transfer than as a shift of the skill distribution.

7.6.3 Phase III: Mass Displacement

In the third phase, $\bar{\alpha} \leq \alpha < \bar{\bar{\alpha}}$, and AI intensity is high; over the reported years of this phase non-participation runs from 52 to 64 percent and the labor share from 30 down to 22 percent. The rent rate reaches its limit of 100 percent in year 24 and stays there, and the normal-return rate is zero; the effective AI rate runs from 70 to 78 percent, and from year 24 it moves only with the rent share of the AI return; the top-tercile labor rate runs from 85 to 87 percent and b_t from 39 to 44 percent of per-capita GDP.

7.6.4 Phase IV: AI-Saturated Steady State

In the fourth phase, $\alpha \geq \bar{\bar{\alpha}}$, and AI intensity approaches its ceiling. Most routine cognitive work is automated; employment is concentrated in AI-complementary high-skill roles and in AI-resistant sectors. The rates converge to the steady state, the final column of Table 1, with non-participation of 91 percent, UBI of 54 percent of per-capita GDP, and an effective AI rate of 86 percent. In this steady state $d\alpha/dt \to 0$, the link between the pace of adoption and the optimal rates is inactive, and policy concerns long-run redistribution and the insurance function of UBI.

7.7 A Calibrated Transition Path

To make the dynamic analysis concrete, we present a calibrated transition path. The parameterization is $\sigma = 1.8$ (AI and labor are gross substitutes), $\nu = 0.6$ (outer-nest substitutability between AI–labor composite and non-AI capital), $\xi_X = 0.70$ (outer-nest weight on the AI–labor composite), $\gamma = 1.2$ (slight skill super-linearity), $\varphi = 0.05$ (adoption of 5 percent of the remaining scope per year at year 0), $\eta = 0.5$ (moderate wage–rental sensitivity), $\psi = 0.3$ (moderate tax sensitivity of adoption), $\rho = 0.03$ (social discount rate), $\delta = 0.10$ (AI capital depreciation reflecting rapid obsolescence), $\varepsilon_{K_A} = 0.5$ at α_0 declining to 0.3 at $\bar{\bar{\alpha}}$, $\kappa_F = 0.05$ (human-capital adjustment speed), $\xi_F = 3.5$ (the shift of the location of the skill distribution per unit of retraining spending over output), and $\alpha_0 = 0.15$ (illustrative starting value). The skill distribution is logit-normal with mode at $\theta = 0.4$.

Three of these values are compared with evidence. In the model a dollar of retraining raises earnings in its first year by $\kappa_F \, \xi_F \, e_\mu$ dollars, where $e_\mu = 0.529$ is the response of output to the location of the skill distribution at year 0 of the path, prices given; at the parameterization above, this first-year gain is 9.3 cents. Evaluations of retraining for displaced workers at its full social cost report ratios of benefits to costs between 1.07 and 1.71, each the present value of earnings gains per dollar of cost. Each source assumes that the gain persists over the remaining working life it considers, 25 or 20 years. Under that assumption the ratios correspond to a first-year gain between 6.8 and 11.5 cents[21] and, at $\kappa_F = 0.05$, to ξ_F between 2.6 and 4.3. The speed κ_F is the rate at which retrained skill leaves the workforce: exit over those working lives and the depreciation

[21] Louis Jacobson, Robert LaLonde, and Daniel G. Sullivan, "Estimating the Returns to Community College Schooling for Displaced Workers," IZA Discussion Paper 1017 (Bonn: Institute for the Study of Labor, 2004), 25, 31, published in *Journal of Econometrics* 125, no. 1–2 (2005): 271–304: a social benefit-to-cost ratio of roughly 1.07 for the mix of community college schooling completed by displaced workers, for a participant who works 25 further years, with gains discounted at 4 percent; Louis S. Jacobson, Robert J. LaLonde, and Daniel G. Sullivan, "Policies to Reduce High-Tenured Displaced Workers' Earnings Losses Through Retraining," Hamilton Project Discussion Paper 2011-11 (Washington, DC: Brookings Institution, 2011), 15, 34: total benefits of two years of training while working full time against total social costs without and with child-care expenses, ratios of 1.71 and 1.21, for a worker of forty with 20 further working years, discounted at 3 percent. The first-year gain is each ratio divided by the present value of one dollar a year over the source's horizon at its discount rate; the sources assume that the gain persists for the remaining working life.

of human capital while employed, which the literature puts between zero and 2.7 percent a year, together give 4.0 through 7.7 percent a year.[22] The adoption ceiling $\alpha_{\max} = 0.75$, 87.8 percent of tasks under equation (5), is a scenario: no source estimates a long-run ceiling on the share of tasks AI performs. Measures of current technical exposure, which count tasks whose time AI could reduce rather than tasks it performs, put between 47 and 56 percent of tasks within reach of large language models with software built on them[23] and between 60 and 70 percent of employees' time in activities current technologies could automate.[24] The lower ceilings 0.53, 0.59, and 0.62, or 55, 66, and 71 percent of tasks, are solved as variants (Appendix C).

The path solves the dynamic planner's problem of Section 7.3 by simultaneous direct collocation (Appendix C), in the following form. The static equilibrium conditions of Section 4.5 hold at every date as algebraic constraints, with the AI capital stock on its supply schedule (42). The adoption law (40) holds as a defect constraint between dates. In place of the law (43) for the whole skill distribution, a law for its location μ_t holds as a defect constraint: the location adjusts toward the target that retraining sets, $d\mu_t/dt = \kappa_F\,(\mu_0 + \xi_F I_{F,t}/Y_t - \mu_t)$, where μ_0 is the year-0 location. The government budget holds at every date, the horizon is finite with a stationary terminal value (Appendix C), and the instruments are chosen jointly over the horizon rather than date by date. All 11 initializations of the problem reach one optimum.

The computed optimal path features α rising from 0.15 to 0.49 over 30 years, at a pace that slows from 0.032 to 0.005 a year as the remaining scope narrows. Non-participation rises from 16 percent at year 0 to 64 percent at year 30, in every year and without a plateau; the year-0 value already exceeds the laissez-faire rate of 5 percent because optimal UBI lifts the reservation wage, so the series measures displacement and policy-induced non-participation together. The labor share falls from 61 to 22 percent. UBI rises from 19 to 44 percent of per-capita GDP. The effective AI capital tax rises from zero at year 0 to 43 percent at year 5 and 78 percent at year 30; the normal-return rate is zero throughout, and the rent rate rises from zero at year 0 to its limit of 100 percent in year 24 and stays there. The top-tercile labor rate is 80 percent at year 0 and 87 percent at year 30, and the middle-tercile rate rises from 58 to 77 percent; the rate set for the bottom tercile, which the marginal worker faces, is zero throughout. Halving the mesh and doubling the horizon move each reported rate by at most 1.5 percentage points and non-participation by at most 0.1 percentage points (Appendix C). Figure 1 visualizes the trajectory.

[22] Jacobson, LaLonde, and Sullivan, "Estimating the Returns," 31, and "Policies to Reduce," 15, for the remaining working lives; James J. Heckman, Lance Lochner, and Christopher Taber, "Explaining Rising Wage Inequality: Explorations with a Dynamic General Equilibrium Model of Labor Earnings with Heterogeneous Agents," *Review of Economic Dynamics* 1, no. 1 (1998): 19, who set the depreciation of human capital to zero; Rodolfo E. Manuelli and Ananth Seshadri, "Human Capital and the Wealth of Nations," *American Economic Review* 104, no. 9 (2014): 2746, Table 1, whose calibrated depreciation rate of human capital is 0.0270.

[23] Tyna Eloundou et al., "GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models" (working paper, arXiv:2303.10130v5, August 2023), abstract and Table 3, published as "GPTs Are GPTs: Labor Market Impact Potential of LLMs," *Science* 384, no. 6702 (2024): 1306–1308. Their measure counts a task as exposed where large language models, with software built on them, could reduce the time it takes by at least half at the same quality.

[24] Michael Chui et al., *The Economic Potential of Generative AI: The Next Productivity Frontier* (McKinsey & Company, June 2023), 3.

Table 2: Computed Optimal Tax Path

Year	α	u (%)	τ_{K_A} (%)	$\tau_{K_A}^{\text{rent}}$ (%)	τ_L^{top} (%)	b/y (%)
0	0.15	16	0	0	80	19
5	0.26	29	43	71	83	27
10	0.34	39	55	81	85	32
15	0.39	49	67	94	85	37
20	0.43	55	73	98	85	40
25	0.46	60	76	100	86	43
30	0.49	64	78	100	87	44

Note: τ_{K_A} is the effective tax rate on the AI return, the AI tax as a share of $r_A K_A$; $\tau_{K_A}^{\text{rent}}$ is the rate on the quasi-rent; the normal-return rate is zero in every year. τ_L^{top} is the income-weighted average labor income tax rate of the top tercile of the employed. b/y is UBI as a percentage of per-capita GDP. The path is the direct-collocation solution of the planner's problem of Section 7.3 (Appendix C); α is a share; the tax-rate, non-participation, and UBI columns are percent, rounded to the whole percent. UBI = universal basic income; GDP = gross domestic product.

The series of Figure 1 measure different quantities: tax rates, the share of the population not participating, UBI as a share of per-capita GDP, and AI intensity, a share scaled by 100 for display; their crossings have no economic meaning.

First, the effective AI tax rises in every reported year and does not peak within them: the rent rate climbs to its limit by year 24, and thereafter the effective rate rises with the rent share of the AI return, which moves from 36 percent at year 0 to 78 percent at year 30. Second, the path taxes AI capital less than the static optimum at the same AI intensity until the rent rate reaches its limit. The static optimum at the AI intensities of years 0, 5, and 10 carries effective rates of 36, 60, and 68 percent, which exceed the path's rates by 36, 17, and 13 percentage points. The difference is the dynamic term of equation (47) together with the congestion term of Section 7.5. The co-state on AI intensity is positive at year 0: a marginally higher initial AI intensity raises the planner's objective. The planner therefore lowers the rate to speed adoption, although adoption outpaces the absorption rate ι in years 0 through 3 and the congestion term works in the opposite direction. Third, non-participation does not plateau within the reported years: it rises in every year, toward the steady-state value of 91 percent in the final column of Table 1. Retraining investment is zero throughout the path.

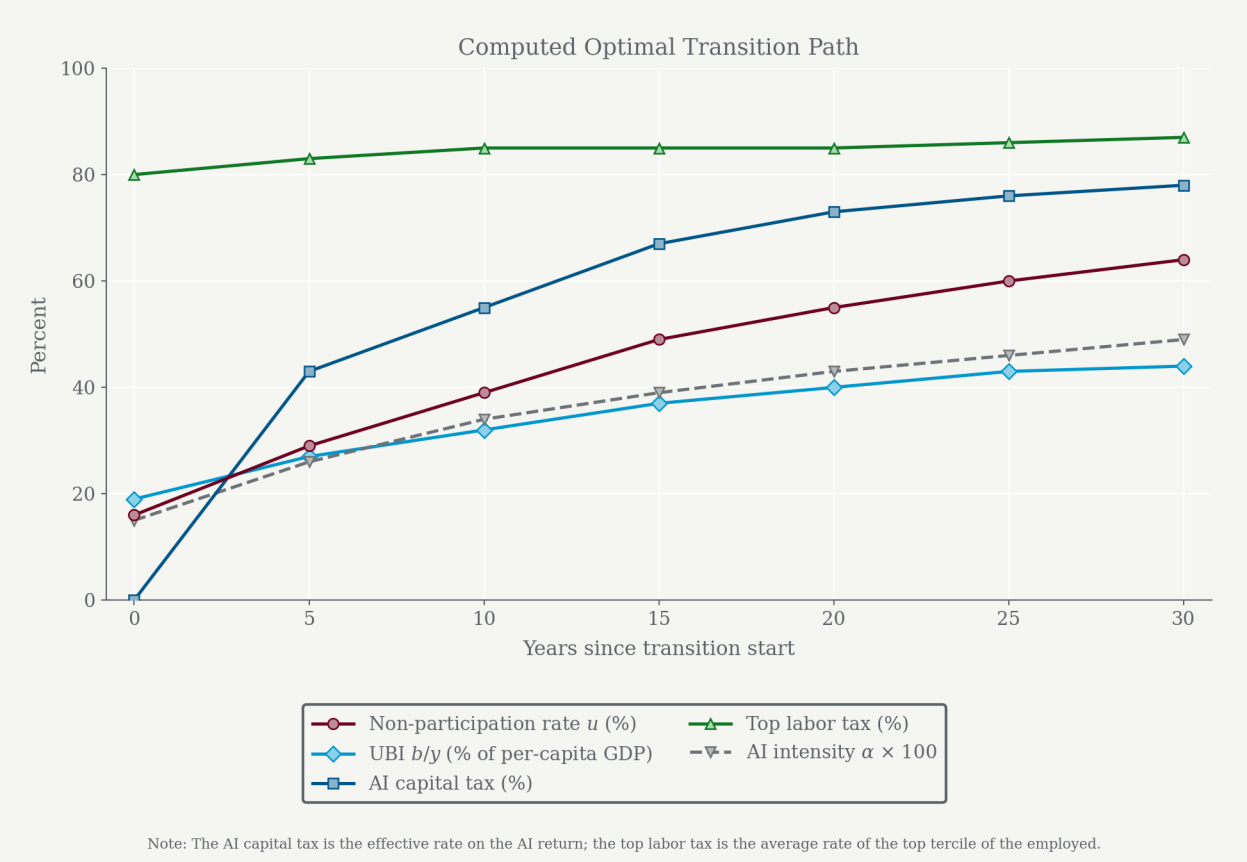


Figure 1: Computed Optimal Transition Path, Years 0 through 30. Horizontal axis: years since the start of the transition; vertical axis: percent. Non-participation rate u (maroon, circles); UBI b/y as a percent of per-capita GDP (bright blue, diamonds); AI capital tax rate τ_{K_A} (deep blue, squares); top labor income tax rate (green, upward triangles); AI intensity α_t, a share, times 100 (gray, dashed, downward triangles). Marker fills are pastel tints of the line hues. Source: author's calculations; values as in Table 2.

7.8 Sensitivity to Four Parameters

The computed path of Section 7.7 rests on a particular parameter vector, and several of those parameters are difficult to pin down empirically. Figure 2 documents how the path depends on four of them, one at a time: the elasticity of substitution σ, the tax-sensitivity of adoption ψ, the AI capital supply elasticity ε_K, and the return to retraining ξ_F. The baseline is $\sigma = 1.8$, $\psi = 0.3$, $\varepsilon_K = 0.5$, and $\xi_F = 3.5$. Each panel re-solves the planner's problem with one parameter moved across the range on its horizontal axis and the others at baseline, and plots three outcomes at year 30 of the re-solved path: non-participation, UBI as a percentage of per-capita GDP, and the effective AI capital tax rate. The three series within a panel are in different units (rates and a GDP share), so their crossings have no economic meaning. The steady state is the static optimum at the adoption ceiling (Section 6.6); ψ and ξ_F do not enter it, so its sensitivity is reported in the text for σ and ε_K alone.

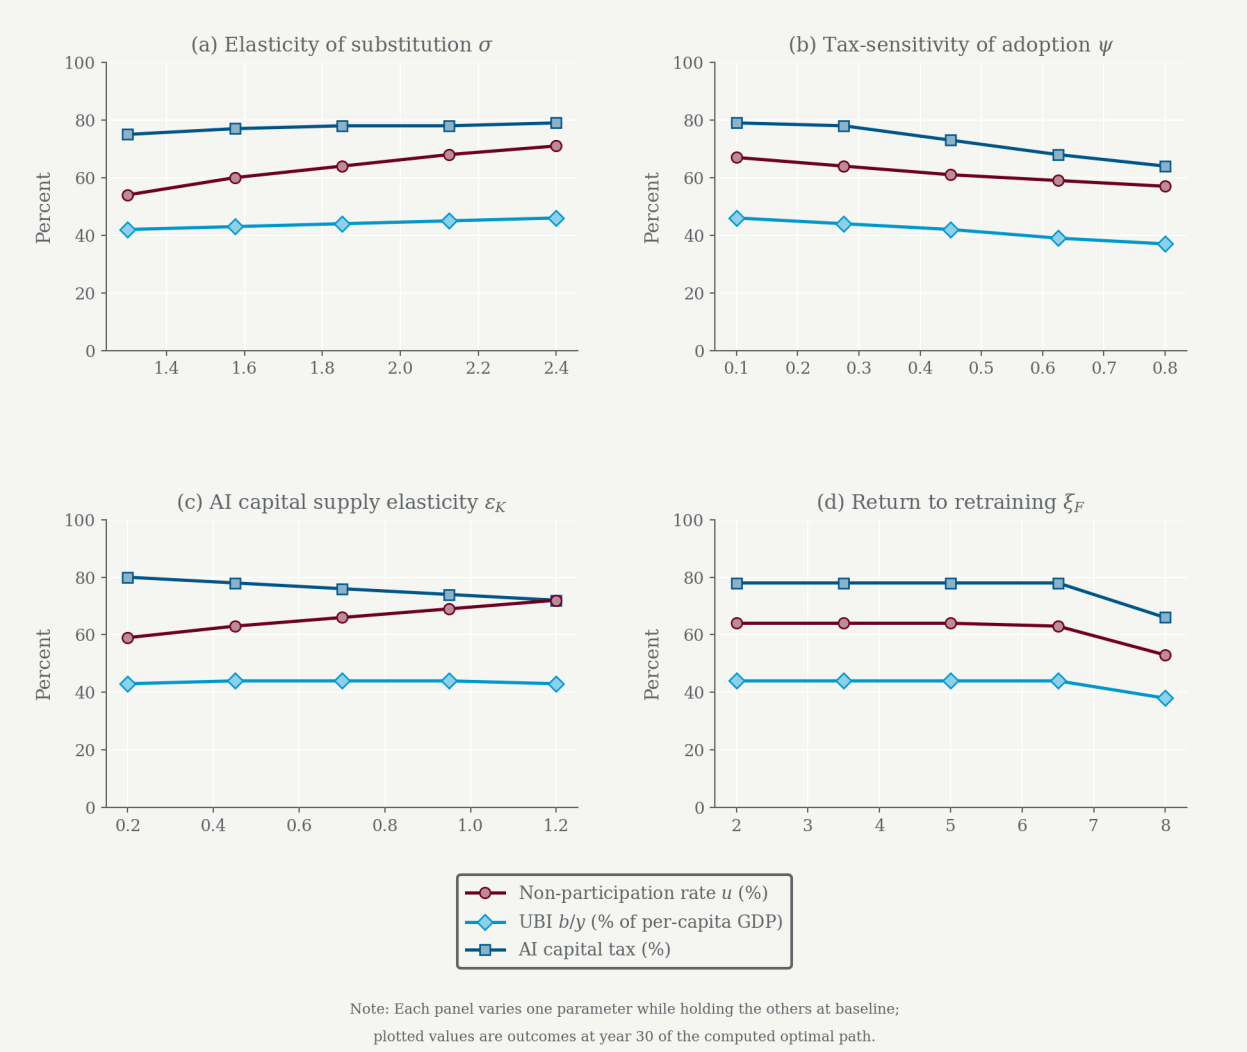


Figure 2: Sensitivity of the Computed Path at Year 30 to Four Parameters. Panels vary (a) σ, (b) ψ, (c) ε_K, and (d) ξ_F over the range on each horizontal axis, others at baseline; vertical axes: percent, common scale. Non-participation rate (maroon, circles); UBI as a percent of per-capita GDP (bright blue, diamonds); AI capital tax rate (deep blue, squares); marker fills are pastel tints. Source: author's calculations.

The computed sensitivities deliver four conclusions. Panel (a): non-participation at year 30 rises from 54 percent at $\sigma = 1.3$ to 71 percent at $\sigma = 2.4$, while UBI rises from 42 to 46 percent of per-capita GDP and the effective AI tax from 75 to 79 percent; in the steady state the same variation moves non-participation from 81 to 96 percent. The displacement result holds across the range, and its size depends on σ. Panel (b): the more adoption responds to the tax, the lower the rate the planner sets and the slower the transition. From $\psi = 0.1$ to $\psi = 0.8$ the effective AI tax at year 30 falls from 79 to 64 percent, non-participation from 67 to 57 percent, and UBI from 46 to 37 percent of per-capita GDP. Panel (c): a more elastic AI capital supply raises non-participation at year 30 from 59 to 72 percent and lowers the effective AI tax from 80 to 72 percent, because elastic supply shrinks the rent share of the AI return, consistent with Section 4.3; UBI stays between 43 and 44 percent of per-capita GDP. In the steady state the same variation moves non-participation from 89

to 93 percent and the effective AI tax from 88 to 80 percent. Panel (d): at the baseline return and below it, retraining is not used and the path does not depend on ξ_F. Retraining enters the optimum between $\xi_F = 5.47$ and $\xi_F = 5.56$, a first-year gain between 14.5 and 14.7 cents on the dollar, above the gain between 6.8 and 11.5 cents that the evidence of Section 7.7 implies. At $\xi_F = 6.5$ retraining is used in years 0 through 3, at up to 6 percent of output, and non-participation at year 30 is 63 percent against 64 percent at the baseline; at $\xi_F = 8.0$ retraining is used through year 22 at up to 18 percent of output a year, and non-participation at year 30 falls to 53 percent, with UBI and the effective AI tax also lower. The return to retraining therefore decides whether retraining enters the optimum, and the baseline value lies below the threshold. Where the computed optimum uses retraining, its peak share of output runs from 0.4 percent just above the threshold to 18.3 percent at the highest return considered, against public spending on training of 0.01 through 0.34 percent of GDP across the 17 countries of one tabulation;[25] the returns of the evidence were measured at that scale of spending, and the retraining technology of the model is linear in spending.

7.9 An Earnings Subsidy at the Participation Margin

The instrument set of the computed solution confines every tax rate to the interval from zero through 100 percent, and the rate at the participation margin is at its lower limit in every year of the path (Section 6.5). This subsection reports the problem with that one rate free below zero, so that the planner can subsidize earnings at the participation margin, the other rates keeping their limits. The extended problem has more than one local optimum. Of four initializations, two converge to two distinct optima, one each, and the rest do not converge. Neither optimum is established as the global one, so the results are stated as bounds over the optima found (Section C.7).

Over those optima, non-participation lies between 1.0 and 4.4 percent in every year from 0 through 30, against 16 through 64 percent along the path of Section 7.7. The rate set at the participation margin lies between -251 and -45 percent, UBI between 7 and 22 percent of per-capita GDP, and the effective AI tax at year 30 between 66 and 77 percent, with the normal-return rate at zero. The planner's objective exceeds that of the path of Section 7.7 by the equivalent of 19.7 through 23.5 percent of consumption.

Two cautions qualify these results. First, the welfare function counts the consumption of a household that does not participate and attaches no value to the outside option against which the participation condition of Section 4.2 compares earnings; a subsidy that draws workers into employment is valued accordingly, and a criterion that valued time out of work would weigh it differently. Second, why non-participation settles at the low level the optima share is not established here. The results of the paper are therefore stated for the instrument set that excludes subsidies.

7.10 Decomposing the Welfare Gain from Optimal Policy

The welfare gain from optimal policy, relative to laissez-faire, decomposes across the mechanisms the paper has introduced as follows. Let W^{opt} denote social welfare under the optimal path of

[25] Japan Institute for Labour Policy and Training, *Databook of International Labour Statistics 2025* (Tokyo: Japan Institute for Labour Policy and Training, 2025), 263, Table 9-6, from Organisation for Economic Co-operation and Development (OECD) data on public expenditure on labour market programmes, fiscal year 2021 (2020 for Italy).

Section 7.7 and W^{lf} welfare under laissez-faire (no UBI, conventional tax rates, no Pigouvian correction). The difference admits the decomposition

$$W^{\text{opt}} - W^{\text{lf}} = \underbrace{\Delta W^{\text{redis}}}_{\text{redistribution}} + \underbrace{\Delta W^{\text{ins}}}_{\text{insurance}} + \underbrace{\Delta W^{\text{pig}}}_{\text{Pigouvian correction}} + \underbrace{\Delta W^{\text{decel}}}_{\text{deceleration / congestion}}, \tag{54}$$

where ΔW^{redis} is the mean-income transfer from AI capital owners to displaced workers (the standard Ramsey–Mirrlees gain), ΔW^{ins} is the insurance gain identified in Section 5, ΔW^{pig} is the internalization of the displacement externality MSC(α) from Section 6.4, and ΔW^{decel} is the congestion gain from slowing adoption when institutional adjustment lags (Section 7.5). The decomposition is an accounting of mechanisms. The computed solution maximizes their sum and does not separate the four terms, so no shares are reported. Two statements follow from the definitions of the terms and the computed path. The Pigouvian and deceleration flow components vanish as $d\alpha/dt \to 0$, so they accrue during the transition. And the congestion term is active only in years 0 through 3 of the computed path, the years in which adoption outpaces the absorption rate.

7.11 A Comparative-Statics Summary

Table 3 collects the signs of the main comparative-statics responses, as computed. The rows for α, σ, ε_K, and ownership concentration are finite differences of the static optimum at $\alpha = 0.65$, with the policy re-optimized. The rows for ψ and ξ_F, which do not enter the static optimum, give the direction of the year-30 values of the transition path over the grids of Section 7.8; a response that is not monotone over its grid is marked n.s. The wage entry of the first row agrees with Proposition 1.

Table 3: Comparative-Statics Summary: Sign of Endogenous Responses

	$\partial w^*/\cdot$	$\partial u^*/\cdot$	$\partial s_L/\cdot$	$\partial \tau^*_{K_A}/\cdot$	$\partial b^*/\cdot$
AI intensity α	−	+	−	+	+
Substitution elasticity σ	−	+	−	+	+
Tax-sensitivity ψ	n.s.	−	+	−	−
Supply elasticity ε_K	+	+	−	−	+
Return to retraining ξ_F	−	−	+	−	n.s.
Ownership concentration	+	−	−	−	+

Note: Entries report the sign of the response of the column variable to an increase in the row variable, other parameters at baseline: for α, σ, ε_K, and ownership concentration, at the static optimum at $\alpha = 0.65$ with the policy re-optimized; for ψ and ξ_F, at year 30 of the transition path over the grids of Figure 2. w^* is the equilibrium wage, u^* the non-participation rate, s_L the labor share, $\tau^*_{K_A}$ the optimal effective AI capital tax rate, b^* the optimal UBI level. Ownership concentration is the exponent of the ownership schedule (Table C.1). n.s. marks a response that is not monotone over the grid.

7.12 Hysteresis and Irreversibility

The dynamic path has the property of partial irreversibility. AI capital, once installed, does not uninstall itself when tax rates change. The state variable α evolves through investments in models, infrastructure, and organizational adaptation with sunk-cost components. This has an implication for the timing of tax policy announcements. Under the instantaneous no-arbitrage condition (42), only the current τ_{K_A} affects the current investment decision; in a model with adjustment costs on investment, any anticipated future increase in τ_{K_A} depresses current investment through

a Tobin's-q channel. Under either reading, the policy design question is about the sequencing of distortion over time, not its elimination: a commitment to raise τ_{K_A} only after a specified level of adoption has been reached extracts revenue from the eventual stock with minimum distortion on pre-commitment investment and concentrated distortion on post-commitment investment—a profile that is welfare-preferred when displacement is the primary externality. The credibility of such a commitment is the binding constraint.

Partial irreversibility also implies that delayed policy response creates welfare losses that cannot be recovered ex post. If UBI implementation lags AI displacement, the consumption shortfalls experienced by displaced workers during the gap are not reversed by later transfers. Under the expected-utility welfare criterion of Section 5, the cumulative welfare loss grows more than proportionally with the length of the delay: each additional year of lag adds another cohort of displaced workers and the per-worker loss rises with α through the MSC schedule.

8 The International Dimension

A single-country analysis assumes AI capital taxation can be set independently of tax choices elsewhere. In practice AI capital is mobile across jurisdictions. This section develops a two-country extension that formalizes the constraint imposed by capital mobility and identifies the welfare gains from coordination.

8.1 A Two-Country Model

Consider two countries $\mathcal{A}$ and $\mathcal{B}$, each with the domestic production technology (6)–(7). AI capital is mobile across the two countries subject to frictions summarized by an adjustment cost $\Psi(K_A^i - K_A^{*i})$ that is zero at the frictionless allocation K_A^{*i} and rises quadratically away from it. Labor and non-AI capital are immobile. The no-arbitrage condition for AI capital across countries is

$$(1 - \tau_{K_A}^{\mathcal{A}})\, r_A^{\mathcal{A}} - \Psi'(K_A^{\mathcal{A}} - K_A^{*\mathcal{A}}) = (1 - \tau_{K_A}^{\mathcal{B}})\, r_A^{\mathcal{B}} - \Psi'(K_A^{\mathcal{B}} - K_A^{*\mathcal{B}}), \tag{55}$$

which pins down the allocation of AI capital between the two countries given tax rates. The quasi-rent term Π_Z^i / K_A^i is now country-specific because the complementary inputs Z^i (compute, electricity, specialized labor) are geographically fixed.

Scarce complementary inputs shape this extension: when $Z^{\mathcal{A}}$ and $Z^{\mathcal{B}}$ are distinct, the two countries have different quasi-rent structures, and the allocation of AI capital between them reflects both tax rates and complementary-input availability. This removes the boundary case of perfectly elastic supply (under which all capital flows to the lowest-tax jurisdiction) and gives each country a nontrivial tax base.

8.2 Nash Equilibrium in AI Capital Taxes

Each country's planner chooses $\tau_{K_A}^i$ to maximize domestic social welfare, taking the other country's rate as given. Capital mobility alters the closed-economy first-order condition in two ways: it raises the effective supply elasticity of the domestic base (an outflow margin), and it makes each country's revenue depend on the other country's rate. A reduced form for the resulting best response is

$$\tau_{K_A}^{i,\mathrm{BR}}(\tau_{K_A}^{-i}) = (1 - m)\, \tau_{K_A}^{i,\mathrm{closed}} + m\, \tau_{K_A}^{-i} - \Delta(m), \tag{56}$$

where $\tau_{K_A}^{i,\text{closed}}$ is the closed-economy optimum derived in Section 6, $m \in (0,1)$ is the capital-mobility sensitivity arising from the substitution channel in (55), and $\Delta(m) > 0$, increasing in m, is the level shift induced by the outflow margin: the marginal revenue lost to capital relocation at any common rate, which each country's first-order condition internalizes while the corresponding revenue gain abroad is external to it. In the symmetric Nash equilibrium, $\tau_{K_A}^{\text{Nash}} = \tau_{K_A}^{\text{closed}} - \Delta(m)/(1-m) < \tau_{K_A}^{\text{closed}}$, with the gap increasing in m and decreasing in the tightness of the complementary-input constraint (higher Π_Z/K_A shrinks Δ, because the tax falls partly on the immobile rent and the outflow margin is correspondingly weaker).

The size of the gap in the calibrated model is computed in Section 8.3.

8.3 Destination-Based Taxation and Coordination

Two corrective mechanisms partially restore welfare losses from tax competition. The first is destination-based taxation: taxing AI services where they are consumed rather than where they are produced. Under destination-based taxation, the tax base becomes location-independent (analogous to VAT on goods), and the mobility sensitivity m and the level shift $\Delta(m)$ in (56) fall toward zero. Second is minimum-rate coordination: a multilateral agreement establishing $\tau_{K_A}^{i} \geq \tau_{K_A}^{\min}$ for all participating countries. Building on the Organisation for Economic Co-operation and Development (OECD) 15 percent corporate minimum tax, such coordination establishes a floor that prevents the race-to-the-bottom outcome of (56).

The coordination gain is

$$\Delta W = \sum_i W^i(\tau_{K_A}^{i,\text{coop}}) - \sum_i W^i(\tau_{K_A}^{i,\text{Nash}}) > 0, \tag{57}$$

and depends on the gap between the uncoordinated rate and the cooperative one. The computed two-country solution (Appendix C) gives the magnitude. With a mobility share of $m = 0.5$, the symmetric Nash equilibrium carries a mean effective AI tax of 36 percent over the reported years, 23 percentage points below the mean effective AI tax of the cooperative optimum. At year 30 the effective AI tax is 57 percent under Nash against 78 percent under cooperation, and UBI is 34 percent of per-capita GDP against 44 percent. Lower taxation of AI capital also speeds adoption: AI intensity at year 30 is 0.51 under Nash against 0.49. The welfare loss of the Nash outcome is equivalent to a reduction of 2.1 percent in every household's consumption at every date of the cooperative path. The practical implication is that the financing structure derived here depends on international coordination of AI capital taxation.

9 Implementation Challenges

This section collects the matters that separate the optimal structure from an enacted one: measurement of the tax base, political economy, dimensions of welfare beyond income, endogenous technical change, demographic heterogeneity, and the scope of the framework.

9.1 Measurement and Definition

Operationalizing the terms "AI capital" and "AI rents" in the tax code requires workable definitions. Cloud computing revenue is partially AI-related, partially not. Software companies enhance their products with AI to varying degrees. Traditional firms deploy AI internally without generating distinct AI revenue. Any definition invites reclassification of activity across its boundary.

A practical approach combines sector codes, activity tests (what fraction of value-added comes from AI-specific processes), and safe-harbor thresholds for small players. An alternative is to tax the AI production function's inputs directly, such as compute and specialized chips, rather than the outputs; input-based taxation is easier to enforce but introduces distortions of its own. The quasi-rent tax on complementary inputs Z introduced in Section 4.3 is suited to input-based implementation because compute capacity, electricity, and specialized labor markets all generate measurable flows.

9.2 Political Economy

A tax structure raising 40 through 64 percent of output requires sustained political support. UBI at 23 percent of per-capita GDP or more is a standing political target whose generosity will be contested in every election, creating uncertainty that undermines its insurance value (cf. Section 5). Embedding UBI in constitutional or supermajority-protected institutions, analogous to central bank independence, could stabilize the system but raises its own concerns about democratic legitimacy.

The political economy of AI capital taxation is challenging for a separate reason. AI owners are concentrated, wealthy, and politically influential. The historical pattern of capital taxation in advanced economies, namely declining effective rates over decades despite rising inequality, suggests that political pressure toward lower AI capital taxation would persist. Maintaining the theoretical optimum requires countervailing institutional design: dedicated UBI financing vehicles, hypothecated taxes, and democratic coalitions built around universal benefits.

9.3 Welfare Beyond Income

The welfare framework of Sections 5 and 6 assumes utility depends on consumption. Employment generates welfare beyond income: meaning, social identity, structure, and connection. UBI addresses the income dimension but not these others. A policy response that reaches these dimensions may supplement UBI with alternative sources of structured activity, such as subsidized civic participation, public employment options, or explicit jobs guarantees.

This consideration modifies the optimal structure at the margin: a share of revenue that would otherwise fund higher UBI instead funds complementary programs. The balance between the two depends on empirical parameters the model does not carry, and no split is computed here.

9.4 Endogenous Technical Change

We have treated the evolution of α as governed primarily by factor prices, policy, and exogenous technological progress. The direction of AI development is itself responsive to policy. Taxation of labor-displacing AI can redirect research and development (R&D) toward labor-complementing

AI, reducing the magnitude of the displacement the tax is designed to address. The Acemoglu-style directed technical change effect implies that moderate, targeted AI taxation can attain higher values of the social welfare function in the long run than the static optimum alone suggests.

Whether domestic taxation can redirect global AI R&D depends on the geographic distribution of AI development. Since AI development is currently concentrated in a small number of countries and firms, the answer depends on whether the major AI-producing countries coordinate. This reinforces the argument of Section 8 for international coordination on AI capital taxation.

9.5 Demographic Heterogeneity

The model treats workers symmetrically conditional on skill θ. In practice, age matters: a 55-year-old worker's θ is less elastic to retraining than a 25-year-old's. An overlapping-generations (OLG) extension would tie the speed of adjustment in F_t to the age distribution: older cohorts carry displacement with them until they exit the workforce, so displacement would persist rather than recover. In the computed path retraining investment is zero and the skill distribution does not adjust, so the age structure of adjustment has nothing to act on; an OLG treatment is an extension of the framework.

9.6 Scope and Limitations of the Framework

Seven limitations bound the conclusions of the framework.

First, the quantitative conclusions of the calibrated path (Section 7.7) are conditional on parameter choices that are not identified from data; the sensitivity analysis in Section 7.8 documents how the path responds to variation in σ, ψ, ε_K, and ξ_F, but does not resolve the identification question. Nor do the qualitative results all survive variation of the quantities the model fixes outside its stated calibration. Of the 11 quantities varied, five change a qualitative outcome of the path when moved to half or to one and a half times their value (Appendix C). With the ownership exponent halved, the rent rate stands at 46 percent at year 30 and the UBI share falls below an earlier value in years 1 through 27. With the reservation-wage response or the non-AI capital rate at one and a half times their values, or with other government spending or the hours elasticity halved, the UBI share stands below its year-0 value in the first year or two, by less than one percentage point; with the non-AI capital rate raised, the rent rate also stands 0.5 percentage points below its limit at year 30. In every accepted variant of the robustness check of Section C.6 (acceptance as defined in Section C.3), non-participation rises in every year and reaches no plateau; with the adoption ceiling halved, to $\alpha_{\max} = 0.375$, it rises to 38 percent at year 30.

Second, the deterministic adoption law (40) abstracts from uncertainty about the pace of AI capability development. A version with stochastic ζ_t would generate option-value considerations for tax policy under uncertainty, whose direction the analysis under certainty does not determine.

Third, the framework is closed with respect to aggregate demand: output is determined by factor supplies and technology, with no role for aggregate demand shortfalls. In an environment of rapid displacement and slow institutional adjustment, the Keynesian channel in which displaced workers' reduced spending depresses aggregate output is relevant, and would interact with the Pigouvian component of the optimal AI capital tax in ways this framework does not capture.

Fourth, the welfare criterion is utilitarian with concave individual utility. A prioritarian or Rawlsian criterion would raise the weight on the worst-off and would generically push the optimal UBI

level above the utilitarian optimum; a Nash-bargaining or median-voter criterion would bring the political-economy considerations of Section 9.2 inside the welfare function rather than outside it. Whether the structure of the computed optimum, in which the quasi-rent to AI capital bears the AI tax in full, carries over to these alternative criteria is not examined here; the magnitudes would differ.

Fifth, AI capital is modeled as a single aggregate factor. In reality, AI capital is composed of heterogeneous assets (frontier models, inference infrastructure, and application-layer software) with differing substitution patterns against labor and differing international mobility. A disaggregated treatment is an extension that we do not attempt here.

Sixth, the transition values of Table 2 and the sensitivities of Figure 2 are the solution of the dynamic planner's problem by direct collocation (Appendix C). Every initialization reaches the same optimum, and the reported values move by at most 1.5 percentage points under a halved mesh and a doubled horizon. The solution nonetheless rests on the quantities the model fixes outside its stated calibration (Table C.1), and on an instrument set whose limits bind: the rent rate at its upper limit from year 24, and the normal-return rate and the labor rate at the participation margin at zero, their lower limits, in every year. Where one of these lower limits binds, the planner would choose a subsidy if the instrument set admitted one, and the reported optimum is conditional on its exclusion; with an earnings subsidy at the participation margin admitted, non-participation stays at or below 4.4 percent over the optima found (Section 7.9).

Seventh, the resource cost of AI capital (20) is charged as an annuity on the stock in place, not as investment when it is made. The two have the same present value along a path where the interest rate equals the discount rate, as it does here, and differ in their timing: the annuity charges an early expansion of the stock over the years that follow it. The cost amounts to 8.3 percent of output at year 0 and 9.5 percent at year 30; the optimum under a charge at the time of investment is not computed.

10 Conclusion

This paper has developed a formal framework for analyzing AI-induced unemployment and the optimal tax structure for financing universal basic income as a policy response. The framework combines a task-based microfoundation for the aggregate production function, a three-factor structure that permits differentiated tax treatment of AI capital, an upward-sloping AI capital supply arising from scarce complementary inputs, an endogenous skill distribution, and an expected-utility welfare criterion that captures the insurance value of UBI. The central conclusions are as follows.

First, mass unemployment from AI adoption is not inevitable but emerges under a specific constellation of six conditions: substitutability between AI and labor in the region identified by Proposition 1, rising AI intensity, an AI capital supply elasticity that is finite but not extreme, a skill distribution with nonnegligible density at the employment margin, skill adjustment that is slow relative to AI adoption, and weak reinstatement. Whether this constellation holds in the current technological moment is an empirical question that the theoretical framework does not settle.

Second, for an instrument set that excludes subsidies, the computed optimal financing structure places the AI capital tax on the quasi-rent to scarce complementary inputs and sets the tax on normal AI returns to zero: in the static optimum the rent rate is at its limit at every AI intensity considered, and along the path it rises from zero at year 0 to that limit in year 24. The labor income tax is progressive across terciles, from a zero rate at the participation margin to a top-tercile rate of 80

through 87 percent along the path. The consumption tax is not identified as a separate instrument in the model, so the case for a VAT rests on the enforcement and measurement considerations of Section 6.3. The structure finances UBI of 19 through 44 percent of per-capita GDP along the path.

Third, the optimal tax structure is time-varying and state-contingent. The computed solution shows the effective AI tax rising with AI intensity in every reported year, and below the static optimum at the same intensity until the rent rate reaches its limit, because the planner values faster adoption at the margin. The system is self-stabilizing: the same economic transition that generates the need for UBI also generates the tax base to finance it.

Fourth, the financing structure depends on international tax coordination. The two-country Nash analysis in Section 8 shows that the uncoordinated mean effective AI tax sits 23 percentage points below the cooperative mean, at a welfare cost equivalent to 2.1 percent of consumption. Destination-based taxation and minimum-rate coordination act on the margin on which that loss arises.

Several questions remain open. The welfare decomposition of Section 7.10 is stated and not computed. With an earnings subsidy at the participation margin admitted, the computed optima use it and non-participation stays at or below 4.4 percent (Section 7.9); a global optimum of that problem is not established, and a subsidy to the normal return on AI capital is not examined. The return to retraining is compared with evidence on displaced workers at the scale of existing programs, while the computed optimum, where it uses retraining, spends between 0.4 and 18.3 percent of output on it at its peak; a retraining technology with diminishing returns is not examined. Welfare beyond income is acknowledged but not formalized. The overlapping-generations extension is sketched but not developed. Endogenous technical change in the Acemoglu tradition, connecting AI capital taxation to the direction of R&D, is discussed but not integrated into the main analysis.

The framework provides a theoretical foundation for analyzing one of the policy questions that could arise over the coming decades if AI adoption follows a path similar to the one modeled here. Whether AI will in fact generate mass displacement is uncertain; conditional on it doing so, the framework identifies which policy instruments are most effective at different phases of the transition and which institutional constraints are the binding ones.

Bibliography

Acemoglu, Daron. "Directed Technical Change." *Review of Economic Studies* 69, no. 4 (2002): 781–809.

Acemoglu, Daron. "Harms of AI." In *The Oxford Handbook of AI Governance*, edited by Justin B. Bullock et al. Oxford: Oxford University Press, 2023.

Acemoglu, Daron, and David Autor. "Skills, Tasks and Technologies: Implications for Employment and Earnings." In *Handbook of Labor Economics*, vol. 4B, edited by David Card and Orley Ashenfelter, 1043–1171. Amsterdam: Elsevier, 2011.

Acemoglu, Daron, and Pascual Restrepo. "Automation and Rent Dissipation: Implications for Wages, Inequality, and Productivity." *Quarterly Journal of Economics* 141, no. 2 (2026): 1521–1579.

Acemoglu, Daron, and Pascual Restrepo. "The Race Between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment." *American Economic Review* 108, no. 6 (2018): 1488–1542.

Acemoglu, Daron, and Pascual Restrepo. "Robots and Jobs: Evidence from US Labor Markets." *Journal of Political Economy* 128, no. 6 (2020): 2188–2244.

Acemoglu, Daron, and Pascual Restrepo. "Tasks, Automation, and the Rise in U.S. Wage Inequality." *Econometrica* 90, no. 5 (2022): 1973–2016.

Atkinson, Anthony B., and Joseph E. Stiglitz. "The Design of Tax Structure: Direct versus Indirect Taxation." *Journal of Public Economics* 6, no. 1–2 (1976): 55–75.

Autor, David H. "Work of the Past, Work of the Future." *AEA Papers and Proceedings* 109 (2019): 1–32.

Autor, David H., Frank Levy, and Richard J. Murnane. "The Skill Content of Recent Technological Change: An Empirical Exploration." *Quarterly Journal of Economics* 118, no. 4 (2003): 1279–1333.

Banerjee, Abhijit, Paul Niehaus, and Tavneet Suri. "Universal Basic Income in the Developing World." *Annual Review of Economics* 11 (2019): 959–983.

Beraja, Martin, and Nathan Zorzi. "Inefficient Automation." *Review of Economic Studies* 92, no. 1 (2025): 69–96.

Chamley, Christophe. "Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives." *Econometrica* 54, no. 3 (1986): 607–622.

Chui, Michael, Eric Hazan, Roger Roberts, Alex Singla, Kate Smaje, Alex Sukharevsky, Lareina Yee, and Rodney Zemmel. *The Economic Potential of Generative AI: The Next Productivity Frontier*. McKinsey & Company, June 2023.

Costinot, Arnaud, and Iván Werning. "Robots, Trade, and Luddism: A Sufficient Statistic Approach to Optimal Technology Regulation." *Review of Economic Studies* 90, no. 5 (2023): 2261–2291.

Diamond, Peter A., and James A. Mirrlees. "Optimal Taxation and Public Production I: Production Efficiency." *American Economic Review* 61, no. 1 (1971): 8–27.

Diamond, Peter A., and Emmanuel Saez. "The Case for a Progressive Tax: From Basic Research to Policy Recommendations." *Journal of Economic Perspectives* 25, no. 4 (2011): 165–190.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. "GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models." Working paper, arXiv:2303.10130v5, August 2023. Published as "GPTs Are GPTs: Labor Market Impact Potential of LLMs." *Science* 384, no. 6702 (2024): 1306–1308.

Ghatak, Maitreesh, and François Maniquet. "Universal Basic Income: Some Theoretical Aspects." *Annual Review of Economics* 11 (2019): 895–928.

Guerreiro, Joao, Sergio Rebelo, and Pedro Teles. "Should Robots Be Taxed?" *Review of Economic Studies* 89, no. 1 (2022): 279–311.

Heckman, James J., Lance Lochner, and Christopher Taber. "Explaining Rising Wage Inequality: Explorations with a Dynamic General Equilibrium Model of Labor Earnings with Heterogeneous Agents." *Review of Economic Dynamics* 1, no. 1 (1998): 1–58.

Hoynes, Hilary, and Jesse Rothstein. "Universal Basic Income in the United States and Advanced Countries." *Annual Review of Economics* 11 (2019): 929–958.

Ide, Enrique, and Eduard Talamàs. "Artificial Intelligence in the Knowledge Economy." *Journal of Political Economy* 133, no. 12 (2025): 3762–3800.

Jacobson, Louis, Robert LaLonde, and Daniel G. Sullivan. "Estimating the Returns to Community College Schooling for Displaced Workers." IZA Discussion Paper 1017. Bonn: Institute for the Study of Labor, February 2004. Published in *Journal of Econometrics* 125, no. 1–2 (2005): 271–304.

Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan. "Policies to Reduce High-Tenured Displaced Workers' Earnings Losses Through Retraining." Hamilton Project Discussion Paper 2011-11. Washington, DC: Brookings Institution, November 2011.

Japan Institute for Labour Policy and Training. *Databook of International Labour Statistics 2025*. Tokyo: Japan Institute for Labour Policy and Training, 2025.

Jones, Charles I. "AI and Our Economic Future." *Journal of Economic Perspectives* 40, no. 3 (2026): 3–22.

Jones, Charles I., and Christopher Tonetti. "Nonrivalry and the Economics of Data." *American Economic Review* 110, no. 9 (2020): 2819–2858.

Jones, Charles I., and Christopher Tonetti. "Past Automation and Future A.I.: How Weak Links Tame the Growth Explosion." Unpublished manuscript, 2026. Cited in Charles I. Jones, "AI and Our Economic Future," *Journal of Economic Perspectives* 40, no. 3 (2026): 11.

Judd, Kenneth L. "Redistributive Taxation in a Simple Perfect Foresight Model." *Journal of Public Economics* 28, no. 1 (1985): 59–83.

Korinek, Anton, and Joseph E. Stiglitz. "Artificial Intelligence and Its Implications for Income Distribution and Unemployment." In *The Economics of Artificial Intelligence: An Agenda*, edited by Ajay Agrawal, Joshua Gans, and Avi Goldfarb, 349–390. Chicago: University of Chicago Press, 2019.

Manuelli, Rodolfo E., and Ananth Seshadri. "Human Capital and the Wealth of Nations." *American Economic Review* 104, no. 9 (2014): 2736–2762.

Mirrlees, James A. "An Exploration in the Theory of Optimum Income Taxation." *Review of Economic Studies* 38, no. 2 (1971): 175–208.

Piketty, Thomas, Emmanuel Saez, and Stefanie Stantcheva. "Optimal Taxation of Top Labor Incomes: A Tale of Three Elasticities." *American Economic Journal: Economic Policy* 6, no. 1 (2014): 230–271.

Ramsey, Frank P. "A Contribution to the Theory of Taxation." *Economic Journal* 37, no. 145 (1927): 47–61.

Saez, Emmanuel. "Using Elasticities to Derive Optimal Income Tax Rates." *Review of Economic Studies* 68, no. 1 (2001): 205–229.

Straub, Ludwig, and Iván Werning. "Positive Long-Run Capital Taxation: Chamley–Judd Revisited." *American Economic Review* 110, no. 1 (2020): 86–119.

Thuemmel, Uwe. "Optimal Taxation of Robots." *Journal of the European Economic Association* 21, no. 3 (2023): 1154–1190.

Vivalt, Eva, Elizabeth Rhodes, Alexander W. Bartik, David E. Broockman, Patrick Krause, and Sarah Miller. "The Employment Effects of a Guaranteed Income: Experimental Evidence from Two U.S. States." *Quarterly Journal of Economics* 141, no. 4 (2026): 2675–2734.

Appendix A Derivation of the Aggregate CES from Task-Level Primitives

This appendix establishes the task-to-aggregate mapping used in Section 3. Starting from the task-level production (1)–(2) and the cost-minimization assignment, total effective labor employed in task $j \in \mathcal{J}_L$ is $\ell(j) = y(j)/\phi_L(j)$; total AI capital employed in task $j \in \mathcal{J}_K$ is $k(j) = y(j)/\phi_K(j)$. Under the CES task aggregator, the demand for task j at unit cost $p(j)$ satisfies $y(j) \propto p(j)^{-\eta_T}$; with $p(j) = w/\phi_L(j)$ for labor tasks and $p(j) = r_A/\phi_K(j)$ for AI tasks, equilibrium task output is $y(j) \propto \phi_\cdot(j)^{\eta_T}$, and the corresponding factor input is $\ell(j) = y(j)/\phi_L(j) \propto \phi_L(j)^{\eta_T - 1}$ (analogously for $k(j)$). Aggregating task-level labor and AI inputs yields

$$L = \int_{\mathcal{J}_L} \frac{y(j)}{\phi_L(j)}\, dj, \qquad K_A = \int_{\mathcal{J}_K} \frac{y(j)}{\phi_K(j)}\, dj.$$

The aggregate-level CES of equation (3) is then obtained by substitution, with the distributional weights

$$\Gamma_L(\mathcal{J}_L) = \left(\int_{\mathcal{J}_L} \phi_L(j)^{\eta_T - 1}\, dj \right)^{1/\eta_T}, \quad \Gamma_K(\mathcal{J}_K) = \left(\int_{\mathcal{J}_K} \phi_K(j)^{\eta_T - 1}\, dj \right)^{1/\eta_T},$$

the weights on $L^{(\eta_T-1)/\eta_T}$ and $K_A^{(\eta_T-1)/\eta_T}$ at a given assignment: within the AI tasks the input is allocated in proportion to $\phi_K(j)^{\eta_T-1}$, so that $\int_{\mathcal{J}_K} \big(\phi_K(j)\, k(j)\big)^{(\eta_T-1)/\eta_T} dj = \Gamma_K\, K_A^{(\eta_T-1)/\eta_T}$, and likewise for labor. With equal task productivities the weights are n_K^{1/η_T} and $(1 - n_K)^{1/\eta_T}$, $n_K = |\mathcal{J}_K|$, which gives equation (5). The aggregate substitution elasticity is of the additive form

$$\sigma(\eta_T, \mathcal{J}_K) = \eta_T + \Theta_B(\mathcal{J}_K), \qquad \Theta_B \geq 0,$$

where Θ_B is the extensive-margin contribution: it is proportional to the value-added-weighted density of tasks at the assignment boundary—tasks for which the unit costs of the two factors are approximately equal, $w/\phi_L(j) \approx r_A/\phi_K(j)$—times the reassignment response of those tasks to a change in relative factor prices. The intensive margin alone (reallocation within a fixed assignment) yields $\sigma = \eta_T$; reassignment of boundary tasks adds substitutability, so $\sigma \geq \eta_T$, with equality when no tasks are near the margin. As $\mathcal{J}_K$ expands into regions where labor's comparative advantage is weak, the mass of tasks near the assignment boundary rises, and Θ_B—and with it the aggregate σ—rises. This is the mechanism by which the aggregate elasticity of substitution increases with AI's task reach, the property used in the Condition 1 discussion of Section 4.7.

Appendix B Notation Summary

The symbols are grouped by the part of the model in which they enter.

B.1 Factors and Production

K_A	AI capital stock
K_N	Non-AI (traditional) capital stock
L	Aggregate effective labor
Y	Final-good output

X	AI–labor composite (inner nest)
σ	Elasticity of substitution between K_A and L
ν	Elasticity of substitution between X and K_N
α	AI intensity in the inner nest
ξ_X	Outer-nest weight on X
A	Hicks-neutral total factor productivity
Z	Scarce complementary inputs to AI capital
Π_Z	Quasi-rent on Z
c_A	Resource cost of the AI capital stock, borne by its owners
$K_{A,0}$	AI capital stock inherited at the start of the transition
ε_K	AI capital supply elasticity
s_X	Outer-nest income share of the composite X
$\tilde{\sigma}$	Effective inner-nest elasticity, $1/\tilde{\sigma} = 1/\sigma - (1 - s_X)/\nu$
ϱ	Inner-nest exponent, $(\sigma - 1)/\sigma$

B.2 Tasks

$j \in [0,1]$	Task index
η_T	Task-level elasticity of substitution
$\mathcal{J}_K, \mathcal{J}_L$	Tasks performed by AI, by labor
n_K	Share of tasks performed by AI, $\lvert\mathcal{J}_K\rvert$
$\phi_K(j), \phi_L(j)$	Task-level productivities of AI, labor
Γ_K, Γ_L	Task-aggregation weights
Θ_B	Extensive-margin (boundary-task) contribution to σ

B.3 Labor, Skills, and Unemployment

$\theta \in [0,1]$	Worker skill type
$F_t(\theta), f_t(\theta)$	Cumulative and density of skill distribution
γ	Skill-in-production convexity
w, w_r	Wage, reservation wage
θ^*	Employment threshold
$u = F(\theta^*)$	Non-participation rate
κ_F	Skill-adjustment speed
I_F	Education / retraining investment
μ_t	Skill location: location of the skill distribution at t, μ_0 at year 0 (distinct from the distributional welfare weights μ_i of B.5)
ξ_F	Shift of the skill location per unit of I_F/Y
e_μ	Response of output to the skill location, prices given

B.4 Policy Instruments and Welfare

$\tau_L(\theta)$	Labor income tax (nonlinear)
$\tau_{K_A}^{norm}$	AI capital tax on normal returns
$\tau_{K_A}^{rent}$	AI capital tax on quasi-rents Π_Z
τ_{K_N}	Non-AI capital tax
τ_Π	Tax on pure economic rents
τ_C	VAT (uniform consumption tax)
b	Per-capita UBI transfer

G	Non-UBI government expenditure
$\omega(\theta)$	Share of AI capital owned by type θ
$\mathrm{MSC}(\alpha)$	Marginal social cost of AI deployment (consumption units)
$\mathrm{MSC}_K(\alpha)$	MSC per unit of AI capital
χ_H, χ_I	Human-capital, institutional externalities

B.5 Ramsey System

λ	Marginal cost of public funds
μ_i	Distributional welfare weight on base i
ε_i^{own}	Compensated own-elasticity of base i
$\varepsilon_{j,i}$	Cross-elasticity of base j with respect to p_i
Y_i	Cross-elasticity correction factor
$\mathcal{E}_i$	Money-metric marginal external benefit of instrument i

B.6 Dynamic Environment

ρ	Social discount rate
δ	AI capital depreciation rate
r	Risk-free interest rate
φ	Baseline adoption speed
η	Wage–rental elasticity of adoption
ψ	Tax-sensitivity of adoption
ζ_t	Exogenous technological progress
κ	Congestion-externality parameter
ι	Institutional absorption rate
$\lambda_\alpha, \lambda_{K_A}, \lambda_F$	Co-state variables

B.7 Risk and Insurance

$D(\alpha, \theta)$	Expected skill-devaluation rate
$s(\theta)$	Idiosyncratic volatility
$\mathcal{W}_t$	Standard Wiener process
γ_R	Coefficient of relative risk aversion (CRRA)
π	Proportional risk premium
CV	Coefficient of variation of consumption
ς^2	Log-variance of the consumption distribution

B.8 International Dimension

$\mathcal{A}, \mathcal{B}$	Country labels
$\Psi(\cdot)$	Capital-reallocation adjustment cost
m	Capital-mobility sensitivity
$\Delta(m)$	Nash level shift from the outflow margin
ΔW	Coordination welfare gain

Appendix C Numerical Solution and Calibration

This appendix states the numerical method, the welfare function and the instrument set, the checks of the solution's accuracy, the choice of the tax base for the quasi-rent, the normalization and the ceilings of the adoption law, the quantities fixed outside the stated calibration, the earnings-subsidy variant, and the two-country solution.

C.1 Method

The dynamic planner's problem of Section 7.3 is solved, in the form Section 7.7 states, by simultaneous direct collocation: the states (AI intensity α_t and the skill distribution, which the computed problem represents by its location μ_t), the controls, and the algebraic equilibrium variables (the wage and the AI capital stock on its supply schedule, equation (42)) are discretized on a Hermite–Simpson mesh of 0.5-year intervals over 120 years with a stationary terminal value and the skill location stationary at the final node. The laws of motion enter as defect constraints and the equilibrium conditions and the government budget as equality constraints at every node and midpoint; the instruments are linear between nodes, except the transfer, which is free at each midpoint so that the budget is met there by the transfer and not by the tax rates. The resulting nonlinear program is solved by an interior-point method with exact first and second derivatives obtained by automatic differentiation. Skill integrals use 32-node Gauss–Legendre quadrature on $[\theta^*, 1]$ after the change of variable $\theta = \theta^* + (1-\theta^*)s$, so every aggregate is smooth in θ^*. The static optimum of Section 6.6 is the same economy, instrument set, and welfare function at a given AI intensity without the laws of motion; at the adoption ceiling it is the steady state of the dynamic problem, and the terminal node of the path agrees with the static optimum at its own AI intensity to within 10^{-5}. The year-0 constants A, K_N, $K_{A,0}$, and w_r are identified from the targets $Y_0 = 1$, a labor share of 60 percent, a laissez-faire non-participation rate of 5 percent, and an AI capital income share of 15 percent, with residuals at machine precision.

C.2 Welfare and Instruments

The objective is the discounted utilitarian sum of constant-relative-risk-aversion utility of consumption, net of the disutility of hours and of the congestion term of Section 7.5. The terms χ_H and χ_I of equation (37) are set to zero, the model giving them no functional form, so the marginal social cost of displacement reduces to its fiscal term, which the planner's budget carries. The instruments are the normal-return and rent rates on AI capital, the three tercile labor rates, the transfer, and retraining investment. Tax rates are confined to the interval from zero through 100 percent, which excludes subsidies; the consumption tax is held at zero (Section 6.3); the non-AI capital rate is fixed. The tercile rates are joined by a smooth step whose width is fixed in units of the population's skill distribution, so the schedule smooths as employment contracts. With the width set in units of the employed distribution instead, at three values of its steepness, the static optimum's non-participation, effective AI tax, and UBI share move by at most 1.3 percentage points at every AI intensity of Table 1, and the bottom- and middle-tercile average rates by up to 20 percentage points; those two rows of Table 1 depend on the smoothing at high AI intensity, and the rate at the participation margin is zero under each.

C.3 Accuracy of the Solution

For every reported solution the solver reports a local optimum, with constraint residuals of order 10^{-8}. A solution is accepted where the solver reports a local optimum and its path stays off the bounds the solver imposes for numerical stability. The baseline problem is solved from 11 initializations. Of these, five are built without an earlier solution: the static optima interpolated along the adoption path (the path of AI intensity that the adoption law (40) gives with its bracket held at one), with a flat wage and the inherited stock as the starting values of the wage and the stock; the same initialization, solved with a weighted penalty on the change of the instruments between dates and then with the weight lowered to zero in steps; every instrument at zero, the laissez-faire policy; the tax rates and the transfer of the static optimum at the third phase

threshold, held at every date; and the same instruments moving linearly over the horizon from the year-0 static optimum to the one at that threshold. The solutions from the two best of these are each perturbed by raising and by lowering every tax rate by 5 percentage points, which gives four more, and the remaining two are the solutions at $\xi_F = 6.5$ and 8.0 of the sensitivity grid of Section 7.8, which use retraining. All reach one optimum, with objectives that agree to within 10^{-7}. Halving the mesh to 0.25 years and doubling the horizon to 240 years move the reported rates by at most 1.5 percentage points and non-participation by at most 0.1 percentage points over years 0 through 30. The closed-form direct-substitution term of Proposition 1 agrees with the numerical partial derivative at every α tested. The static optimum that Section 7.7 compares with the path is computed at the path's AI intensity in each tabulated year. No re-solved path of the sensitivity grids reaches a bound the solver imposes for numerical stability, and neither does the baseline solution at either horizon.

C.4 Rent Base

The quasi-rent is the stock times the excess of the supply price over the normal return, $\Pi_Z = K_A(s - r - \delta)$, equation (16). With the normal-return base $(r+\delta)K_A$ and its tax it partitions the AI factor payment, $r_A K_A = \tau_{K_A}^{norm}(r+\delta)K_A + (r+\delta)K_A + \Pi_Z$, at every year of the path. The area above the supply schedule, $K_A\, s/(1+\varepsilon_K)$, is a different object: at year 0 that area and the normal-return base together exceed the factor payment by 5.2 percent of output, so it is not used as a tax base.

C.5 Adoption Law

The adoption ceiling $\alpha_{\max} = 0.75$ is a scenario (Section 7.7). The lower ceilings 0.53, 0.59, and 0.62 are solved as variants, each initialized from the baseline solution: every qualitative outcome of the path holds at each, retraining is not used, and non-participation at year 30 is 49, 54, and 56 percent against 64 percent at the ceiling $\alpha_{\max} = 0.75$. With the ceiling at one, the planner's problem has no interior long-run regime: from year 50 employment falls below one percent of the population and the wage reaches the solver's upper bound, because the wage–rental term of the law rises without bound as employment vanishes and adoption has scope left. The solution is therefore not accepted (Section C.3) and is reported for this property alone. The progress term is the constant $\zeta = (r_{A,0}/w_0)^{\eta}$, which makes φ the year-0 adoption rate per unit of remaining scope; without that normalization the bracket of the law at the year-0 laissez-faire prices would be 2.96 in place of one, and AI intensity would reach 0.65 at year 30 in place of 0.49. The ceiling $\alpha_{\max} = 0.75$ and the progress term are rows of Table C.1.

C.6 Quantities Fixed Outside the Stated Calibration

Table C.1 lists the 15 quantities the model fixes outside its stated calibration, each with its value, its source or the absence of one, and whether the robustness check described here varied it. Each of the 11 varied quantities except the adoption ceiling was moved to half and to one and a half times its value; the ceiling was halved and raised to 0.90, since one and a half times its value exceeds one, the upper limit of AI intensity. Seven qualitative outcomes of the path over years 0 through 30 were recorded: whether non-participation, the effective AI tax, and the UBI share rise in every year; whether the labor share falls in every year; whether non-participation plateaus; whether the normal-return rate is zero; and whether the rent rate is at its limit by year 30; a series rises in every year where it never stands half a percentage point or more below an earlier value. Non-participation and the effective AI tax rise in every year, the labor share falls, non-participation reaches no plateau, and the normal-return rate is zero in every accepted variant. The UBI share does not rise in every year in five variants: the reservation-wage response and the non-AI capital rate at one and a half times their values, and the ownership exponent, other government spending, and the hours elasticity at half theirs. In two of them, the ownership exponent halved and the non-AI capital rate raised, the rent rate is also below its limit at year 30. Two variants are not accepted and are not counted: the adoption ceiling raised to 0.90, whose path reaches the solver's upper bound on the wage beyond the reported years, and the skill spread at half its value, whose continuation did not converge.

Table C.1: Quantities the Model Fixes Outside Its Stated Calibration. "Varied" marks quantities included in the robustness check of Section C.6 (each moved to half and to one and a half times its value, except the adoption ceiling, halved and raised to 0.90). Values are those the baseline solution uses.

Quantity	Value	Source	Varied
Skill-distribution spread (logit scale)	0.6	unsourced	yes
Retraining shift of the skill location per unit I_F/Y	3.5	evidence range	yes
Reservation-wage response to UBI	1	unsourced	yes
Ownership concentration exponent	3	unsourced	yes
Non-AI capital tax rate	0.225	unsourced	yes
Tax rate on pure rents outside AI (zero base in the model)	0.8	unsourced	no
Other government spending, share of year-0 output	0.15	unsourced	yes
Deceleration cost curvature	5	unsourced	yes
Institutional absorption rate of adoption, per year	0.02	unsourced	yes
Adoption ceiling (saturation factor $\alpha_{\max} - \alpha$)	0.75	scenario	yes
Progress term of the adoption law, constant at $(r_{A,0}/w_0)^\eta$	0.338	year-0 normalization	yes
Interest rate, equal to the discount rate ρ	0.03	stipulated	no
Mobility share, two-country extension	0.5	stipulated	no
AI capital income share at year 0 (calibration target)	0.15	unsourced	no
Intensive-margin hours elasticity	0.25	Saez, "Using Elasticities"	yes

C.7 Earnings-Subsidy Variant

The variant of Section 7.9 replaces the lower limit of zero on the rate set at the participation margin with a bound of −1,000 percent, imposed for numerical stability, that must not bind, and keeps every other limit. It is solved from four initializations. The first is the baseline solution, with the bound moved from zero to its value in steps. The second starts from the first's solution. In each year whose non-participation is at least half its year-0 value, it replaces the instruments, the wage, and the stock by values interpolated between the years in which non-participation is below half its year-0 value; a year outside the span of those years takes the values of the nearest of them. The third consists of the variant's static optima, interpolated along the adoption path, with a flat wage and the inherited stock as the starting values of the wage and the stock. The fourth consists of the variant's static optima recomputed on a grid of AI intensities and interpolated along the adoption path with each optimum's own wage and stock. The optimum at the year-0 intensity is solved from the baseline's static optimum there, and each optimum above it from the one below. The solution from an initialization is counted among the optima found where the solver reports a local optimum and the path stays off the bound of −1,000 percent; each such solution also stays off the other bounds the solver imposes for numerical stability, so it is accepted in the sense of Section C.3.

C.8 Two-Country Solution

Each country's AI capital stock on its supply schedule is reduced by the fraction $m(\tau - \tau')$ of itself, where τ and τ' are the home and foreign effective AI tax rates and $m = 0.5$ is the mobility share, so that relocation is zero in the symmetric equilibrium. Each planner solves the dynamic problem of Section 7.3 taking the other country's path of effective rates as given, and the symmetric Nash outcome is the fixed point of a damped best response on that path; the iteration stops when the path moves by less than half a percentage point at every reported year, which it does after seven iterations. The cooperative optimum is the closed-economy solution. The welfare loss is expressed as a consumption equivalent: the uniform proportional reduction of every household's consumption along the cooperative path that lowers the planner's objective to its Nash value.