

The AI Deflation Dividend

Ha Nguyen*

Rachel Yuting Fan*

May 2026

Abstract

Will artificial intelligence compress wages? We build an analytically tractable general equilibrium model with an oligopolistic AI sector selling labor-augmenting services to competitive production sectors. A nested labor bundle distinguishes augmentable from unaugmentable tasks – those requiring physical presence, human judgment, or legal accountability – and delivers a natural wage floor proportional to each sector’s unaugmentable share. Two channels act in opposition: a price-output channel that partially supports wages through the revenue effect, and a task-bundle channel through which cheap AI substitution compresses the marginal product of augmentable labor. The paper delivers two central results. First, the unaugmentable share drives a wedge between goods price deflation and nominal wage compression: because AI lowers output prices across all sectors while wages are anchored by the tasks AI cannot perform, consumer prices fall faster than paychecks shrink, raising real wages economy-wide. This price-deflation channel is the paper’s central mechanism. Second, the model delivers a clear policy ranking: removing adoption barriers is approximately Pareto improving in the calibrated economy, while capping the oligopoly markup merely redistributes without growing the surplus. Calibrated to actual AI usage data (Anthropic Economic Index, 2026), barrier removal raises real wages by 25.2 percent relative to the no-AI baseline, 17.3 percentage points above the current-adoption equilibrium. A markup cap adds only 1.6 percentage points to real wages above the current equilibrium.

*Ha Nguyen is a Senior Economist at the IMF. Rachel Yuting Fan was an Economist at the World Bank’s MENAAP Chief Economist’s Office at the time of writing. The authors thank Anton Korinek, Andrew Berg, and Felipe Zanna for helpful comments. The authors used AI for editing and coding assistance. All errors are our own. All views expressed in this paper are those of the authors and do not necessarily reflect the position of the IMF, the World Bank, their Executive Boards, or their management.

1 Introduction

Artificial intelligence could reduce wages and destroy jobs. The concern is real: as AI services become cheaper, firms substitute AI for the tasks it can handle, reducing the scarcity of human augmentable hours and compressing their marginal product. In AI-intensive sectors, workers effectively compete with an increasingly cheap substitute – one whose price falls through learning-by-doing while their own wages do not. At the same time, the AI sector is dominated by a small number of firms with high fixed costs, who set prices above marginal cost and extract rents that flow to their owners rather than to the workforce. These two forces together – task-bundle compression and oligopoly extraction – may generate a distributional outcome that a standard productivity narrative obscures: aggregate output rises, but workers may bear the wage compression while capitalists capture the surplus.

Yet in a general equilibrium framework, other forces are at work. When firms adopt AI, goods prices fall – and if prices fall faster than nominal wages, workers are better off in real terms even as their paychecks shrink. This paper shows that prices can fall faster, and identifies the structural reason why. A fraction of tasks in every sector requires physical presence, irreducible human judgment, or legal accountability – tasks beyond current AI’s reach. This unaugmentable residual acts as a natural wage floor: wages are less likely to be driven down without limit, because some human labor remains essential to production. The floor is the mechanism that widens the gap between price deflation and nominal wage compression. In sectors where most tasks require physical presence – Construction, Healthcare Support, Food Preparation – AI lowers output prices by raising productivity on a small augmentable slice, while the large unaugmentable share prevents wages from tracking prices downward. These sectors can contribute to economy-wide deflation without contributing to wage compression. Workers across the economy consume cheaper goods from low-exposure sectors while their own wages compress by much less. The result is a deflation dividend: this paper characterizes when it arises, which policies unlock it, and what prevents workers from capturing it today. This price-deflation mechanism is absent from the otherwise closely related frameworks of Korinek and Suh (2024) and Restrepo (2025), which work in one-good or numeraire-unit environments. In those frameworks wages rise when labor remains scarce – but can fall sharply once automation advances far enough, with no goods-price offset. Our mechanism is distinct: the price channel generates a real wage gain independently of labor scarcity, as long as the unaugmentable share remains positive.

We build an analytically tractable general equilibrium model in which competitive production sectors use AI services alongside labor and capital to produce consumption goods. AI is supplied by a small number of firms with very high fixed costs – training and infrastructure – who set prices above marginal cost and extract a markup from users. Two types of households face AI adoption differently: workers, who earn only wages, and capitalists, who own both physical capital and the AI firms. Under separable CRRA preferences, whether employment responds positively to real wage gains depends on the income-effect parameter γ ; the distinction between nominal and real wage changes matters directly for worker welfare regardless of the employment margin. The

model is deliberately static: each equilibrium is indexed by the prevailing AI cost level rather than solved dynamically, which keeps the mechanisms – the wage floor, the surplus split, the barrier policy ranking – transparent and analytically tractable. We calibrate the model to actual AI usage data from the Anthropic Economic Index (AEI) and use it to decompose wage, employment, and welfare outcomes across a set of counterfactual scenarios. A structural model is only as credible as its calibration strategy, and most existing work relies on potential-exposure indices – measures of what AI *could* do to occupations based on task descriptions and engineering assessments of AI capabilities. These indices illuminate the technological frontier but do not reveal what firms are actually doing. Unlike potential-exposure indices, the AEI records what AI *is* doing: actual Claude conversations classified to O*NET task categories, with empirical estimates of worker-hours saved per interaction. This makes two structural primitives directly observable – the AI productivity weight ϕ_s (worker-years saved per conversation) and the adoption intensity x_s/L_s (annual conversations per worker) – at the sector level. The adoption barriers τ_s that currently constrain firms below their frictionless optimum are then identified from the gap between observed adoption and the level that the AI price would justify. The AEI calibration therefore converts abstract structural parameters into quantities grounded in realized firm behavior, and the deflation dividend it recovers is not a theoretical possibility but a dividend already visible in the data.

Two channels determine how wages respond to AI adoption, and together they imply a structural limit on how far wages can fall (Propositions 1–2). The *price-output channel*: AI raises output and lowers goods prices, but nominal revenue rises under elastic goods demand, so this channel partially supports wages. The *task-bundle channel*: as firms buy more AI – attracted by its lower cost – they flood the augmentable bundle with a cheap substitute, reducing the scarcity of human augmentable hours; the marginal product of those hours falls precisely because AI, not the human hour, is now the marginal unit covering augmentable tasks. This channel compresses wages directly through the production function, independently of goods-market effects. Nominal wages fall when the task-bundle compression exceeds the price-output support.

The structural limit on wage compression – the natural wage floor – is set by the share of tasks in each sector that AI cannot do. Construction workers, nurses, and farmhands perform tasks that require physical presence; current AI software cannot readily substitute for them. Workers in these sectors are largely protected. Even in Computer & Math, where more than half of all tasks are AI-augmentable, wages cannot fall below 49 percent of their pre-AI level from the task-bundle channel alone. This wage floor has no analog in the standard Acemoglu–Restrepo framework, which allows wages to fall without limit as AI deepens. It is structurally related to the “thick tail” condition in Korinek and Suh (2024): their unbounded task-complexity distribution ensures labor always has harder tasks to perform, keeping wages from collapsing; our discrete unaugmentable share $(1 - f_s)$ plays the same anchoring role. The mechanisms differ – theirs operates through task scarcity in a one-good model, ours through the price-deflation channel in a multi-sector general equilibrium – but both identify the persistence of human-essential work as the critical structural condition.

Two policy instruments address the oligopoly distortion. *Barrier-removal* policies – worker

training, procurement reform, API standardization – reduce the organizational and informational frictions that slow AI adoption below its frictionless optimum. A *markup cap* forces AI firms to price at or near marginal cost, eliminating gross markup rents and transferring the savings to users through lower AI prices. The model delivers a decisive ranking of these instruments that hinges on whether physical infrastructure costs will fall far enough to allow competitive entry. Barrier removal is approximately Pareto improving in the calibrated economy (Proposition 4): AI suppliers receive higher gross markup payments, capital owners gain from higher output, and workers gain through higher real wages – because it expands the surplus pie rather than redistributing it. A price cap, by contrast, merely transfers rents from AI firms to users without growing the pie, and when entry is feasible it is ultimately redundant – as compute costs fall through learning-by-doing, new firms erode the incumbent markup and workers capture the competitive dividend without regulatory intervention. When data-center costs – land, power, and physical infrastructure – form a binding and permanent floor, however, the calculus reverses: the duopoly retains its markup indefinitely, and the markup cap becomes the only remaining lever for the gap that market forces cannot close.

The model is calibrated to the Anthropic Economic Index (AEI), which measures actual AI deployment by classifying real Claude conversations to O*NET occupational categories and estimating worker-hours saved per interaction – the first data source to reveal what AI is doing rather than what it could do. With a duopoly AI supplier, the model predicts that under current adoption with existing barriers, AI deflates the consumer price index by 9.4 percent while compressing nominal wages by 2.3 percent, leaving real wages 7.9 percent higher. Eliminating adoption barriers expands real wages by 25.2 percent; under the model’s separable CRRA preferences ($\gamma = 2$), employment falls 5.5 percent as the income effect dominates the substitution effect. This real wage gain far exceeds what capability improvements alone could achieve: expanding the AI-augmentable task frontier by half raises real wages by only 4.3 percentage points, because capability gains intensify both the price-reduction and wage-compression channels in roughly equal measure and largely cancel. What matters for workers is not how capable AI is, but whether firms can actually deploy it at the scale its cost justifies.

AI adoption need not cause net job loss at the macro level. Real wages rise because AI deflates consumer prices faster than it compresses nominal wages. Under the model’s separable CRRA preferences ($\gamma = 2$), the Marshallian elasticity $\varepsilon_M = -0.25$ means the income effect dominates: real wage gains induce workers to reduce hours, so employment falls across all scenarios. This paper treats the employment response as a secondary, specification-dependent result. The real wage result itself – the AI deflation dividend – is robust: sustaining the CPI-wages gap requires only that not all tasks are augmentable by AI, i.e., that $f_s < 1$ in every sector. Such tasks exist for structural reasons unlikely to disappear – physical presence cannot be delegated to software, and legal accountability structures require a human to bear responsibility for high-stakes decisions. The paper’s predictions therefore do not require strong assumptions about the permanent limits of AI capability. The condition $f_s < 1$ is also the analog of Korinek and Suh’s (2024) bounded-complexity scenario: if f_s were to converge to one – all tasks eventually becoming augmentable –

Proposition 3(iii) shows the wage floor dissolves and the deflation dividend narrows. The durability of the dividend thus depends on the pace at which the unaugmentable frontier erodes, which is an empirical question this paper cannot resolve but which our physical-capacity analysis (Appendix E) shows is far from binding at current adoption levels.

The paper is organized as follows. Section 2 reviews the literature. Section 3 presents the model and equilibrium definition. Section 4 describes the AEI data. Section 5 calibrates the model. Section 6 establishes the two headline wage propositions (Propositions 1–2) and two supporting lemmas. Section 7 presents Proposition 3 (Equilibrium Wage Path), Corollary 1 (Persistent Infrastructure Cost), and decomposes wage changes across six counterfactual scenarios. Section 8 analyzes price regulation and Proposition 4 (Barrier Policy), showing that barrier reduction approximately Pareto-dominates price caps in the calibrated economy. Section 9 discusses the key modeling choices. Section 10 concludes.

2 Related Literature

Task-based models of automation and augmentation. The task-based framework of Acemoglu and Restrepo (2018) and Acemoglu and Restrepo (2019) is the closest structural antecedent. They show that automation displaces tasks from labor, lowering the labor share unless new tasks are created to reinstate it. We adopt their task-partition structure: in sector s , only fraction f_s of tasks are augmentable by AI, and the remaining $1 - f_s$ are performed by human labor alone. Our key departure is the aggregation technology: where A-R place all tasks in a single CES bundle (appropriate for robots that can substitute for physical tasks), we use a nested Cobb-Douglas/CES that makes non-augmentable labor essential. This delivers a natural wage floor absent from the A-R specification, and is the appropriate structure when AI is modeled as software augmenting cognitive/data-processing tasks rather than as a robot substituting for physical effort. Zeira (1998) is the earlier model of labor-saving technology with machines; we generalize his insight to a nested-bundle framework with an oligopolistic technology supplier. Eloundou et al. (2023) map the task-based exposure logic to large language models, showing that a majority of US workers are in occupations where AI could affect a substantial fraction of tasks. Their measure captures potential exposure; our AEI-based calibration captures actual adoption, using realized conversations to pin down the structural primitives f_s and ϕ_s .

Recent frameworks for the transition to AGI. Korinek and Suh (2024) develop a compute-centric model in which tasks differ in computational complexity and an automation index grows with Moore’s Law. Their central result is that wages follow a hump-shaped trajectory during the transition: automation initially raises wages as productivity gains dominate, but once the capital-to-labor ratio crosses a threshold, labor becomes abundant and wages collapse to the level of capital returns. In scenarios where the task distribution is bounded, this collapse occurs in finite time – within 15–20 years under their baseline calibration – and the labor share converges to zero. Our unaugmentable share $(1 - f_s)$ is the discrete analog of their thick-tail condition: it prevents the

automation frontier from eliminating all human-essential work and pins a positive wage floor at every adoption level. The critical difference is the welfare mechanism: in Korinek and Suh (2024), the single output good is normalized as numeraire, so nominal and real wages are definitionally identical; wages rise only when labor remains scarce and fall sharply once it does not. In our model, goods prices fall faster than nominal wages as long as $f_s < 1$, generating a real wage gain independently of labor scarcity. The policy contrast is equally sharp: Korinek and Suh (2024) show that slowing automation maximizes wages at large output cost (Section 4.3). Our model reverses this conclusion because more adoption lowers consumer prices; the output cost of restrictions is simultaneously a real-wage cost their framework cannot capture.

Restrepo (2025) distinguishes bottleneck work – essential to output expansion – from supplementary work that AI can replicate, and shows that wages anchor to the compute-equivalent cost of replacing human skills. The transition dynamics depend on automation speed relative to capital accumulation: if automation outpaces capital, wages collapse immediately to the compute floor; if capital keeps pace, wages grow during the transition. In the AGI limit real wages may nominally exceed pre-AGI levels, but the labor share converges to zero. Our results are consistent in that long-run direction – if $f_s \rightarrow 1$, the wage floor disappears – but we characterize the empirically relevant case where f_s is well below one and adoption barriers are the binding constraint on real wage gains. The sharpest structural contrast is that neither Restrepo (2025) nor Korinek and Suh (2024) models barriers or a markup-bearing AI sector, so the Pareto-improving barrier-removal result and the barrier-versus-cap ranking are specific to our framework.

Technology and market structure. The fixed-cost structure of the AI sector connects to the literature on endogenous market structure with increasing returns. With free entry and sunk costs, the number of firms is determined by a zero-profit condition, as in Dixit and Stiglitz (1977) and the entry-deterrence literature. Our Lemma 2 formalizes how a falling fixed cost maps into rising entry, lower markups, and declining rents, with a clean uniqueness proof. The Cournot-with-free-entry framework delivers the Lerner index $\mathcal{L} = 1/(N\eta^{\text{agg}})$ as a measurable summary of market power, where η^{agg} is the aggregate demand elasticity derived from the nested adoption condition.

AI, rents, and inequality. Korinek and Stiglitz (2021) and Bessen (2022) argue that concentration in the AI industry amplifies income inequality, but these are policy arguments rather than formal models. We provide the general equilibrium foundations: a two-class model in which the rent channel and the wage-compression channel together drive the distributional outcome, and a theorem (Lemma 1) that gives the exact splitting rule for the augmentation surplus. The formal policy ranking – barrier removal Pareto-dominates a markup cap, unless infrastructure costs form a permanent floor – does not appear in this literature and is the paper’s central policy contribution.

Empirical measurement of AI diffusion. Micro-level evidence that AI augments labor productivity motivates the calibration of ϕ_s . Noy and Zhang (2023) show in a randomized experiment that access to ChatGPT raises output quality and cuts task completion time for professional writing tasks, with larger gains for lower-skill workers. Brynjolfsson et al. (2023) find that generative AI raises the productivity of customer support agents by 14 percent on average, again with larger

gains for novice workers – consistent with the model’s prediction that the augmentable-task bundle benefits workers with lower baseline productivity. These papers establish that $\phi_s > 0$ is empirically well-grounded and that the productivity gains vary systematically across workers and tasks, as the nested bundle structure assumes. Fan (2026) and Fan and Nguyen (2026) construct the Anthropropic Economic Index, which provides the sector-level AI adoption data that we calibrate. Our contribution is to use the AEI’s task-level time-savings estimates as a structural input: the hours-saved-per-conversation variable maps to the technology primitive ϕ_s in the inner CES bundle, and the cross-sectional variation in conversations per worker identifies θ from the model’s first-order condition.

3 Model

3.1 Environment

Time is static. The economy consists of S sectors indexed by $s = 1, \dots, S$, two classes of households (workers and capitalists), competitive firms in each sector, and an oligopolistic AI sector with N symmetric firms. Sectors correspond to the 22 SOC-2 occupation groups in the Anthropropic Economic Index.

3.2 Households

There are two types of households, reflecting the empirical reality that ownership of AI firms is highly concentrated and largely separate from the workforce.

Workers (mass λ). Workers are hand-to-mouth: they hold no capital and have no claim on firm or AI profits. A worker in sector s chooses consumption C_s and hours h_s to maximize separable CRRA utility:¹

$$U_s^W = \frac{C_s^{1-\gamma}}{1-\gamma} - \chi \frac{h_s^{1+\eta}}{1+\eta}, \quad \chi > 0, \eta > 0, \gamma > 0 \quad (1)$$

subject to the budget constraint $P C_s = w_s h_s$, where P is the aggregate CES price index and w_s is the sector- s wage. Here $C_s \equiv (\sum_{j=1}^S c_{j,s}^W)^{\sigma(\sigma-1)/\sigma}$ is the CES consumption aggregate over all S goods for a worker employed in sector s , with $c_{j,s}^W$ denoting consumption of good j ; the subscript s on C_s indicates the worker’s employment sector, not the good consumed (contrasting with the capitalist’s c_s^C in equation (5)). For $\gamma = 1$, interpret $C_s^{1-\gamma}/(1-\gamma)$ as $\ln C_s$.

Substituting the budget constraint $C_s = (w_s/P)h_s$, the first-order condition gives

$$h_s = \chi^{-1/(\eta+\gamma)} \left(\frac{w_s}{P} \right)^{\varepsilon_M}, \quad \varepsilon_M \equiv \frac{1-\gamma}{\eta+\gamma} \quad (2)$$

¹Separable CRRA is the standard workhorse in quantitative macro – Bewley–Aiyagari, RBC, and DSGE frameworks all use it as their baseline. Unlike quasi-linear GHH preferences, separable CRRA generates income effects on labor supply through the curvature parameter γ , making the employment response to real wage gains ambiguous. The paper therefore treats employment as a secondary, specification-dependent result. The purchasing-power gain is robust: worker welfare is strictly increasing in w_s/P for all $\gamma > 0$ by the envelope theorem, so the central result does not depend on the sign of the employment response.

where ε_M is the Marshallian labor supply elasticity. Its sign depends on γ : positive when $\gamma < 1$ (substitution effect dominates), zero when $\gamma = 1$ (log utility; income and substitution effects cancel), and negative when $\gamma > 1$ (income effect dominates). Since macro calibrations commonly use $\gamma \in [1, 4]$, the employment response to real wage gains is ambiguous in general. The central result of the paper – the purchasing-power gain – does not depend on this sign. By the envelope theorem, $dV_s^W/d(w_s/P) = (C_s^*)^{-\gamma} h_s^* > 0$ for all $\gamma > 0$: worker welfare is strictly increasing in the real wage regardless of the CRRA parameter. The indirect utility function is:

$$V_s^W = \frac{\eta + \gamma}{(1 - \gamma)(1 + \eta)} \chi^{-(1-\gamma)/(\eta+\gamma)} \left(\frac{w_s}{P}\right)^\mu, \quad \mu \equiv \frac{(1 - \gamma)(1 + \eta)}{\eta + \gamma} \quad (3)$$

for $\gamma \neq 1$. Because $V_s^W < 0$ when $\gamma > 1$, neither $\Delta \ln V_s^W$ nor the ratio $V_s^{W,1}/V_s^{W,0}$ is a valid welfare comparison in general. The clean object for scenario comparisons is the positive welfare index $\mathcal{W}_s \equiv [(1 - \gamma)V_s^W]^{1/(1-\gamma)}$, which satisfies:

$$\Delta \ln \mathcal{W}_s = \frac{1 + \eta}{\eta + \gamma} \Delta \ln \left(\frac{w_s}{P}\right) \quad (4)$$

The coefficient $(1 + \eta)/(\eta + \gamma) > 0$ for all $\eta, \gamma > 0$, so worker welfare rises whenever w_s/P rises – regardless of the CRRA parameter. For $\gamma = 1$, $\Delta V_s^W = \Delta \ln(w_s/P)$ directly.

Capitalists (mass $1 - \lambda$). Each capitalist is endowed with capital $k = K/(1 - \lambda)$, supplies no labor, and holds equal shares in all N AI firms. Capitalists receive capital income and AI oligopoly profits:

$$U^C = \left(\sum_{s=1}^S c_s^C \frac{\sigma-1}{\sigma} \right)^{\frac{\sigma}{\sigma-1}} \quad (5)$$

subject to $\sum_s p_s c_s^C = rk + \Pi_{AI}/(1 - \lambda)$, where r is the return to capital and $\Pi_{AI} = N\pi_i$ is total AI profit, with N the number of AI firms.

This two-class structure is the key departure from a representative-agent model. Workers and capitalists both consume all S goods, but their incomes come from entirely different sources. Wage compression falls entirely on workers, while capitalists benefit from higher capital returns and AI rents.

3.3 Firms and the Nested Labor Bundle

Each sector s is served by a competitive representative firm. Output is produced using an effective labor aggregate and capital:

$$Y_s = Z_s \ell_s^\alpha K_s^{1-\alpha}, \quad \alpha \in (0, 1) \quad (6)$$

The nested labor bundle. Labor in sector s has two components that play fundamentally different roles:

- **Unaugmentable labor** $(1 - f_s)L_s$: tasks requiring human physical presence, direct human interaction, or human accountability – situations where the consequences of error require a human to be responsible – for which there is no AI substitute. Examples include construction, patient care, field work, direct personal service, and high-stakes decisions carrying legal or professional liability. Without this component, sector output is zero.
- **Augmentable labor** $f_s L_s$: cognitive, analytical, and data-processing tasks that AI can perform alongside human labor. Examples include drafting, scheduling, reporting, research, and code generation.

The distinction is important for AI, though not for robots: a robot might plausibly replace physical construction tasks, but AI software cannot. A single economy-wide wage w_s is paid per unit of L_s , since workers perform both types of tasks within a sector and the split f_s is a technology parameter, not a worker characteristic.

The effective labor aggregate ℓ_s is defined by a two-level nest.

Inner bundle (AI-augmented component). Within the augmentable fraction, AI services x_s and augmentable labor $f_s L_s$ are CES substitutes:

$$\ell_s^{\text{aug}} = [(\phi_s x_s)^\rho + (f_s L_s)^\rho]^{1/\rho}, \quad \rho < 1, \rho \neq 0 \quad (7)$$

with elasticity of substitution $\theta = 1/(1 - \rho)$. Here $\phi_s \geq 0$ is the **AI productivity weight** – worker-years of effective labor per conversation in sector s , calibrated from AEI hours-saved data divided by 2080. With x_s in annual conversations and L_s in workers, $\phi_s x_s$ is in worker-years, directly commensurate with $f_s L_s$.

Outer bundle (aggregating essential and augmented labor). The overall effective labor aggregate combines unaugmentable labor and the augmented inner bundle via a Cobb-Douglas:

$$\ell_s = [(1 - f_s)L_s]^{1-f_s} \cdot [\ell_s^{\text{aug}}]^{f_s} \quad (8)$$

The exponents $(1 - f_s, f_s)$ match the task-composition shares and have a direct interpretation: a fraction f_s of the labor aggregate’s productivity comes from the AI-augmented bundle, and $1 - f_s$ from unaugmentable tasks.

Essential role of unaugmentable labor. From equation (8), $\ell_s = 0$ whenever $(1 - f_s)L_s = 0$ (provided $f_s < 1$). Hence sector output $Y_s = 0$ without unaugmentable labor, regardless of how much AI is deployed. This is the key property that distinguishes the nested structure from the A-R flat CES, in which x_s could in principle substitute for all of L_s .

Limiting cases.

- $f_s = 0$ (fully non-augmentable sector): $\ell_s \rightarrow L_s$, $Y_s = Z_s L_s^\alpha K_s^{1-\alpha}$ – standard Cobb-Douglas, no AI.

- $f_s \rightarrow 1$ (nearly fully augmentable): $\ell_s \rightarrow \ell_s^{\text{aug}}$, the inner CES alone.
- $x_s = 0$ (no AI adoption): $\ell_s^{\text{aug}} = f_s L_s$, so $\ell_s = [(1 - f_s)L_s]^{1-f_s} [f_s L_s]^{f_s} = f_s^{f_s} (1 - f_s)^{1-f_s} L_s \equiv \delta_s L_s$, where $\delta_s = f_s^{f_s} (1 - f_s)^{1-f_s} \in (0, 1)$ for $f_s \in (0, 1)$.

The parameter $f_s \in (0, 1)$ is sector s 's **task susceptibility**: the fraction of O*NET tasks in sector s with any AEI conversations, capturing the extensive margin of AI adoption. In the data, Computer & Math has $f_s \approx 0.51$ and Farming has $f_s \approx 0.02$.

Marginal products. Differentiating the nested bundle with respect to L_s and x_s :

$$\frac{\partial \ell_s}{\partial L_s} = \frac{\ell_s}{L_s} [(1 - f_s) + f_s \kappa_s^\rho] \quad (9)$$

$$\frac{\partial \ell_s}{\partial x_s} = f_s \cdot \frac{\ell_s}{\ell_s^{\text{aug}}} \cdot \left(\frac{\phi_s x_s}{\ell_s^{\text{aug}}} \right)^{\rho-1} \cdot \phi_s \quad (10)$$

where

$$\kappa_s \equiv \frac{f_s L_s}{\ell_s^{\text{aug}}} \in (0, 1] \quad (11)$$

is the **augmentable-labor intensity** in the inner CES bundle – the share of human augmentable labor in the total augmentable bundle: $\kappa_s = 1$ when $x_s = 0$ (no AI, so humans supply all augmentable efficiency units) and $\kappa_s < 1$ whenever $x_s > 0$ (AI contributes additional efficiency units, diluting the human share).

See Appendix C for the derivation of equations (9)–(10).

Euler's theorem. Since ℓ_s is homogeneous of degree 1 in (x_s, L_s) :

$$L_s \frac{\partial \ell_s}{\partial L_s} + x_s \frac{\partial \ell_s}{\partial x_s} = \ell_s \quad (12)$$

AI adoption condition. The firm maximizes:

$$\max_{L_s, K_s, x_s \geq 0} p_s Y_s - w L_s - r K_s - p_{AI} x_s \quad (13)$$

The first-order conditions are standard for capital ($r = p_s(1 - \alpha)Y_s/K_s$) and:

$$w = p_s \cdot \alpha \frac{Y_s}{\ell_s} \cdot \frac{\partial \ell_s}{\partial L_s} \quad (14)$$

$$p_{AI} = p_s \cdot \alpha \frac{Y_s}{\ell_s} \cdot \frac{\partial \ell_s}{\partial x_s} \quad (15)$$

Dividing (15) by (14) and letting $t \equiv \phi_s(x_s/L_s)$, the productivity-weighted AI-to-labor ratio (effective AI labor units per worker):

$$\frac{p_{AI}}{w} = \frac{f_s \phi_s t^{\rho-1}}{(1 - f_s) t^\rho + f_s^\rho} \quad (16)$$

This is an implicit equation for the AI-to-labor ratio $r_s \equiv x_s/L_s = t/\phi_s$.

Remark: no closed-form AI-to-labor ratio. Unlike the A-R CES specification, equation (16) cannot be solved explicitly for r_s for general ρ . The $(1 - f_s)t^\rho$ term in the denominator – which arises precisely from the unaugmentable labor channel in the MPL – prevents separation. In the degenerate limit where unaugmentable labor carries zero weight ($f_s = 1$), the equation reduces to $p_{AI}/w = \phi_s t^{\rho-1}$, which does solve for t in closed form. For interior $f_s \in (0, 1)$, $r_s(p_{AI}/w, \phi_s, f_s)$ is defined implicitly and is computed numerically in calibration.

Two limiting cases of adoption.

- *Low adoption* ($t \ll f_s$): denominator $\approx f_s^\rho$, so equation (16) gives $p_{AI}/w \approx f_s^{1-\rho} \phi_s t^{\rho-1}$, and solving for t : $t \approx [f_s^{1-\rho} \phi_s w / p_{AI}]^{1/(1-\rho)} = f_s \phi_s^\theta (w/p_{AI})^\theta$, so $r_s \approx f_s \phi_s^{\theta-1} (w/p_{AI})^\theta$. This coincides with the factor-augmenting model.
- *High adoption* ($t \gg f_s$): denominator $\approx (1 - f_s)t^\rho$ and $p_{AI}/w \approx f_s \phi_s / [(1 - f_s)t]$, giving $r_s \approx f_s \phi_s w / [(1 - f_s) p_{AI}]$. AI productivity ϕ_s enters only through this expression; the adoption ratio saturates and θ plays no role in the limit.

The wage function. Combining (14) and (9):

$$w_s = \frac{\alpha p_s Y_s}{L_s} \cdot \Gamma(r_s) \quad (17)$$

where

$$\Gamma(r_s) \equiv \frac{(1 - f_s)(\phi_s r_s)^\rho + f_s^\rho}{(\phi_s r_s)^\rho + f_s^\rho} \in [1 - f_s, 1] \quad (18)$$

and $r_s \equiv x_s/L_s$ is the AI-to-labor ratio (units of AI service per worker). $\Gamma(r_s)$ is the task-bundle compression factor: it scales the wage relative to the sector's nominal output per worker, capturing how much AI adoption has eroded the marginal product of human augmentable labor. When $r_s = 0$, AI plays no role and $\Gamma = 1$ – workers receive the full Cobb-Douglas share. As r_s rises, AI crowds out the augmentable labor component and Γ falls, compressing the wage directly through the production function, independently of any goods-price effect. Two properties follow immediately:

- (i) $\Gamma(0) = 1$: with no AI, $w_s = \alpha p_s Y_s / L_s$, the standard Cobb-Douglas MPL.
- (ii) $\Gamma(r_s) \rightarrow (1 - f_s)$ as $r_s \rightarrow \infty$: even with unlimited AI, $w_s \rightarrow (1 - f_s) \cdot \alpha p_s Y_s / L_s > 0$.

Property (ii) is the natural wage floor established below.

3.4 The AI Oligopoly

The AI sector consists of N symmetric firms competing in quantities (Cournot). Each firm i bears a fixed cost $F > 0$ and a marginal cost $mc > 0$ (inference). The fixed cost has two components: a compute cost F_c (model training, R&D), which declines over time through learning-by-doing, and

a physical infrastructure cost F_d (data-center land, power, cooling), which is tied to real assets and does not fall with the chip learning curve. In the static model $F = F_c + F_d$ is treated as given; Section 7 embeds it in a dynamic path $F_t = F_0 \delta^t$ and examines what happens when $F_d > 0$ prevents F_t from reaching zero.

Demand. Total AI demand is:

$$X(p_{AI}) = \sum_{s=1}^S L_s \cdot r_s(p_{AI}) \quad (19)$$

where $r_s(p_{AI})$ is defined implicitly by (16).

Sector-level and aggregate AI demand elasticities. In the flat CES of Acemoglu–Restrepo, the adoption condition solves explicitly for $r_s \propto (w_s/p_{AI})^\theta$, yielding a constant elasticity $\eta_s = \theta$ regardless of adoption intensity. In the nested model, the adoption condition (16) is implicit, so the elasticity varies with the adoption level. Differentiating (16) with respect to $\ln(w_s/p_{AI})$ at fixed $\{w_s, \phi_s, f_s\}$ and inverting: The **sector-level AI demand elasticity** η_s measures how responsive the AI-to-labor ratio r_s is to a change in the relative price of AI versus labor:

$$\eta_s(t_s) \equiv \frac{d \ln r_s}{d \ln(w_s/p_{AI})} = \frac{(1 - f_s) t_s^\rho + f_s^\rho}{(1 - f_s) t_s^\rho + (1 - \rho) f_s^\rho}, \quad t_s \equiv \phi_s r_s \quad (20)$$

Three properties follow directly.

- (i) *Flat-CES limit.* As $t_s \rightarrow 0$ (low adoption), $\eta_s \rightarrow 1/(1 - \rho) = \theta$. When AI adoption is negligible, the unaugmentable labor in the denominator of (16) plays no role and the nested model behaves like the flat CES: the sector’s AI demand elasticity equals the substitution parameter θ between AI and augmentable labor.
- (ii) *Elasticity falls with adoption.* $\partial \eta_s / \partial t_s < 0$ for $\rho \in (0, 1)$: as AI deepens in sector s , demand for AI services becomes less price-elastic. The reason is that the unaugmentable component $(1 - f_s)$ becomes an increasingly binding constraint as r_s rises – firms cannot substitute further into AI even if the price falls, because the production function requires human labor for the remaining tasks. In the limit $t_s \rightarrow \infty$, $\eta_s \rightarrow 1$: AI demand approaches unit elasticity, the minimum consistent with profit maximization.
- (iii) *Strict lower bound.* $\eta_s(t_s) > 1$ for all finite t_s , since the numerator exceeds the denominator by $\rho f_s^\rho > 0$. This guarantees that AI demand is always sufficiently elastic for the Cournot equilibrium to be well-defined and the markup to remain finite at any adoption level.

The aggregate demand elasticity entering the Cournot FOC is the AI-quantity-weighted average:

$$\eta^{\text{agg}}(p_{AI}) \equiv -\frac{d \ln X}{d \ln p_{AI}} = \frac{\sum_s L_s \eta_s(t_s) r_s(p_{AI})}{\sum_s L_s r_s(p_{AI})} \quad (21)$$

Since each $\eta_s > 1$, we have $\eta^{\text{agg}} > 1$ at any finite adoption level. Since η_s is decreasing in t_s and t_s decreases with p_{AI} , η^{agg} is strictly increasing in p_{AI} : lower AI prices deepen adoption and push each sector's elasticity below θ , amplifying the markup.

Cournot equilibrium. In the symmetric equilibrium the Lerner index satisfies:

$$\mathcal{L} \equiv \frac{p_{AI}^* - mc}{p_{AI}^*} = \frac{1}{N \eta^{\text{agg}}(p_{AI}^*)} \quad (22)$$

At low adoption $\eta^{\text{agg}} \approx \theta$, recovering the flat-CES approximation $\mathcal{L} = 1/(N\theta)$. At higher adoption $\eta^{\text{agg}} < \theta$, so the true Lerner *exceeds* $1/(N\theta)$: the flat-CES formula underestimates the markup as AI deepens. At the calibrated AEI R5 adoption level, $\eta^{\text{agg}} \approx \theta$ and the bias is negligible; it grows materially in the counterfactual scenarios of Section 7 as adoption rises.

Free entry drives per-firm profit to zero:

$$\pi_i = (p_{AI}^* - mc) \frac{X}{N} - F = 0 \quad (23)$$

which pins $N^*(F, \{w_s\}, \{\phi_s, f_s, L_s\})$.

The costs mc and F are denominated in units of the numeraire consumption good and represent real resource draws – predominantly compute capital services – supplied by capitalists who own the underlying infrastructure. Under this convention the global resource constraint clears exactly: the sum of household budget constraints equals total final goods output net of AI sector input requirements, satisfying Walras' law at every equilibrium.

3.5 Equilibrium

Definition 1 (Equilibrium). *An equilibrium is a price vector $\{p_s^*\}$, factor prices (w^*, r^*) , AI price p_{AI}^* , and quantities $\{Y_s^*, L_s^*, K_s^*, x_s^*\}$ and consumption allocations $\{c_s^{W*}, c_s^{C*}\}$ such that:*

- (i) *Each firm satisfies FOCs (14)–(15);*
- (ii) *Each AI firm satisfies (22) and (23); N^* is determined by free entry;*
- (iii) *Workers and capitalists maximize utility subject to their budget constraints; sector- s employment is $L_s^* = \lambda h_s^* = \lambda \chi^{-1/(\eta+\gamma)} (w_s^*/P^*)^{(1-\gamma)/(\eta+\gamma)}$;*
- (iv) *Labor market clears: $\sum_s L_s^* = \lambda \sum_s h_s^*$;*
- (v) *Capital market clears: $\sum_s K_s^* = K$;*
- (vi) *Each goods market clears: $Y_s^* = \lambda c_s^{W*} + (1 - \lambda) c_s^{C*}$.*

Employment outcomes in the scenario table are computed under the separable-CRRA household block with Marshallian elasticity $\varepsilon_M = (1 - \gamma)/(\eta + \gamma) = -0.25$ ($\gamma = 2$, $\eta = 2$). Since every scenario

raises real wages and $\gamma > 1$ places the calibration in the income-effect-dominant regime, employment falls in all scenarios; the paper presents employment as a secondary, specification-dependent result.

Discussion of key assumptions. We treat labor as sector-specific in the static model: workers cannot costlessly reallocate across sectors. A single wage w_s clears each sector’s labor market, since the split of tasks between augmentable and non-augmentable is a technology parameter rather than a worker attribute – the same worker performs both types of tasks. We require $\theta > 1$ (gross substitutes) for a well-defined interior Cournot equilibrium.

The model is static by design. This is a feature, not a limitation. The static framework delivers four propositions with analytical proofs – the natural wage floor (Proposition 2), the surplus split (Lemma 1), the equilibrium wage path (Proposition 3), and the Pareto-improving barrier policy (Proposition 4) – none of which require numerical simulation. The mechanisms are transparent: each result traces back to a small number of structural parameters (θ , f_s , ϕ_s , $\mathcal{L}$) with clear economic interpretations. The dynamic extension of Section 7 is tractable precisely because the static block is analytically clean: embedding $F_t = F_0\delta^t$ in a clean static equilibrium yields Proposition 3 without numerical methods. A fully dynamic model with endogenous capital accumulation and household heterogeneity would add realism but obscure the mechanisms – the wage floor, the deflation dividend, and the policy ranking – that are this paper’s contribution. The static framework is the right tool for identifying those mechanisms cleanly.

4 Data

4.1 The Anthropic Economic Index

The Anthropic Economic Index (AEI) is a repeated cross-section of actual Claude AI conversations, released by Anthropic in partnership with the Federal Reserve Bank of St. Louis (Fan, 2026). Each wave draws a random sample of conversations, strips identifying information, and classifies each conversation to an O*NET occupational task. This paper uses wave R5 (February 2026) for the baseline calibration, and waves R3 (August 2025) and R4 (November 2025) for the dynamic analysis.

Three features of the AEI make it well suited for calibrating the model. First, the data come from *actual* AI usage, capturing real-world adoption decisions. Second, the O*NET task classifier allows aggregation to the 22 SOC-2 groups in the model. Third, each wave records estimates of human task duration (`human_only_hrs`) and AI time savings (`pct_saved`), from which task-level hours saved are computed.²

The model and data are organized by 22 SOC-2 occupation groups rather than industries. Two considerations justify this choice. First, the AEI classifies conversations to O*NET tasks, which are inherently occupation-specific: a nurse performs nursing tasks whether employed in a hospital,

²For example, the task “interpret laws, rulings, and regulations for individuals and businesses” (SOC-2 Legal) records a solo duration of 1.6 hours and an AI time savings of 89 percent, implying 1.4 hours saved per conversation – achieved in 11 minutes of AI-assisted work.

a school, or a corporate wellness program, and the AI adoption pattern follows the task, not the employer’s industry. Occupation is therefore the natural unit at which f_s , ϕ_s , and x_s/L_s are measured. Second, at the 22-group level, SOC-2 occupation groups are broad enough to function like industries: Management produces coordination and decision services, Healthcare Practitioner produces clinical care, Construction produces physical infrastructure. The within-group variation in industry that is lost by working at the occupation level is smaller than the across-group variation in AI exposure that the task mapping captures. Working at the industry level would force us to aggregate heterogeneous occupational task structures into a single industry index, discarding exactly the variation that identifies the wage floor.

4.2 Mapping AEI Variables to Model Primitives

Table 1 shows how each model primitive maps to an empirical counterpart. The task susceptibility f_s and AI productivity weight ϕ_s come from AEI wave R5, which provides the most recent cross-section of task-level coverage and time savings. The AI intensity x_s/L_s is constructed from all three waves (R3–R5) to trace the adoption path needed for barrier calibration. Wages draw on two BLS OEWS vintages: 2019 provides the pre-AI calibration baseline, and 2025 provides the post-AI wages used to estimate θ . Employment shares use BLS May 2019 OEWS, consistent with the pre-AI baseline. The substitution elasticity θ is identified from AEI R3 and BLS data using the pre-trend design described in Section 5.2.

Table 1: AEI and BLS variables and their model counterparts.

Model object	Empirical proxy	Source
f_s (task susceptibility)	Fraction of O*NET tasks in sector s with ≥ 1 AEI conversation	AEI R5; O*NET task list
ϕ_s (AI productivity weight)	Conv.-wtd. avg. hrs saved per conversation $\div$ 2080 (worker-years per conversation)	AEI R5; O*NET task-level
x_s/L_s (AI intensity)	Annual US AI conversations in sector $s \div$ BLS employment (conversations per worker per year)	AEI waves R3–R5; BLS May 2025 OEWS
w_s (sector wage)	Median annual wage by SOC-2	BLS May 2019 & 2025 OEWS (θ est.); BLS May 2019 (calibration table)
L_s (employment share)	Employment share by SOC-2	BLS May 2019 OEWS
θ (AI–labor subst.)	Estimated from adoption condition (Sec. 5.2)	AEI R3; BLS

5 Calibration

5.1 Calibrating the AI Productivity Weight ϕ_s

The parameter ϕ_s measures effective labor units per unit of AI service in the augmentable CES bundle (7). For dimensional consistency with $f_s L_s$ – which is in workers – and x_s measured in conversations, ϕ_s must be in worker-years per conversation. We calibrate it as the conversation-weighted average hours saved per AI interaction in sector s , divided by 2080 (annual hours per worker):

$$\phi_s = \frac{1}{2080} \cdot \frac{\sum_{t \in s} h_t c_t}{\sum_{t \in s} c_t} \quad (24)$$

where h_t is hours saved per conversation on task t (human task time minus AI-assisted completion time, from the AEI), c_t is conversation volume weight, and the sum is over all O*NET tasks t classified to sector s . The division by 2080 converts hours into the worker-year units required by the CES bundle, so that $t_s \equiv \phi_s r_s$ – the product of worker-years saved per conversation and annual conversations per worker – is a dimensionless fraction of annual labor directly comparable to f_s .

The ϕ_s values range from 0.00047 worker-years per conversation (Farming, 0.97 hrs saved per conversation) to 0.00173 (Computer & Math, 3.60 hrs saved per conversation), a 3.7-fold range driven by the duration and cognitive intensity of the underlying tasks. A single Computer & Math conversation therefore saves nearly four times as many worker-hours as a Farming conversation – reflecting AI’s comparative advantage in cognitive, data-intensive tasks over those requiring physical judgment or environmental assessment. The parameter ϕ_s thus encodes cross-sector heterogeneity in how effectively each AI interaction augments worker time. The average conversation in the AEI data replaces 2.14 hours of solo work in 12.5 minutes of AI-assisted time – a tenfold reduction – so even a small number of annual conversations per worker represents a meaningful labor equivalent. Both f_s and ϕ_s are computed from all occupationally-matched conversations irrespective of use-case tag, since both measure technological capacity – what tasks AI can augment and how much time it saves – rather than the volume of AI deployed in production.

5.2 Estimating the Substitution Elasticity θ

We exploit a pre-trend design to identify θ , using five waves of BLS OEWS (May 2015–2019) to establish sector-specific wage trends before AI adoption began, and BLS May 2025 wages combined with AEI Round 3 (August 2025) adoption data as the post-AI observation. The pre-AI window (2015–2019) predates meaningful AI adoption and avoids the COVID-19 disruptions that contaminate 2020–2021 wages. For each sector s we estimate a linear pre-AI wage trend by OLS:

$$\ln w_{s,t} = \alpha_s + \beta_s \cdot t, \quad t \in \{2015, \dots, 2019\}$$

and construct the 2025 residual $\varepsilon_s \equiv \ln w_{s,2025} - (\hat{\alpha}_s + \hat{\beta}_s \cdot 6)$, which measures how much each sector’s 2025 wage deviated from its pre-AI trajectory. This strips out secular wage growth differentials

– high-skill sectors were already growing faster before AI arrived – so only the AI-attributable component of wage dynamics identifies θ . All 18 estimation sectors have positive residuals in 2025, meaning wages exceeded their pre-AI trend, consistent with the deflation dividend story.

In the low-adoption limit ($t \ll f_s$), the denominator of (16) collapses to f_s^ρ and the adoption condition log-linearizes exactly to

$$\ln(x_s/L_s) \approx \ln f_s + (\theta - 1) \ln \phi_s + \theta \ln(w_s/p_{AI}),$$

which differs from the flat-CES limit by replacing $\theta \ln \phi_s$ with $(\theta - 1) \ln \phi_s$ and adding $\ln f_s$. Controlling for $\ln f_s$ and replacing the post-AI wage level with the pre-trend residual ε_s , the coefficient on ε_s identifies θ , while the coefficient on $\ln \phi_s$ identifies $(\theta - 1)$. We impose the theoretical restriction that both coefficients satisfy the same underlying θ , which identifies θ from whether adoption-intensive sectors deviated above or below their pre-AI trajectories in 2025. θ measures how readily firms substitute AI for human labor: sectors where wages ran above trend (making labor more expensive relative to AI) should have adopted more AI, and the steepness of this relationship – controlling for f_s (the share of augmentable tasks) and ϕ_s (AI effectiveness per conversation) – pins down θ . Restricting the coefficient on $\ln \phi_s + \varepsilon_s$ to equal θ (as implied by the theory), using HC3 standard errors, over the 18 sectors with $f_s \geq 0.05$:³

$$\hat{\sigma}_{LA} = 2.079 \quad (\text{se } 0.439, \quad 95\% \text{ CI: } [1.219, 2.940]), \quad \hat{\rho} = 0.519, \quad R^2 = 0.425 \quad (25)$$

The confidence interval now excludes 1 from below, confirming that AI and labor are gross substitutes ($\theta > 1$) – the maintained assumption for a well-defined interior Cournot equilibrium – at conventional significance levels. We adopt $\hat{\theta} = 2.079$ and $\hat{\rho} = 0.519$ as the baseline.

The economic content of $\hat{\theta} = 2.079$ goes beyond a technical parameter. It speaks directly to a question that readers and policymakers care about: are AI and human labor substitutes or complements? The evidence points to substitutes, and meaningfully so. $\theta > 1$ means that as AI becomes cheaper relative to labor, firms shift the augmentable bundle toward AI and away from human augmentable hours – the task-bundle channel operates and compresses wages. Had $\theta < 1$, AI and augmentable labor would be complements: cheaper AI would raise the marginal product of human augmentable hours and push wages up. The calibration points away from that scenario. The estimate of 2.079 also implies moderate substitutability – not the near-perfect substitutability ($\theta \rightarrow \infty$) of a linear model, but well above the Cobb-Douglas benchmark ($\theta = 1$) that separates the two regimes. This places AI software in a different category from physical capital ($\sigma_{KL} \approx 1$) and closer to the high-substitutability range found for skilled-task automation in the prior literature.

³The four excluded sectors – Farming, Construction, Production, and Transport – have fewer than 5% of their O*NET tasks classified as AI-augmentable by the AEI. Their conversations-per-worker are effectively at the noise floor, so they contribute no signal to identifying θ and only add variance. These sectors are nonetheless fully retained in the general equilibrium model and all counterfactual scenarios: their large unaugmentable task shares ($1 - f_s > 0.95$) anchor the natural wage floor, and their output price deflation – driven by AI productivity gains on a small augmentable slice – contributes materially to the economy-wide CPI decline that generates the deflation dividend.

Section 7.3 provides robustness checks for all scenario outcomes at $\theta = 1.219$ and $\theta = 2.940$, the lower and upper bounds of the 95% confidence interval.

5.3 AI Productivity, Wages, and Employment by Sector

Table 2 reports ϕ_s , f_s , w_s , and L_s for all 22 SOC-2 groups. The AI adoption intensity index is $r_s/r_{C\&M} \times 100$, where $r_s = x_s/L_s$ is solved numerically from the nested adoption condition (16) at $\hat{\theta} = 2.079$, with p_{AI} absorbed into a common normalization constant that cancels across sectors.

The qualitative pattern is the same as in the task-based version: Computer & Math ranks first, with high f_s , high ϕ_s , and high w_s all reinforcing each other. Sectors with low f_s (Farming, Construction) face a near-unit wage floor: even unlimited AI leaves their wages close to the no-AI MPL. This is a new prediction of the nested structure.

Table 2: Baseline calibration: AI productivity weight, task susceptibility, wages, and employment by SOC-2 sector.

Sector	ϕ_s (hrs/conv.)	f_s (task frac.)	w_s (\$000) (BLS 2019)	L_s share (BLS 2019)	AI index (C&M = 100)
Management	3.53	0.227	105.7	0.055	38.1
Business & Fin.	3.27	0.320	69.8	0.056	35.2
Computer & Math	3.60	0.513	88.3	0.031	100.0
Arch. & Eng.	3.36	0.159	81.4	0.018	18.1
Life/Phys. Sci.	3.35	0.248	68.2	0.009	24.6
Community & Social	2.29	0.293	46.1	0.015	16.1
Legal	2.84	0.254	81.8	0.008	30.5
Education	3.00	0.391	50.8	0.061	29.8
Arts & Media	2.92	0.330	51.1	0.014	23.6
Healthcare Pract.	1.87	0.165	68.2	0.059	12.8
Healthcare Supp.	1.30	0.103	28.5	0.044	1.9
Protective Svc.	1.11	0.060	41.6	0.024	1.8
Food Prep.	1.15	0.138	24.2	0.092	1.8
Building & Grounds	1.12	0.055	28.3	0.030	0.9
Personal Care	1.06	0.121	26.2	0.022	1.7
Sales	1.32	0.386	29.6	0.098	9.0
Office & Admin.	1.66	0.250	37.6	0.133	8.6
Farming	0.97	0.019	27.2	0.003	0.3
Construction	1.85	0.044	47.4	0.042	1.9
Install./Repair	1.39	0.057	46.6	0.039	2.2
Production	2.98	0.026	36.0	0.062	0.9
Transport	1.11	0.045	32.4	0.085	0.9

Notes: ϕ_s = conv.-weighted avg. hrs saved per AI conversation in sector s (AEI R5). f_s = fraction of O*NET tasks in sector s with AEI coverage (task susceptibility). w_s = median annual wage, BLS May 2019. L_s = employment share, BLS May 2019. AI index = $(r_s/r_{C\&M}) \times 100$, where $r_s = x_s/L_s$ is solved numerically from the nested adoption condition at $\hat{\sigma}_{LA} = 2.079$; C&M = 100.

5.4 Baseline Calibration and Key Moments

Table 3: Baseline parameters and calibrated model moments.

Parameter	Value	Source	Moment
N (AI firms)	2	Claude/OpenAI duopoly	$\mathcal{L} = 24.1\%$
θ	2.079	OLS factor-augmenting (Sec. 5.2)	Interior Cournot price
α (labor share)	0.67	NIPA, US	Capital share = 33%
σ (cross-sector CES)	2.0	Broda-Weinstein (2006)	$\beta = 1.67$
ε (Frisch elas.)	0.5	Chetty (2012) macro estimates	$\eta = 1/\varepsilon = 2$; enters GE wage fixed point
γ (CRRA curvature, welfare)	2.0	Standard macro calibration	Welfare index exponent = $(1 + \eta)/(\eta + \gamma) = 0.75$
p_{AI}	\$0.34/conv.	$\approx 34K$ tokens $\times$ Sonnet blended rate [†]	
mc	\$0.26/conv.	Cournot-consistent (compute + amortized training)	Markup = 24.1%
<i>Derived moments</i>			
Lerner $\mathcal{L} = 1/(N\eta^{\text{agg}})$	0.241	Eq. (22); $\eta^{\text{agg}} \approx \theta$ at R5	24.1% markup
Marshallian elas. $\varepsilon_M = (1 - \gamma)/(\eta + \gamma)$	-0.25	CRRA $\gamma = 2, \eta = 2$	Employment falls when real wages rise
Combined nom. wage elas. (low adoption)	≈ -0.12		$1 - 1/\beta - \rho \approx -0.12$
Combined real wage elas. (low adoption)	+0.28		Real wages still rise on impact
Wage floor: Computer & Math	$\geq 49\%$		$(1 - 0.51) \times \text{MPL}$
Wage floor: Construction	$\geq 96\%$		$(1 - 0.04) \times \text{MPL}$

[†]AEI R5 work conversations average approximately 34,000 tokens (prompt + completion). At Claude Sonnet API list prices of \$3/MTok input and \$15/MTok output, a 50/50 split gives a blended rate of \$9/MTok and a per-conversation cost of $34,000 \times \$9/10^6 \approx \0.31 , rounded to \$0.34. The Lerner markup $\mathcal{L} = 1/(N\theta) = 24.1\%$ and marginal cost $mc = p_{AI}(1 - \mathcal{L}) \approx \0.26 follow from the Cournot equilibrium.

The natural wage floor is the key new quantitative prediction. For Computer & Math ($f_s = 0.51$), even with unlimited AI, wages cannot fall below 49% of the no-AI MPL from the task-bundle channel. For Farming ($f_s = 0.02$), the floor is 98%. The floor is binding only asymptotically; at moderate AI adoption levels, both channels compress wages above the floor. The calibrated r_s values for AEI R5 suggest that most sectors are still far from their floors – adoption is not yet deep enough for the floor to bind – but the floor will become increasingly relevant as AI capabilities and f_s expand.

6 Equilibrium Analysis

6.1 Wage Compression

Proposition 1 (Wage Compression). *Suppose $\theta > 1$ ($\rho \in (0, 1)$) and $\sigma > 1$. Let $\beta \equiv 1 + \alpha(\sigma - 1)$. The wage in sector s satisfies:*

$$w_s = \frac{\alpha p_s Y_s}{L_s} \cdot \Gamma(r_s), \quad \Gamma(r_s) = \frac{(1 - f_s)(\phi_s r_s)^\rho + f_s^\rho}{(\phi_s r_s)^\rho + f_s^\rho} \quad (26)$$

Two channels act on wages in opposing directions as AI adoption deepens (r_s increases):

- (i) **Price-output channel.** AI adoption raises ℓ_s and sector output Y_s , lowering the goods price p_s . The elasticity of sectoral revenue $p_s Y_s$ with respect to effective labor is:

$$\frac{\partial \ln(p_s Y_s)}{\partial \ln \ell_s} = \frac{\beta - 1}{\beta} > 0 \quad (27)$$

This channel alone contributes $1 - 1/\beta > 0$ to the wage elasticity – it raises nominal wages whenever $\beta > 1$, which holds iff $\sigma > 1$. The intuition is that elastic goods demand ($\sigma > 1$) means the output expansion induces enough new demand that nominal revenue $p_s Y_s$ rises even as the unit price falls – the quantity effect dominates the price effect. At the calibrated value $\sigma = 2.0$ (Broda–Weinstein, 2006), this condition holds comfortably. If instead $\sigma < 1$, revenue would fall with ℓ_s and this channel would compress wages in absolute terms as well.

- (ii) **Task-bundle channel.** As r_s rises (deeper AI adoption), $\kappa_s = f_s L_s / \ell_s^{\text{aug}}$ falls, reducing $\Gamma(r_s)$. The elasticity is:

$$\frac{d \ln \Gamma}{d \ln r_s} = -\rho \cdot f_s^{1+\rho} \cdot \frac{(\phi_s r_s)^\rho}{\Gamma(r_s) [(\phi_s r_s)^\rho + f_s^\rho]^2} < 0 \quad \text{for } \rho > 0 \quad (28)$$

This channel compresses wages unconditionally whenever $\theta > 1$ (equivalently, whenever $\rho > 0$, since $\theta = 1/(1 - \rho)$). For $\theta < 1$ ($\rho < 0$), Γ rises with r_s : AI raises wages through complementarity within the augmentable bundle.

The two channels combine into the total intensive-margin wage elasticity:

$$\left. \frac{d \ln w_s}{d \ln \ell_s} \right|_{L_s} = 1 - \frac{1}{\beta} - \rho \cdot \frac{\kappa_s^\rho}{\Gamma(r_s)} \quad (29)$$

Nominal wages fall overall iff $\rho \cdot \kappa_s^\rho / \Gamma(r_s) > 1 - 1/\beta$; real wages fall iff $1 - (1 - \alpha)/\beta - \rho \cdot \kappa_s^\rho / \Gamma(r_s) < 0$. The real-wage condition replaces $1/\beta$ with the smaller $(1 - \alpha)/\beta$ because CPI deflation – sector prices fall at rate α/β per unit of $\ln \ell_s$ as AI raises productivity – offsets part of the nominal wage compression. Real wages therefore always exceed nominal wages in elasticity terms: the deflation dividend α/β raises the threshold at which real wages turn positive. At the baseline calibration, nominal wages fall (−0.12) while real wages still rise (+0.28) at low adoption, precisely because the

deflation more than compensates for the nominal decline. The nominal condition pits two opposing forces against each other. The right-hand side, $1 - 1/\beta > 0$, is the price-output support: AI raises sector output, which raises revenue and pushes nominal wages up through the goods market. The left-hand side, $\rho \cdot \kappa_s^\rho / \Gamma(r_s)$, is the task-bundle compression: as AI crowds out human augmentable hours, κ_s falls and Γ falls, pulling wages down directly through the production function. Nominal wages fall when compression dominates support – when AI is sufficiently substitutable for labor (ρ large) and the human augmentable share has not yet been fully eroded (κ_s still close to one).

At the baseline calibration ($\theta = 2.079$, $\alpha = 0.67$, $\sigma = 2.0$, $\beta = 1.67$, $\rho = 0.519$): the price-output channel alone contributes $1 - 1/\beta = +0.401$ (always positive); the task-bundle channel contributes $-\rho\kappa_s^\rho/\Gamma$, which at low adoption ($\kappa_s^\rho/\Gamma \approx 1$) equals -0.519 . The combined intensive-margin elasticity is approximately -0.12 : nominal wages fall on impact even at low AI adoption, with the task-bundle channel dominating the price-output channel.

Proof. See Appendix A. □

Intuition. The nested structure separates the two channels cleanly. Channel (i) is the goods-market mechanism: AI expansion lowers goods prices, but with elastic goods demand ($\sigma > 1$), nominal revenue rises and this channel partially *supports* wages rather than compressing them. Channel (ii) is new to the nested structure: it operates within the production function itself. As AI deepens in the augmentable bundle, augmentable labor’s marginal product falls relative to its share in the wage (because AI absorbs more of the marginal product of the bundle); this feeds into the total wage through the $\Gamma(r_s)$ factor.

Remark (threshold for real wage gains). The general condition for real wages to rise is equation (29) with positive sign: $1 - (1 - \alpha)/\beta > \rho\kappa_s^\rho/\Gamma_s(r_s)$. Since $\kappa_s^\rho/\Gamma_s(r_s)$ declines as adoption deepens in the calibrated substitutable regime ($\rho > 0$), this condition becomes easier to satisfy as AI deepens. At low adoption ($\kappa_s^\rho/\Gamma \approx 1$), the condition simplifies to the low-adoption approximation $\theta < \beta/(1 - \alpha)$. At the calibrated parameters ($\alpha = 0.67$, $\beta = 1.67$), this low-adoption threshold is:

$$\theta^* = \frac{\beta}{1 - \alpha} = \frac{1.67}{0.33} \approx 5.1$$

Real wages rise on impact whenever $\theta < 5.1$ – a permissive condition satisfied by our point estimate of 2.079 and by every value within the 95% confidence interval [1.219, 2.940]. The deflation dividend (α/β per unit of $\ln \ell_s$) is large enough relative to the task-bundle compression to deliver real wage gains across the entire empirically supported range of substitution elasticities. The general condition is even more permissive away from low adoption: as $\kappa_s^\rho/\Gamma_s(r_s)$ falls below one, the right-hand side of the real-wage condition shrinks and the deflation dividend strengthens relative to compression as the unaugmentable share increasingly dominates the bundle.

Proposition 2 (Natural Wage Floor). *For any $r_s \geq 0$ and any $\theta > 1$:*

$$w_s \geq (1 - f_s) \cdot \frac{\alpha p_s Y_s}{L_s} \tag{30}$$

The bound is tight as $r_s \rightarrow \infty$: in the limit of unlimited AI, wages converge to exactly $(1 - f_s)$ times the Cobb-Douglas MPL. The wage floor is strictly increasing in $1 - f_s$: sectors with more unaugmentable tasks are better protected.

Proof. See Appendix A. □

Implications. The wage floor is an inherent structural protection from the unaugmentable task margin. Even in Computer & Math ($f_s = 0.51$), wages cannot fall below 49% of the no-AI MPL from the task-bundle channel alone (before goods-price adjustment). For Construction ($f_s \approx 0.04$), the floor is 96%. This result has no analog in the A-R flat CES: in that specification, a sufficiently large AI input could in principle drive wages arbitrarily low. In contrast, the nested structure predicts that sectors differ not only in the *rate* of wage compression but in the *depth* to which compression can ultimately reach.

6.2 Surplus Distribution

Lemma 1 (Surplus Split). *Let $\Delta(\cdot) \equiv (\cdot)^* - (\cdot)^0$ denote the change from the no-AI counterfactual, and define gross final-goods output $\text{GDP}^G \equiv \sum_s p_s Y_s$. The augmentation surplus is split as follows:*

$$\Delta(rK) = (1 - \alpha) \Delta \text{GDP}^G > 0 \quad (31)$$

$$\Pi_{AI}^* = (p_{AI}^* - mc) X^* - NF \geq 0 \quad (32)$$

$$\sum_s \Delta(w_s L_s) = \alpha \Delta \text{GDP}^G - \Pi_{AI}^* - \Delta(mc X + NF) \quad (33)$$

Under the two-class structure, gains (31) and (32) accrue exclusively to capitalists. Worker welfare changes by the nominal wage change plus the purchasing-power gain from lower goods prices.

Proof. See Appendix A. □

Lemma 1 gives a simple but powerful accounting identity: every dollar of GDP gain from AI adoption is divided among exactly three claimants. Capital owners receive their standard factor share $(1 - \alpha)$ of gross output gain – the same return they would earn from any productivity improvement. The AI sector captures the markup on every unit of AI service sold, which is the source of gross markup payments flowing to owners of AI infrastructure and model-development assets. Workers receive the remainder: the labor share α of gross output, minus net AI profits $\Pi_{AI}^* = (p_{AI}^* - mc)X^* - NF$, minus the real resource costs (compute and infrastructure) the AI sector draws from the economy. At calibrated adoption levels the resource-cost term $\Delta(mc X + NF)$ is small relative to gross markup rents and the labor share, so the dominant wedge between workers' gross share and their net income is the markup. In the fixed- N counterfactual scenarios, gross markup rents are the main distributional wedge; under free entry, these same gross markup payments finance fixed costs rather than appearing as net profits. This accounting identity makes precise why workers can be harmed even when aggregate output rises – their residual claim is compressed from above by the markup and from below by task displacement, while capital and AI owners capture the bulk

of the surplus. The result holds for any labor bundle that is homogeneous of degree one in (x_s, L_s) , making it robust to the nested specification.

Two objects in equation (32) deserve careful distinction. **Gross markup rents** are $(p_{AI} - mc)X$: total payments by user firms to AI suppliers above inference cost. **Net AI profits** are $\Pi_{AI} = (p_{AI} - mc)X - NF$: what remains after fixed costs are covered. Under the free-entry condition $\pi_i = (p_{AI} - mc)X/N - F = 0$, aggregate net profits are zero and gross markup rents equal exactly NF , financing the fixed costs of compute and infrastructure. Nevertheless, gross markup rents remain economically meaningful: they represent payments by user firms to the owners of scarce AI infrastructure and model-development assets. The distributional wedge facing workers in equation (33) is therefore best interpreted as gross AI markup payments plus resource costs $\Delta(mcX + NF)$, not pure profit after free entry. In the counterfactual scenarios below, N is held fixed at its baseline value of 2, so net profits can be non-zero when X or p_{AI} change; the scenario table reports gross markup rents before fixed costs, as noted.

6.3 The Role of Fixed Cost and Oligopoly Power

Lemma 2 (Fixed Cost, Entry, and Rents). *Let $N^*(F)$ and $p_{AI}^*(F)$ denote the equilibrium number of AI firms and AI price as functions of the fixed cost F , with $\theta > 1$. Assume the regularity condition $p_{AI}^* < mc(1 + \eta^{\text{agg}})/(\eta^{\text{agg}} - 1)$, which ensures the reduced-form equilibrium function $\tilde{G}(p_{AI})$ is strictly increasing over the relevant price range and the equilibrium is unique (verified at calibrated values in Appendix A).*

- (i) $\partial N^*/\partial F < 0$ and $\partial p_{AI}^*/\partial F > 0$.
- (ii) $\partial \mathcal{L}/\partial F > 0$: the Lerner index rises with F .
- (iii) As $F \rightarrow 0$: $N^* \rightarrow \infty$, $p_{AI}^* \rightarrow mc$, $\mathcal{L} \rightarrow 0$.
- (iv) A larger aggregate labor market supports a larger N^* , reducing $\mathcal{L}$.

Proof. See Appendix A. □

Lemma 2 characterizes how the structure of the AI industry – specifically, how high fixed costs are – determines how much rent the sector can extract. High fixed costs act as a barrier to entry: only a small number of firms can cover them at the prevailing markup, so the oligopoly remains tight, prices stay high, and rents flow to AI owners. As fixed costs fall – through cheaper compute, more efficient training, or hardware improvements – more firms can enter profitably, competition erodes the markup, and prices converge toward marginal cost. In the limit, free entry drives rents to zero and the full productivity gain passes to users. Part (iv) adds a market-size dimension: a larger economy supports more firms at the same fixed cost, suggesting that AI markets in larger economies should exhibit lower concentration and smaller markups – a prediction consistent with the cross-country pattern of AI pricing. The lemma connects the technological cost trajectory of AI infrastructure directly to the distributional outcome, making the path of compute costs a first-order determinant of who ultimately captures AI’s surplus.

7 Equilibrium Wage Path and Scenario Analysis

The static model of Sections 3–6 characterizes how wages and surplus are distributed at a given AI price and adoption level. It delivers the wage floor, the surplus split, and the conditions for real wage gains – all analytically, without specifying how AI costs evolve. This section has two parts. The first traces the frictionless theoretical benchmark: it embeds the static equilibrium in a dynamic path by letting fixed costs fall geometrically ($F_n = F_0\delta^n$), with adoption barriers absent ($\tau_s = 1$), and derives the equilibrium wage trajectory (Proposition 3) and the infrastructure floor that can permanently halt long-run recovery (Corollary 1). The second part moves to the calibrated economy: it estimates sector-level adoption barriers τ_s from the three AEI waves and uses them to decompose wage changes across a set of counterfactual scenarios. The two parts are complementary – Proposition 3 establishes the qualitative direction and long-run logic; the scenario analysis quantifies the magnitudes under realistic frictions.

Proposition 3 (Equilibrium Wage Path). (Frictionless benchmark: adoption barriers absent, $\tau_s = 1$.) *Index the static equilibrium by n with $F_n = F_0\delta^n$, $\delta \in (0, 1)$, and $f_{s,n}$ non-decreasing. Under $\theta > 1$:*

- (i) **Compute-cost channel.** *Falling F_n (fixed cost per AI firm) drives $p_{AI,n}^* \downarrow$, deeper AI adoption, higher $\ell_{s,n}$ (the augmented labor bundle), and real wages upward at the combined rate $1 - (1 - \alpha)/\beta - \rho\kappa_s^\rho/\Gamma$ per unit of $\ln \ell_s$ – the price-output support net of the concurrent task-bundle compression (Proposition 1). At the baseline calibration, this net rate is approximately $1 - 0.33/1.67 - 0.519 \approx +0.28 > 0$ at low adoption, so real wages rise monotonically from this channel.*
- (ii) **Task-bundle channel.** *As $p_{AI,n}^*$ falls, $r_{s,n}$ (the AI-to-labor ratio in sector s) rises and $\Gamma(r_{s,n})$ (the task-bundle compression factor, eq. 18) falls. As $r_{s,n}$ rises with each successive equilibrium, $\kappa_{s,n}$ (the share of human augmentable hours in the augmented bundle) falls and Γ declines further, tracing out deepening wage compression across equilibria. At early periods, for sectors where adoption deepens rapidly, this channel can dominate.*
- (iii) **Long-run recovery.** *Two sub-cases arise depending on the path of marginal inference cost.*
 - Case A: $mc > 0$ fixed. As $F_n \rightarrow 0$: $N_n^* \rightarrow \infty$, $p_{AI,n}^* \rightarrow mc$, and $\Pi_{AI,n} \rightarrow 0$. Real wages converge to a finite competitive-AI ceiling – the equilibrium at which AI is priced at marginal cost and gross markup rents are fully eliminated. This ceiling strictly exceeds the current oligopoly equilibrium and corresponds to Scenario S4 in the calibrated analysis. Wages do not diverge to infinity: with $p_{AI} \rightarrow mc > 0$, adoption converges to a finite level and real wages to a finite limit.*
 - Case B: $mc_n \rightarrow 0$ (falling inference costs). If marginal inference cost also declines – through algorithmic efficiency or hardware improvements – then $p_{AI,n}^* \rightarrow 0$, adoption $r_{s,n} \rightarrow \infty$, and real wages grow without bound, provided $f_{s,n} < 1$. The natural wage floor prevents nominal*

wages from collapsing while continuous AI-driven deflation raises real wages. This stronger result requires the additional assumption $mc_n \rightarrow 0$ beyond fixed-cost decline alone.

Both cases require $f_{s,n} < 1$: if AI eventually augments every task, the wage floor disappears and neither limit result need hold.

- (iv) **Wage floor at all n .** By Proposition 2, the natural floor $w_{s,n} \geq (1 - f_{s,n}) \cdot \alpha p_{s,n} Y_{s,n} / L_s$ holds at every period, bounding the depth of the compression phase.

Proof. See Appendix A. □

Proposition 3 assumes that total fixed costs fall over time, driven by the compute component $F_c(n)$. This is a reasonable approximation for chip and training costs, which have followed steep learning curves. It is less plausible for the physical infrastructure component of AI – land, power, cooling, and data-center construction – which is tied to real assets and does not benefit from the same technological progress. If this infrastructure floor is large enough, fixed costs never fall to zero, entry remains limited, and the long-run wage recovery of part (iii) fails to materialize. The following corollary formalizes this.

Corollary 1 (Persistent Infrastructure Cost). *Decompose the fixed cost as $F_n = F_c(n) + F_d$, where $F_c(n) = F_{c,0} \delta^n \rightarrow 0$ is the compute component and $F_d \geq 0$ is a constant infrastructure floor reflecting land, power, and data-center capacity costs. Under $\theta > 1$:*

- (i) *As $n \rightarrow \infty$: $F_n \rightarrow F_d$, so the equilibrium number of firms converges to $\bar{N}^* \equiv N^*(F_d) < \infty$ and the markup converges to $\bar{\mathcal{L}} = 1/(\bar{N}^* \eta^{\text{agg}}) > 0$.*
- (ii) *The long-run wage recovery of Proposition 3(iii) fails: $\tilde{w}_{s,n}$ converges to a finite ceiling rather than growing without bound.*
- (iii) *The ceiling is decreasing in F_d : a larger infrastructure floor permanently supports a higher markup and a lower long-run wage.*
- (iv) *The natural wage floor of Proposition 2 continues to hold at every n , so workers retain earnings from unaugmentable tasks regardless of F_d .*

Proof. By Lemma 2, $N^*(F)$ is strictly decreasing in F and $\mathcal{L}(F) > 0$ for all $F > 0$. Since $F_n \rightarrow F_d > 0$, part (i) follows directly. Part (ii) follows because Proposition 3(iii) required $F_n \rightarrow 0$ to drive $\Pi_{AI,n} \rightarrow 0$; with $\bar{\mathcal{L}} > 0$ the AI sector permanently extracts rents, bounding real wages below the competitive ceiling. Part (iii) is immediate from the monotonicity in Lemma 2(i)–(ii). Part (iv) follows because Proposition 2 is independent of F . □

Comparison with task-based version. In the task-based A-R CES, long-run wages could in principle fall to zero as $r_s \rightarrow \infty$ (no wage floor). In the nested model, the floor $(1 - f_{s,n}) \cdot \alpha p_{s,n} Y_{s,n} / L_s$ ensures wages converge to a positive multiple of the growing no-AI MPL as $n \rightarrow \infty$. This is a

structurally sharper and more economically defensible prediction: workers always retain earnings from the tasks that are inherently human.

By Lemma 2, N_n^* rises and $p_{AI,n}^*$ falls monotonically, deepening AI adoption across all sectors. Two forces act simultaneously on wages: the *compute-cost channel* (falling $p_{AI,n}^*$ raises adoption and, through the price-output channel, real wages) and the *task-bundle channel* (rising $r_{s,n}$ compresses $\Gamma(r_{s,n})$ and nominal wages). At the baseline calibration the real wage elasticity is approximately +0.28 at low adoption, so real wages rise on impact from the compute-cost channel alone. Proposition 3 and Corollary 1 characterize this frictionless path analytically. The remainder of this section introduces adoption barriers $\tau_s > 1$ – the empirically relevant case – calibrates them from the AEI wave data, and uses them to construct counterfactual scenarios that quantify the sources of wage change in the actual economy.

7.1 Calibrating Adoption Barriers

We calibrate adoption barriers from the three AEI waves (R3–R5) using a two-stage procedure. In Stage 1 we construct $x_{s,t}/L_s$ from AEI conversation counts by combining two traffic components: Claude.ai consumer conversations and Anthropic API traffic.

The AEI publishes separate data files for consumer (Claude.ai) and API (1P API) usage. Each consumer file records approximately 190,000–220,000 US conversations classified by O*NET task category, of which roughly 42 percent are tagged as professional work use; the remainder are personal and coursework. We scale each sector’s matched consumer sample count to annual US work-AI conversations using the wave-specific global Claude.ai volume C_t^{con} (414M, 489M, and 800M conversations per week at R3, R4, and R5 respectively; Fan and Nguyen 2026), the US share of 22 percent, and the work fraction $f_{\text{work}}^{\text{con}} \approx 0.42$.

The API files record global usage only, without a US-specific geographic breakdown, but cover all 22 SOC-2 groups with 86–88 percent of conversations matched to the O*NET taxonomy – the same sector distribution used for the consumer component. We scale the API sample counts to annual US work-API conversations using wave-specific global API volumes C_t^{api} , derived from Anthropic’s annualized revenue (\$5B at R3, \$9B at R4, \$14B at R5) multiplied by the API revenue share (70–75 percent; ?), divided by the per-conversation price $p_{AI} = \$0.34$ and by 52 weeks. The work fraction for API conversations is $f_{\text{work}}^{\text{api}} \approx 0.74$, read directly from the AEI use-case facet in the API files (rounds R4 and R5; R3 uses the R4 proxy). The US share of global API conversations is set to 50 percent, reflecting the enterprise-heavy geography of API customers relative to the 22 percent US share of consumer traffic.⁴

⁴API traffic accounts for 70–75 percent of Anthropic revenue but was absent from prior versions of this calibration. Its inclusion raises the estimated US work conversations per worker by approximately $3\times$ at R3 and $4\times$ at R5. Computer & Math has the highest API concentration – 51 percent of matched API conversations versus 36 percent for consumer traffic – reflecting Claude Code’s dominance among professional API users.

The total $x_{s,t}/L_s$ is the sum of the two components:

$$\frac{x_{s,t}}{L_s} = \frac{n_{s,t}^{\text{con}}}{N_t^{\text{con}}} \cdot \frac{C_t^{\text{con}} \cdot f_{\text{work}}^{\text{con}} \cdot 0.22 \cdot 52}{L_s} + \frac{n_{s,t}^{\text{api}}}{N_t^{\text{api}}} \cdot \frac{C_t^{\text{api}} \cdot f_{\text{work}}^{\text{api}} \cdot 0.50 \cdot 52}{L_s} \quad (34)$$

where $n_{s,t}$ is the matched sample count, N_t is the total matched sample, and L_s is BLS May-2025 sectoral employment.⁵ In Stage 2 we invert the adoption condition (16) numerically (since no closed-form exists in the nested model) to recover the effective AI price $p_{\text{eff},s,q}$ at each wave, then fit $\ln p_{\text{eff},s,q} = \ln \tau_{s,0} + q \ln \delta_{\tau,s}$ by OLS, where $q \in \{0, 1, 2\}$ indexes the three waves (R3, R4, R5), each approximately one quarter apart, so $\delta_{\tau,s}$ is a per-quarter barrier decline rate. A level shift common to all sectors changes $\tau_{s,0}$ but leaves $\delta_{\tau,s}$ unchanged, since the decay rate is identified from the inter-wave growth in $x_{s,q}/L_s$, which reflects the documented growth in all-platform conversation volume C_q across waves.

The model parameter $\tau_{s,0}$ equals $p_{\text{eff},s,0}/w_s$, the effective AI price at wave R3 normalized by the sector wage. Because AI conversations are cheap relative to wages, the economically interpretable object is the barrier multiple $B_s \equiv \tau_{s,0} \cdot w_s/p_{\text{AI}}$, where $p_{\text{AI}} = \$0.34$ per conversation is the all-platform market price. B_s measures how many times more expensive AI appears to firms in sector s than it actually costs at the time of first observation (wave R3, August 2025): a value of $B_s = 500$ means firms are currently adopting as if each conversation costs $500 \times \$0.34 = \170 rather than the market price of $\$0.34$, depressing adoption far below the frictionless optimum. The no-barrier benchmark is $B_s = 1$; when barriers are removed ($\tau_s \rightarrow \tau_s^* = p_{\text{AI}}/w_s$), firms adopt up to the physical capacity ceiling $T_{\text{max},s}$, which is scenario S3. Table 4 reports B_s and $\delta_{\tau,s}$ for all 22 sectors.

⁵For example, for Management (SOC-2 11), the combined $x_{s,t}/L_s$ rises from approximately 16 work AI conversations per manager per year at R3 to 49 at R5 – roughly one conversation every five working days at R5, compared to 5.6 under the consumer-only scaling. The model’s frictionless optimum ($\tau_s = p_{\text{AI}}/w_s$) still implies approximately 89,000 conversations per worker per year, so the barrier remains large.

Table 4: Calibrated barrier multiples B_s and quarterly decline rates $\delta_{\tau,s}$.

Sector	B_s	$\delta_{\tau,s}$	R^2
Management	493	0.748	0.987
Business & Fin.	375	0.810	0.992
Computer & Math	120	0.732	0.986
Arch. & Eng.	139	0.749	0.878
Life/Phys. Sci.	61	0.871	0.511
Community & Social	189	0.813	0.959
Legal	446	0.952	0.075
Education	235	0.823	0.959
Arts & Media	47	0.913	0.792
Healthcare Pract.	279	0.878	0.705
Healthcare Supp.	223	0.808	0.992
Protective Svc.	90	0.744	0.825
Food Prep.	241	0.864	0.998
Building & Grounds	191	0.794	0.954
Personal Care	62	0.861	0.973
Sales	137	0.858	0.996
Office & Admin.	103	0.713	0.999
Farming	5	0.684	0.921
Construction	156	0.796	0.971
Install./Repair [†]	103	1.016	0.026
Production*	48	0.888	0.221
Transport	156	0.956	0.249

B_s is the barrier multiple at wave R3 (August 2025): the ratio of the effective AI price facing firms to the all-platform market price $p_{\text{AI}} = \$0.34$ per conversation (approximately 34,000 tokens per work AI conversation at Claude Sonnet API pricing of $\$3/\15 per MTok input/output). For Management, $B_s = 0.00159 \times \$105,700 / \$0.34 \approx 493$; firms in that sector adopt as if each conversation costs $493 \times \$0.34 \approx \168 rather than $\$0.34$. $\tau_{s,0} = p_{\text{eff},s,0} / w_s$ is calibrated by inverting the FOC at R3; w_s is the BLS May 2019 median wage. The no-barrier benchmark is $B_s = 1$. $\delta_{\tau,s}$ is the per-quarter barrier decline rate; $1 - \delta_{\tau,s}$ is the quarterly erosion rate. *Production excluded from wage simulations ($R^2 = 0.22$). [†]Install./Repair: $\delta_{\tau} > 1$ and $R^2 < 0.03$; B_s and $\delta_{\tau,s}$ are imprecisely estimated.

Management ($B_s = 493$) and Legal ($B_s = 446$) have the largest barriers: at R3, these sectors behave as if each AI conversation costs $\$168$ and $\$152$ respectively, rather than the all-platform market price of $\$0.34$. Farming ($B_s = 5$) is near the frictionless frontier, consistent with its very low augmentable task share leaving little room for AI adoption regardless of barriers. The quarterly decline rates $\delta_{\tau,s}$ show how fast barriers are eroding: $\delta_{\tau,s} = 0.748$ for Management implies a 25% per-quarter decline, so a sector starting at $B_s = 493$ at R3 reaches $B_s \approx 154$ within one year. Most sectors show rapid barrier erosion, suggesting the current gap between observed and frictionless adoption is transitional rather than permanent.

Even if barriers were fully eliminated, adoption would still be bounded by a second limit: the physical capacity of the workday. A worker can converse with AI only during working hours; at the sector-average conversation duration d_s minutes, an 8-hour day accommodates at most $8 \times 60/d_s$ conversations per day, yielding a physical ceiling of $r_{\max,s} = f_s \cdot (8 \times 60/d_s) \times 260$ conversations per worker per year. The worker-year equivalent $T_{\max,s} = \phi_s \cdot r_{\max,s}$ serves as the adoption cap in Scenario S3. Appendix E reports $r_{\max,s}$ and $T_{\max,s}$ for all 22 sectors; Table 9 shows how they compare to observed adoption at R5. Two patterns stand out. First, current adoption falls far short of even the physical ceiling in nearly every sector – the ratio $r_{\max,s}/r_{\text{obs}}$ ranges from approximately $3\times$ (Farming) to approximately $580\times$ (Food Prep), confirming that barriers, not physical capacity, are the binding constraint today. Second, the physical ceiling itself is well below the frictionless economic optimum: at the market price of \$0.34 per conversation, the unconstrained FOC solution exceeds $r_{\max,s}$ for every sector, so removing barriers would drive adoption to the physical ceiling, not an interior solution.⁶

7.2 Six-Scenario Wage Analysis

We decompose the sources of wage change into six scenarios that isolate the effects of task expansion, barrier removal, and competitive pricing. Each scenario holds fixed all elements not being varied, so the scenarios are individually interpretable as controlled counterfactuals. In particular, the number of AI firms is held fixed at $N = 2$ across all scenarios; the Cournot markup $\mathcal{L} = 1/(N\eta^{\text{agg}}) = 24.1\%$ therefore reflects the baseline market structure in S1–S3 and S6, while S4 and S5 replace it with marginal-cost pricing by regulatory cap rather than through endogenous entry.

Scenario definitions.

- **S1 – Current equilibrium** ($f_s = f_s^{\text{cur}}$, $p_{\text{eff},s} = \tau_{s,0} \delta_{\tau,s}^2$).

This scenario represents the model’s prediction for the economy as of AEI R5 (February 2026): task susceptibilities at their currently observed values, and effective AI prices at the R5 level implied by the fitted barrier path in Table 4. All percentage changes – in S1 and in S2–S5 – are measured relative to the no-AI baseline (S0: $f_s = f_s^{\text{cur}}$, $x_s = 0$, wages equal BLS May 2019 medians). S1 answers the question of what the model predicts for wages under today’s technology and frictions; S2–S5 show what changes under counterfactual policy or capability assumptions.

- **S2 – Task expansion** ($f_s \rightarrow 1.5 f_s^{\text{cur}}$, effective price unchanged).

AI capabilities improve so that 50% more tasks in each sector become augmentable, while the effective AI price and barriers remain at R5 levels. This captures the *extensive margin* of AI diffusion: newer, more capable models begin handling tasks previously beyond AI’s reach –

⁶Even the physically feasible ceiling – how many conversations fit in an 8-hour workday – is far below the frictionless optimum of 89,000: at the sector-average duration of 15 minutes per conversation, a manager completes at most $8 \times 60/15 \approx 32$ conversations per day; multiplied by 260 working days and the augmentable task share $f_s = 22.7\%$, this physical ceiling is approximately 1,876 conversations per worker per year.

more nuanced legal reasoning, complex engineering analysis, or clinical judgment. The scenario isolates the effect of capability growth from price or barrier changes. Workers in sectors with already high f_s (Computer & Math, Business & Finance) face the largest absolute expansion of the augmentable frontier and therefore the strongest additional compression.

- **S3 – Barrier removed** ($\tau_s \rightarrow \tau_s^* = p_{AI}/w_s$, adoption bounded by physical capacity $T_{\max,s}$). The informational and organizational barriers τ_s that slow AI adoption beyond what the technology’s price alone would predict are fully eliminated. In practice these barriers reflect workers and managers who have not yet learned to integrate AI into their workflows, procurement and compliance hurdles, and organizational inertia. Removing them is conceptually distinct from cutting the AI price (S4): it allows every firm to adopt AI up to the physical capacity of the workday – how many conversations fit within working hours given the sector-average conversation duration – rather than the artificially depressed level imposed by organizational barriers. This physical ceiling $T_{\max,s}$ (not the unconstrained economic optimum) is the binding constraint once barriers are lifted: for Management, the frictionless economic optimum is approximately 89,000 conversations per worker per year, but the physical workday allows at most 1,876. The large estimated B_s values in Table 4 – 493 for Management, 375 for Business & Finance, 446 for Legal – indicate that most sectors are far from even this physical ceiling at R5, leaving substantial unrealized adoption.
- **S4 – Competitive pricing** ($p_{AI} \rightarrow mc$, τ_s unchanged; $p_{\text{eff},s}^{S4} = p_{\text{eff},s}^{R5} \cdot (mc/p_{AI}) = 0.759 p_{\text{eff},s}^{R5}$). The AI duopoly (Lerner index 24.1%) is replaced by marginal-cost pricing, eliminating gross markup rents by forcing $p_{AI} = mc$ and transferring markup payments from AI suppliers to users. Adoption barriers remain at their R5 levels; only the markup is removed. This scenario captures the effect of antitrust intervention, public AI infrastructure investment, or future entry that erodes duopoly rents. Since $mc/p_{AI} = 0.759$, all firms face a uniform 24.1% price reduction, which raises AI adoption and productivity in all sectors but through a much smaller channel than barrier removal.
- **S5 – All combined** (S2 + S3 + S4 simultaneously). All three frictions are removed at once: the task frontier expands by 50%, adoption barriers disappear, and AI is priced at marginal cost. This is a theoretical upper bound on worker welfare gains – not a policy prescription, but a diagnostic that reveals how much of AI’s potential productivity dividend currently remains locked behind capability limits, adoption frictions, and oligopoly rents.
- **S6 – Persistent infrastructure floor** ($F_d = \bar{F}_d > 0$; $\tau_s \rightarrow 1$, $f_s = f_s^{\text{cur}}$). Compute costs fall as in S3–S5, but data-center costs – land, power, and physical infrastructure – remain at a positive floor $\bar{F}_d$, preventing competitive entry. The duopoly markup converges to $\bar{\mathcal{L}} > 0$ rather than zero. Adoption barriers are fully removed, as in S3. This scenario captures the long-run outcome when infrastructure constraints are binding: barrier removal (S3) is

achievable, but the competitive-pricing gain of S4 is structurally inaccessible. The scenario is parameterized by holding p_{AI} at its current duopoly level permanently, so the static outcomes are identical to S3; the distinction is interpretive – S3 is a waystation toward S5 in the baseline model, whereas here it is the permanent ceiling.

The baseline for all comparisons is the *pre-adoption* state with the same f_s^{cur} task structure but zero AI use ($n \rightarrow 0$), so wages equal BLS May 2019 medians. This baseline is distinct from a no-task-access counterfactual ($f_s = 0$): both have the same task structure, but the baseline has no active AI adoption. At $n = 0$, the labor bundle collapses to $\ell_{s,0} = (1 - f_s)^{1-f_s} f_s^{f_s} \equiv \ell_0$. The CES price index is:

$$\text{CPI} = \left[\sum_s \omega_s \left(\frac{p_s}{p_s^0} \right)^{1-\sigma} \right]^{1/(1-\sigma)}, \quad \frac{p_s}{p_s^0} = \left(\frac{\ell_0}{\hat{\ell}_s} \right)^{\alpha/\beta} \quad (35)$$

where $\hat{\ell}_s \equiv \ell_s/L_s$ is per-worker effective labor and $\beta = 1 + \alpha(\sigma - 1)$. This follows directly from Appendix A Step 1, which establishes $p_s \propto \ell_s^{-\alpha/\beta}$ under mobile capital and elastic goods demand ($\sigma > 1$). Equation (35) is a reduced-form general-equilibrium price relation, not a unit-cost formula; it incorporates mobile capital adjustment, CES goods demand, and the sectoral output response jointly. In the scenario calculations, L_s is held at its baseline sectoral employment share when computing the intensive productivity shock, so the relevant object is effective labor per worker, $\hat{\ell}_s = \ell_s/L_s$. The price ratio equals 1 at baseline (since $\hat{\ell}_s = \ell_0$ when $x_s = 0$) and falls below 1 whenever AI adoption raises effective labor productivity above ℓ_0 . Wages and AI adoption r_s are solved jointly from the adoption FOC (16) and the GE wage relationship (Proposition 1).

Table 5 summarizes the aggregate (employment-weighted) outcomes; the full sector-level table is in Appendix B. Figure 1 shows nominal wage changes with CPI overlaid for eight anchor occupations. Figure 2 shows aggregate nominal and real wage changes.

Table 5: Aggregate outcomes across six scenarios (employment-weighted averages).

Sc.	Description	CPI %	Nom. wage %	Real wage %	Employ. %	Worker welfare %	Capitalist welfare %
S1	Current equilibrium	−9.4	−2.3	+7.9	−1.9	+5.9	+6.0
S2	Task expansion ($1.5f_s$)	−14.6	−4.2	+12.2	−2.8	+9.0	+9.2
S3	Barrier removed ($\tau_s \rightarrow \tau_s^*$, bounded by $T_{\max,s}$)	−21.3	−1.5	+25.2	−5.5	+18.4	+19.4
S4	Competitive pricing ($p_{AI} \rightarrow mc$)	−10.8	−2.4	+9.5	−2.2	+7.1	+7.1
S5	All combined	−27.9	−4.6	+32.3	−6.8	+23.4	+23.4
S6	Persistent infra. floor (long-run ceiling) [†]	−21.3	−1.5	+25.2	−5.5	+18.4	+19.4

All figures are percentage changes relative to the no-AI baseline (S0: $f_s = f_s^{\text{cur}}$, $x_s = 0$, wages = BLS May 2019 medians).

Capitalist welfare % = real capital income + gross AI markup rents, both as a percentage of baseline capital income $(1 - \alpha) \cdot \text{GDP}_0$. Real capital income = real GDP change (Cobb-Douglas factor share). Gross AI markup rents before fixed costs = $(p_{AI} - mc) \cdot \sum_s r_s L_s / [\text{CPI} \cdot (1 - \alpha) \cdot \text{GDP}_0]$; zero in S4–S5 where $p_{AI} = mc$.

[†]S6 holds p_{AI} at the duopoly level permanently ($F_d > 0$ prevents entry); static outcomes equal S3.

S3 is a waystation toward S5 in the baseline; S6 is the permanent ceiling (Corollary 1).

Employ. % uses the Marshallian elasticity $\varepsilon_M = (1 - \gamma)/(\eta + \gamma) = -0.25$ implied by the model’s separable CRRA preferences ($\gamma = 2$, $\eta = 2$); employment falls when real wages rise because the income effect dominates the substitution effect. Worker welfare % is the CRRA welfare index $\mathcal{W}_s \propto (w_s/P)^{(1+\eta)/(\eta+\gamma)}$ with $\gamma = 2$, $\eta = 2$, exponent = 0.75; the index rises whenever w_s/P rises for any $\gamma > 0$.

Key findings.

- (i) **Current AI adoption delivers a small but positive deflation dividend.** In S1, nominal wages fall 2.3% and CPI falls 9.4%, leaving real wages 7.9% above the no-AI baseline. Under the model’s separable CRRA preferences ($\gamma = 2$), the Marshallian elasticity $\varepsilon_M = -0.25$ implies the income effect dominates: employment falls 1.9% as workers voluntarily reduce hours in response to higher real income. Worker welfare nonetheless rises 5.9% under the CRRA welfare index ($\gamma = 2$), because $\mathcal{W}_s \propto (w_s/P)^{0.75}$ rises whenever w_s/P rises regardless of the employment margin. The modest magnitude reflects the current low level of AI adoption: informational barriers at $B_s \approx 50\text{--}400$ (R5) restrict most sectors to roughly 10–20 work conversations per worker per year, so the task-bundle compression channel barely activates and the CPI deflation channel dominates.
- (ii) **Task expansion (S2) amplifies the deflation dividend.** Broadening the task frontier by 50% deepens CPI deflation to −14.6% while nominal wages fall 4.2%, raising real wages to +12.2% above the no-AI baseline. Under CRRA ($\gamma = 2$), employment falls 2.8% as the larger real income gain strengthens the income effect; worker welfare rises 9.0%. The real wage gain relative to S1 reflects larger productivity gains – the broader task frontier raises effective labor more than it deepens compression – consistent with Proposition 1.
- (iii) **Barrier removal (S3) is the dominant single instrument by a large margin.** When $\tau_s \rightarrow \tau_s^* = p_{AI}/w_s$, firms adopt AI up to the physical capacity of the workday ($T_{\max,s}$),

eliminating the organizational friction that holds adoption far below even that ceiling. CPI falls 21.3% and real wages rise 25.2% despite a 1.5% nominal decline. The mechanism is the deflation dividend: AI adoption deflates goods prices faster than it compresses nominal wages, especially in low-exposure sectors where the wage floor is binding. Under the model’s CRRA preferences ($\gamma = 2$), the +25.2% real wage gain strengthens the income effect enough to reduce employment by 5.5%. Worker welfare rises 18.4% under the CRRA welfare index with $\gamma = 2$ ($\Delta \ln \mathcal{W}_s = 0.75 \Delta \ln(w_s/P)$); the direction – positive whenever w_s/P rises – holds for all $\gamma > 0$ by the envelope theorem.

- (iv) **Competitive pricing (S4) adds modestly to the current equilibrium.** Moving to marginal-cost pricing reduces the AI markup from 24.1% to zero, raising real wages to +9.5% – 1.6 pp above S1. The markup ($mc/p_{AI} = 0.759$, a 24.1% price reduction) is too small relative to the adoption barrier ($B_s \approx 50\text{--}400$, a multi-thousand-percent effective overcharge) to generate large adoption or welfare gains. This confirms that the barrier, not the markup, is the binding wedge between the current equilibrium and the full deflation dividend.
- (v) **Combined (S5) delivers transformative welfare gains, primarily through the barrier channel.** Removing all three frictions simultaneously compresses CPI by 27.9%, raising real wages 32.3% above the no-AI baseline and worker welfare +23.4% under the CRRA welfare index ($\gamma = 2$). Employment falls 6.8% under the model’s CRRA preferences ($\gamma = 2$) as income effects dominate at large real wage gains. Barrier removal alone (S3) accounts for +25.2 pp real wage and +18.4 pp welfare out of the S5 total, confirming it as the dominant instrument. S5 is the theoretical ceiling when capability limits, adoption frictions, and oligopoly rents are simultaneously eliminated.
- (vi) **Persistent infrastructure costs (S6) permanently foreclose the competitive dividend.** In the baseline model, the long-run outcome approaches S5 as compute costs fall and entry erodes the markup. Under Corollary 1, if infrastructure costs prevent entry, the economy is permanently at S3 (+25.2% real wages) rather than S5 (+32.3%). The gap – roughly 7.1 percentage points of real wage gains – represents the share of AI’s productivity dividend captured by the AI sector through durable oligopoly rents. In this scenario the markup cap of S4 becomes the only mechanism available to recover the competitive dividend, raising its policy value substantially relative to the baseline ranking.

The gap between CPI deflation and nominal wage compression – 21.3 versus 1.5 percentage points under barrier removal – reflects a sector-heterogeneity asymmetry that the wage floor makes precise. Barrier removal activates AI adoption in sectors that previously had very high frictions: Food Preparation, Healthcare Support, Construction, and Transport. Their output rises substantially when barriers fall, pushing down goods prices and contributing heavily to CPI deflation. But since most tasks in these sectors require physical presence – serving food, lifting patients, operating machinery – the task-bundle channel barely operates and wages remain close to their pre-AI level.

This finding contrasts sharply with Computer & Math, where the task-bundle channel compresses wages by up to 14 percent under the combined scenario. Workers consume the cheaper goods that the low-exposure sectors produce while their own sector wages compress by much less. The wage floor is therefore not only a structural bound on worker loss in high-exposure sectors: it is the mechanism that widens the gap between CPI deflation and nominal wage compression and delivers the real wage dividend.

A natural concern is whether AI adoption causes net job displacement. With $\hat{\theta} = 2.08 > 1$, AI and labor are gross substitutes in production: holding output fixed, a fall in the AI price shifts firms toward AI and away from labor. The relevant general-equilibrium outcome is the balance between this within-bundle *substitution effect* and the economy-wide *scale effect*. As AI reduces production costs, goods prices fall; the resulting real wage gain has two opposing effects on labor supply. First, higher real wages raise the opportunity cost of leisure, inducing additional hours (substitution effect). Second, higher real wages raise real income, inducing workers to consume more leisure (income effect). Under separable CRRA preferences, the net employment response depends on the CRRA parameter γ : employment expands when $\gamma < 1$ (substitution effect dominates), is unchanged at $\gamma = 1$ (log utility), and contracts when $\gamma > 1$ (income effect dominates). Since macro calibrations commonly use $\gamma \in [1, 4]$, net displacement is a genuine possibility even when real wages rise. The employment figures in Table 5 – ranging from -1.9% in S1 to -6.8% in S5 – are computed under the model’s separable CRRA preferences ($\gamma = 2$, $\varepsilon_M = -0.25$), which place the calibration firmly in the income-effect-dominant regime: every scenario raises real wages, and under CRRA with $\gamma > 1$ higher real income induces workers to reduce hours. The paper therefore presents employment as a secondary, specification-dependent result; the deflation-dividend mechanism is primarily production-side and the welfare gain from higher w_s/P is robust to the sign of the employment response.

Capitalists receive two income streams: capital income $(1 - \alpha)\Delta\text{GDP}$ and AI gross markup rents $(p_{AI} - mc) \cdot X$ (before fixed costs; net AI profit is $(p_{AI} - mc)X - NF$), where X denotes total AI usage. The capitalist welfare column in Table 5 adds real gross AI markup rents to real capital income: $+6.0\%$ (S1), $+9.2\%$ (S2), $+19.4\%$ (S3), $+7.1\%$ (S4), and $+23.4\%$ (S5). In S1 and S2 gross AI markup rents are negligible; in S3 barrier removal raises adoption enough that gross markup rents add approximately $+1$ percentage point above the real capital income gain; in S4 and S5 competitive pricing eliminates gross markup rents entirely. These track real wage movements closely, as expected: with a Cobb-Douglas outer production function, capital and labor each claim fixed shares $(1 - \alpha)$ and α of real GDP, so both groups benefit symmetrically from the deflation channel.

The distinguishing factor is gross AI markup rents. In scenarios S1–S3, the AI markup is unchanged ($\mathcal{L} = 24.1\%$); barrier removal (S3) raises total AI usage X without altering p_{AI} , so gross markup rents increase – consistent with Proposition 4’s finding that barrier reduction is approximately Pareto improving in the calibrated economy. In S4 and S5, competitive pricing sets $p_{AI} = mc$ and eliminates gross markup rents entirely, redistributing markup payments to all agents

through lower goods prices. Capitalists therefore face a trade-off in S4: real capital income rises only 7.1% (compared with 18.4% in S3) and gross markup rents are eliminated, so S4 benefits workers through lower goods prices but delivers far less to capital than barrier removal. By contrast, S3 delivers both a large capital income gain (+18.4%) and higher gross markup rents, making barrier policy the instrument that unambiguously improves capitalist welfare. The contrast with workers – who gain most from S3 but also benefit from S4 (which redistributes rents toward them) – illustrates the broader distributional logic of Lemma 1: barrier removal expands the surplus pie, while price regulation only reshuffles slices.

Figure 1 shows that the CPI benefit (dashed lines) is *economy-wide*, whereas nominal wage compression is concentrated in high- f_s sectors (Computer & Math, Education). Workers in selected low- f_s sectors (Construction, Transport) experience near-zero nominal wage change but still benefit from the full CPI decline – they are net winners in real terms under all scenarios. This cross-sectional pattern is consistent with the natural wage floor (Proposition 2).




Figure 1: Nominal wage changes by scenario for eight selected occupations, sorted by task susceptibility f_s (high to low). Dashed lines show the aggregate CPI change under each scenario; a bar lying above (below) its dashed line indicates a real wage gain (loss) for that occupation. Bars in blue (S1) and orange (S4) are nearly identical, confirming that competitive pricing alone has a small effect.




Figure 2: Employment-weighted aggregate nominal wage (solid) and real wage (hatched) changes by scenario. Real wages are computed as nominal divided by CPI from equation (35). Under S1–S4, nominal wage declines are partly or fully offset by CPI declines; S5 shows a large real gain driven by the barrier-removal and task-expansion effects compounding.

7.3 Robustness: Substitution Elasticity at 95% CI Bounds

The five-scenario analysis above uses the OLS factor-augmenting point estimate $\hat{\theta} = 2.079$. We check sensitivity by re-running the full analysis at the 95% confidence interval bounds: $\theta = 1.219$ (lower bound) and $\theta = 2.940$ (upper bound). Both bounds remain in the gross-substitutes regime ($\theta > 1$), consistent with the paper’s maintained assumption. In each case τ_s is re-calibrated via the same FOC inversion so that the effective AI price $p_{\text{eff},s}$ reflects the alternative elasticity throughout.

Table 6 compares aggregate outcomes across all three elasticity values; full sector-level detail is in Appendix D.

Table 6: Robustness: aggregate outcomes at the 95% CI bounds for θ .

Scenario	$\theta = 2.079$ (baseline)			$\theta = 1.219$ (CI lower)			$\theta = 2.940$ (CI upper)		
	CPI %	Nom. %	Real %	CPI %	Nom. %	Real %	CPI %	Nom. %	Real %
S1 Current	−9.4	−2.3	+7.9	−23.3	+6.5	+39.0	−3.7	−2.0	+1.7
S2 Task expansion	−14.6	−4.2	+12.2	−34.8	+9.6	+68.2	−5.9	−3.4	+2.7
S3 Barrier removed	−21.3	−1.5	+25.2	−36.4	+12.8	+77.5	−19.9	−3.4	+20.6
S4 Comp. pricing	−10.8	−2.4	+9.5	−25.2	+7.2	+43.4	−4.5	−2.3	+2.2
S5 All combined	−27.9	−4.6	+32.3	−49.1	+17.3	+130.3	−25.6	−7.6	+24.3

Interpretation. Three findings stand out. First, real wages rise under barrier removal (S3) regardless of θ , and barrier removal dominates competitive pricing (S4) across the entire confidence interval.

Nominal wages fall at the baseline and upper CI, but actually rise at the lower CI bound ($\theta = 1.219$): weaker AI-labor substitutability strengthens the price-output support channel relative to task-bundle compression, so firms' nominal revenue per worker expands rather than contracts. Second, lower substitutability amplifies real wage gains. At $\theta = 1.219$, AI and augmentable labor are weaker substitutes, so AI adoption raises the effective labor bundle ℓ_s by more for a given price decline, passing through to a larger CPI fall. In the current-equilibrium scenario (S1), the CPI falls an additional 13.9 pp at $\theta = 1.219$ relative to baseline, shifting real wages from +7.9% to +39.0%. The Lerner index rises from 24.1% to 41.0% at $\theta = 1.219$, partially attenuating the competitive-pricing gain (S4), but the barrier channel (S3) is large enough that the real gain rises from +25.2 pp to +77.5 pp. Third, higher substitutability at $\theta = 2.940$ attenuates but does not reverse the results: real wages still rise +20.6 pp under barrier removal and +24.3 pp under the combined scenario. The baseline analysis at $\hat{\theta} = 2.079$ is therefore *conservative relative to the lower CI bound*: if the true elasticity lies toward the lower end of our estimate, worker gains are substantially larger.

8 Optimal Regulation

The scenario analysis in Section 7 identifies two distinct frictions that prevent workers from capturing AI's productivity dividend: the oligopoly markup ($\mathcal{L} = 24.1\%$) and the informational adoption barriers τ_s . Their quantitative magnitudes differ sharply – barrier removal (S3) raises real wages to +25.2% above the no-AI baseline (+17.3 percentage points above the current equilibrium), versus +1.6 percentage points above the current equilibrium for competitive pricing – so the regulator's choice of instrument matters as much as the regulator's objective. This section formalizes those scenarios as policy instruments. Scenario S3 ($\tau_s \rightarrow \tau_s^* = p_{AI}/w_s$) is the formal counterpart of Proposition 4: reducing adoption barriers is not merely one scenario among many, but is approximately Pareto improving in the calibrated economy – AI suppliers receive higher gross markup payments, capital owners gain from higher output, and workers gain through higher real wages simultaneously. Scenario S4 ($p_{AI} \rightarrow mc$) corresponds to the Ramsey cap analyzed in Section 8.1. We analyze each instrument in turn and show that, in the calibrated economy, barrier reduction approximately Pareto-dominates price regulation, because it expands the productivity pie rather than redistributing it.

8.1 Price Regulation: The Ramsey Cap

A benevolent regulator sets a price cap $\bar{p}_{AI} \leq p_{AI}^*$ to maximize the weighted welfare:

$$\Omega = \mu \cdot V^W(\bar{p}_{AI}) + (1 - \mu) \cdot V^C(\bar{p}_{AI}), \quad \mu \in [0, 1] \quad (36)$$

subject to each AI firm's participation constraint $(\bar{p}_{AI} - mc) \cdot X(\bar{p}_{AI})/N \geq F$. The **Ramsey price** $\bar{p}_{AI}^R$ sets the constraint binding.

Remark (Price Regulation).

- (i) Workers unambiguously prefer lower AI prices: $\partial V^W / \partial \bar{p}_{AI} < 0$ for all $\bar{p}_{AI} \in (\bar{p}_{AI}^R, p_{AI}^*]$.
- (ii) For μ sufficiently large, the optimal cap equals the Ramsey price $\bar{p}_{AI}^R$.
- (iii) Price regulation is redistributive under Cournot: a cap reduces AI profits ($\partial \Pi_{AI} / \partial \bar{p}_{AI}|_{p_{AI}^*} = X^*(N-1)/N > 0$), so capitalists gain from a cap only if the capital-income gain $(1 - \alpha)|\partial \text{GDP} / \partial \bar{p}_{AI}|$ exceeds the profit loss $X^*(N-1)/N$. This need not hold in general.
- (iv) There exists $\bar{\mu} \in [0, 1)$ such that $\partial \Omega / \partial \bar{p}_{AI}|_{p_{AI}^*} < 0$ for all $\mu > \bar{\mu}$: a sufficiently worker-weighted regulator always imposes the cap.

Proof. See Appendix A. □

Quantitative scope. Even full competitive pricing (S4) yields only +1.6 pp in real wages above the current equilibrium. A price cap can close at most this gap, and does so at the cost of redistributing AI rents from capitalists. As N grows and $\mathcal{L} \rightarrow 0$, price regulation becomes negligible as a welfare instrument.

8.2 Adoption Barrier Policy

The barriers τ_s in Table 4 are not taxes that accrue to any agent: they represent real resource costs of learning, procurement, and organizational adaptation that dissipate without entering any income stream. A policy that reduces τ_s — through worker training, AI literacy programs, regulatory harmonization, or procurement reform — eliminates these frictional costs and releases the productivity dividend to all factor owners simultaneously.

Proposition 4 (Barrier Policy is Approximately Pareto Improving). *Suppose τ_s falls marginally for some sector s , with p_{AI} held fixed. Then:*

- (i) *AI firm profits rise: $d\tau_s < 0 \Rightarrow d\Pi_{AI} > 0$, equivalently $\partial \Pi_{AI} / \partial \tau_s < 0$.*
- (ii) *Capital income rises: $d\tau_s < 0 \Rightarrow d(rK) > 0$, equivalently $\partial(rK) / \partial \tau_s < 0$.*
- (iii) *Worker real welfare rises when the labor-share gain from higher gross output exceeds the additional AI markup and resource-cost claims: $\alpha d\text{GDP}^G > d\Pi_{AI} + d(mcX + NF)$. This condition holds in the calibrated economy: scenario S3 ($\tau_s \rightarrow \tau_s^*$ for all s) raises real wages by 25.2% above the no-AI baseline, with $\alpha \cdot \Delta \text{GDP}^G$ far exceeding $\Delta \Pi_{AI} + \Delta(mcX + NF)$.*

Proof. See Appendix A. □

Contrast with price regulation. A price cap transfers markup from AI firms to users: redistributive, with limited quantitative scope. Barrier policy eliminates frictional deadweight loss and shares the gain across all factor owners: AI suppliers receive higher gross markup payments from higher volume, capital gains from higher output, and workers gain through higher real wages. The contrast explains why price regulation requires a welfare weight $\mu > 0$ to be optimal, while barrier reduction is desirable for all μ .

9 Discussion of Modeling Choices

The model makes four structural choices – sector-specific labor markets, separable CRRA utility, a two-class household structure, and a static framework – each motivated by tractability and empirical plausibility. This section explains each choice and assesses how the central results would change if it were relaxed.

Labor mobility. The model treats sectoral labor markets as separate, with no cross-sector reallocation in response to wage differences. This reflects the short-to-medium run horizon of the analysis: workers do not move freely between SOC-2 occupation groups in response to cyclical wage changes, and the model captures each static equilibrium at a given AI cost level rather than a long-run steady state. If labor were instead perfectly mobile, the real wage result would be preserved and, in fact, strengthened. In the sector-specific model, real wages rise in every sector in S3 – including Computer & Math, where a nominal compression of -1% is more than offset by economy-wide CPI deflation (-20.6%). Since workers are already better off in real terms everywhere they stand, there is no sector to flee from in real terms. Labor mobility would redistribute the real wage gains by equalizing them across sectors – drawing workers out of high- f_s sectors moderates compression there while adding labor supply to low- f_s sectors – but cannot reverse any of them. The effect of mobility is to compress the dispersion of real wage gains, not to change their sign.

Separable CRRA utility. Workers’ preferences use separable CRRA – $U_s^W = C_s^{1-\gamma}/(1-\gamma) - \chi h_s^{1+\eta}/(1+\eta)$ – the standard workhorse in quantitative macro. The *real wage result* is a production-side outcome – CPI deflation outpacing nominal wage compression – and is independent of the utility specification; by the envelope theorem, worker welfare is strictly increasing in w_s/P for all $\gamma > 0$. The *employment result* depends on the CRRA parameter: when $\gamma < 1$ the substitution effect dominates and employment expands with real wages; when $\gamma \geq 1$ the income effect is large enough to mute or reverse the hours response. The employment figures in Table 5 are computed under the model’s own calibration ($\gamma = 2$, Marshallian elasticity $\varepsilon_M = -0.25$), which places the calibration in the income-effect-dominant regime: employment falls in every scenario because every scenario raises real wages. This is the honest implication of the CRRA preference specification at standard macro calibrations of risk aversion. The real wage gain remains the central finding; worker welfare rises regardless of the sign of the employment response, because $\mathcal{W}_s \propto (w_s/P)^{0.75}$ increases whenever w_s/P increases for any $\gamma > 0$.

Two-class household structure. The model divides households into workers (wage income only) and capitalists (capital income plus AI oligopoly profits), reflecting the empirical concentration of corporate equity ownership. While this abstracts from the middle class – workers who hold capital through pension funds or retirement accounts – the main results are robust to partial relaxation. The surplus-split identity (Lemma 1) is an accounting identity that holds for any household structure: every dollar of GDP gain divides among factor income, AI rents, and the labor residual in fixed proportions. Workers with partial capital holdings benefit from both the wage and capital income channels, which would only strengthen the Pareto improvement result

of Proposition 4. The two-class assumption is therefore conservative: it maximizes the apparent conflict between workers and capitalists, making the finding that barrier removal benefits all parties simultaneously more striking rather than less.

Static framework. The model is static by design: it characterizes each equilibrium at a given AI cost level F_n rather than solving an optimization problem with capital accumulation and household Euler equations.⁷ Proposition 3 recovers a wage path by indexing static equilibria across a geometric sequence of cost levels, delivering the three-phase prediction analytically. A fully dynamic model would add capital accumulation (higher early GDP raises the future capital stock, amplifying long-run wage gains), forward-looking behavior (workers anticipate future real wage growth and adjust present labor supply), and endogenous entry timing. These refinements would primarily affect magnitudes and transition speeds rather than qualitative direction. The wage floor (Proposition 2), the CPI–wages asymmetry, and the Pareto-improving nature of barrier removal are period-specific properties of the production structure and market equilibrium that hold in any single period regardless of intertemporal considerations.

Fixed aggregate capital. The model fixes the aggregate capital stock K , abstracting from the endogenous accumulation dynamics of the standard neoclassical model. In that framework, an AI-driven productivity gain would raise the marginal product of capital, induce investment, and drive K to a higher steady state. The channel through which rising K amplifies real wages is capital–labor complementarity: under the Cobb–Douglas production function, the marginal product of labor $\alpha Y_s/\ell_s$ is increasing in K_s , so capital deepening bids up nominal wages directly. The effect on goods prices is more limited: marginal cost $MC_s \propto w^\alpha r^{1-\alpha}$ is roughly invariant to K , because the fall in the rental rate r is offset by the rise in w in the cost function. Real wages therefore rise primarily through higher nominal earnings rather than through additional CPI deflation. Fixed K omits this complementarity channel, so the real wage gains in Table 5 understate what endogenous accumulation would deliver.

10 Conclusion

The central finding of this paper is an optimistic one: AI-driven deflation outpaces nominal wage compression, so real wages can rise even as AI adoption deepens. In our calibration, removing adoption barriers causes the consumer price index to fall by more than fourteen times as much as nominal wages fall – a gap of nearly 20 percentage points that translates directly into large real wage gains. Under the model’s separable CRRA preferences ($\gamma = 2$), higher real income induces workers to reduce hours – employment falls 5.5 percent under barrier removal – because the income

⁷A related simplification concerns the Cournot markup in Scenarios S3 and S5. When barrier removal drives several sectors to their physical capacity ceiling $T_{\max,s}$, those sectors operate on an inelastic margin ($\eta_s \rightarrow 0$), which would technically reduce η^{agg} and raise the equilibrium markup above $\mathcal{L} = 24.1\%$ under full Cournot re-optimization. Holding the markup fixed at its baseline level isolates the adoption-barrier channel from endogenous market-structure responses. If capacity constraints reduce demand elasticity as adoption rises, equilibrium markups would increase under full Cournot re-optimization, dampening the worker gains from barrier removal in S3 and S5 without reversing the sign of the deflation dividend.

effect dominates at standard macro calibrations of risk aversion. This finding is not a coincidence of parameter values. It is a structural consequence of the model’s central departure from the standard task-based framework.

The mechanism is the natural wage floor: because some tasks in every sector require physical presence, irreducible judgment, or legal accountability, wages are bounded below even as AI adoption deepens. This floor is what generates the asymmetry between price deflation (broad) and wage compression (concentrated and bounded). The robustness of this result is striking: real wages rise whenever $\theta < \beta/(1 - \alpha)$, a condition that is far from binding – our point estimate lies well below the threshold, and every value within the 95% confidence interval satisfies it comfortably. The deflation dividend is not a knife-edge result – it holds across the entire empirically supported range of AI-labor substitutability.

The distributional picture is more nuanced than either the pessimistic or optimistic narrative alone. The model’s income-accounting identity shows that every dollar of AI-driven GDP gain is divided among three claimants: capital owners (factor share), AI firms (oligopoly markup), and workers (the residual). Under barrier removal, AI suppliers receive higher gross markup payments, capital owners gain from higher output, and workers gain through higher real wages – the Pareto logic of Proposition 4. But the division is not equal. Capital and workers share proportionally through the Cobb-Douglas outer production function; AI firms extract a significant markup on every conversation. The policy implication is that removing adoption barriers is approximately the first-best instrument: it tends to expand the pie for all parties. Price caps are redistributive – they transfer markup from AI firms to users – and deliver an order of magnitude less real wage gain than barrier removal. The ranking reverses only when infrastructure costs permanently block competitive entry, at which point the rent cap becomes the only available lever.

The gap between the current deflation dividend and the potential one under barrier removal is the sharpest quantitative result of the paper. In the current equilibrium AI delivers a real wage gain of 7.9 percent; barrier removal raises that figure to 25.2 percent, a 3.2-fold difference. This gap is not driven by the oligopoly markup – competitive pricing adds only 1.6 percentage points above the current equilibrium. It is driven almost entirely by organizational and informational frictions that keep firms adopting at a tiny fraction of the frictionless optimum: Management behaves as if each AI conversation costs \$94 rather than the market price of \$0.34, a barrier multiple of 276. Closing even part of this wedge – through worker AI-literacy programs, procurement reform, or API standardization – releases far more of the productivity dividend than any antitrust remedy could.

The speed at which barriers erode differs sharply across sectors, and this heterogeneity determines which workers capture the dividend first. Most sectors show rapid erosion: Management’s barrier multiple falls at roughly 25 percent per quarter, suggesting the current friction is largely transitional. Legal stands apart, with a quarterly decline rate below 1 percent ($\delta_{\tau, \text{Legal}} = 0.972$), implying that without active intervention Legal workers could remain near the current low-dividend equilibrium long after workers in high-adoption sectors have captured the bulk of the deflation

gain. This finding redirects policy attention: the priority is not to regulate AI prices uniformly, but to identify and address the specific institutional and regulatory frictions that slow adoption in particular occupations.

References

- Daron Acemoglu and Pascual Restrepo. The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6): 1488–1542, 2018.
- Daron Acemoglu and Pascual Restrepo. Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2):3–30, 2019.
- James Bessen. The new goliaths: How corporations use software to dominate industries, kill competition, and undermine regulation. *Yale University Press*, 2022.
- Erik Brynjolfsson, Danielle Li, and Lindsey R. Raymond. Generative AI at work. *NBER Working Paper*, (31161), 2023.
- Avinash K. Dixit and Joseph E. Stiglitz. Monopolistic competition and optimum product diversity. *American Economic Review*, 67(3):297–308, 1977.
- Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock. GPTs are GPTs: An early look at the labor market impact potential of large language models. *Science*, 384(6702):1306–1308, 2023.
- Rachel Yuting Fan. Ai adoption intensity and occupational breadth: Evidence from the anthropic economic index. *Working Paper*, 2026.
- Rachel Yuting Fan and Ha Minh Nguyen. Aggregate gains from AI and their distribution: Global evidence from usage data. *IMF Working Paper*, (26/147), 2026.
- Anton Korinek and Joseph E. Stiglitz. Artificial intelligence, globalization, and strategies for economic development. *NBER Working Paper*, (28453), 2021.
- Anton Korinek and Donghyun Suh. Scenarios for the transition to AGI. *Working Paper*, 2024.
- Shakked Noy and Whitney Zhang. Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654):187–192, 2023.
- Pascual Restrepo. We won’t be missed: Work and growth in the AGI world. In Ajay Agrawal, Erik Brynjolfsson, and Anton Korinek, editors, *The Economics of Transformative AI*. University of Chicago Press, 2025.
- Joseph Zeira. Workers, machines, and economic growth. *Quarterly Journal of Economics*, 113(4): 1091–1117, 1998.

A Proofs

Proof of Proposition 1 (Wage Compression)

Proof. Step 1: Price as a function of effective labor. With mobile capital, the capital FOC gives $K_s = (1 - \alpha)p_s Y_s / r$. Substituting into production and combining with CES goods demand $p_s Y_s \propto Y_s^{(\sigma-1)/\sigma}$:

$$p_s \propto \ell_s^{-\alpha/\beta}, \quad \beta = 1 + \alpha(\sigma - 1) \quad (37)$$

Step 2: Revenue in sector s. $p_s Y_s \propto \ell_s^{(\beta-1)/\beta}$.

Step 3: Wage. From (14) and (9):

$$w_s = \frac{\alpha p_s Y_s}{\ell_s} \cdot \frac{\ell_s}{L_s} [(1 - f_s) + f_s \kappa_s^\rho] = \frac{\alpha p_s Y_s}{L_s} \cdot \Gamma(r_s)$$

where $\Gamma(r_s) = (1 - f_s) + f_s \kappa_s^\rho$. Substituting $\kappa_s^\rho = f_s^\rho / [(\phi_s r_s)^\rho + f_s^\rho]$ yields equation (18).

Step 4: Intensive-margin elasticity (part i). With L_s fixed, ℓ_s varies through x_s (so $d \ln r_s = d \ln x_s$). Decompose:

$$\left. \frac{d \ln w_s}{d \ln \ell_s} \right|_{L_s} = \frac{\partial \ln(p_s Y_s)}{\partial \ln \ell_s} + \frac{d \ln \Gamma}{d \ln r_s} \cdot \frac{d \ln r_s}{d \ln \ell_s}$$

The first term is $(\beta - 1)/\beta = 1 - 1/\beta$ from Step 2 (with L_s fixed, $p_s Y_s / L_s$ moves only through $p_s Y_s$). For the second term: from the inner CES, $\partial \ln \ell_s / \partial \ln x_s = f_s(1 - \kappa_s^\rho)$, so $d \ln r_s / d \ln \ell_s = 1/[f_s(1 - \kappa_s^\rho)]$. From equation (28) in Step 5, $d \ln \Gamma / d \ln r_s = -\rho f_s \kappa_s^\rho (1 - \kappa_s^\rho) / \Gamma$. The $(1 - \kappa_s^\rho)$ factors cancel:

$$\left. \frac{d \ln \Gamma}{d \ln \ell_s} \right|_{L_s} = -\frac{\rho \kappa_s^\rho}{\Gamma}$$

Summing gives equation (29). The pure price-output contribution (holding r_s also fixed, so Γ constant) is $1 - 1/\beta > 0$. The real-wage elasticity is $1 - (1 - \alpha)/\beta - \rho \kappa_s^\rho / \Gamma$, using $\partial \ln P / \partial \ln \ell_s = -\alpha/\beta$ from Step 1.

Step 5: Task-bundle channel (part ii). $d\Gamma/dr_s$: Let $u \equiv (\phi_s r_s)^\rho$ so $\Gamma = [(1 - f_s)u + f_s^\rho] / (u + f_s^\rho)$.

$$\frac{d\Gamma}{du} = \frac{(1 - f_s)(u + f_s^\rho) - [(1 - f_s)u + f_s^\rho]}{(u + f_s^\rho)^2} = \frac{-f_s^{1+\rho}}{(u + f_s^\rho)^2} < 0$$

For $\rho > 0$: u is increasing in r_s , so $d\Gamma/dr_s < 0$ (AI compresses wages). For $\rho < 0$: u is decreasing in r_s , so $d\Gamma/dr_s > 0$ (AI raises wages through complementarity). Equation (28) follows from the chain rule. $\square$ $\square$

Proof of Proposition 2 (Natural Wage Floor)

Proof. Since $\Gamma(r_s) \geq 1 - f_s$ for all $r_s \geq 0$ (established in Section 3.3), equation (17) gives $w_s \geq (1 - f_s) \cdot \alpha p_s Y_s / L_s$. Tightness follows from $\Gamma(r_s) \rightarrow 1 - f_s$ as $r_s \rightarrow \infty$, which requires $(\phi_s r_s)^\rho \rightarrow \infty$

(true for $\rho > 0$, i.e., $\theta > 1$). The floor is increasing in $1 - f_s$, with $f_s = 0$ giving a floor equal to the full MPL (no AI effect) and $f_s \rightarrow 1$ giving floor $\rightarrow 0$. $\square$ $\square$

Proof of Lemma 1 (Surplus Split)

Proof. Step 1: Capital income. The capital FOC gives $rK_s = (1 - \alpha)p_s Y_s$ in every sector. Summing: $rK = (1 - \alpha)\text{GDP}^G$. Differencing: $\Delta(rK) = (1 - \alpha)\Delta\text{GDP}^G$.

Step 2: Labor income residual. By Euler's theorem (12): $p_s \alpha Y_s / \ell_s \cdot (L_s \partial \ell_s / \partial L_s + x_s \partial \ell_s / \partial x_s) = p_s \alpha Y_s / \ell_s \cdot \ell_s = \alpha p_s Y_s$. From the FOCs: $w_s L_s + p_{AI} x_s = \alpha p_s Y_s$. Summing over s and using $p_{AI} X = \Pi_{AI} + mc X + NF$:

$$\sum_s w_s L_s = \alpha \text{GDP}^G - \Pi_{AI} - mc X - NF.$$

The national income identity is $\text{GDP}^G \equiv \sum_s w_s L_s + rK + \Pi_{AI} + mc X + NF$, and Walras' law is satisfied. Differencing gives (33). $\square$ $\square$

Proof of Lemma 2 (Fixed Cost, Entry, and Rents)

Proof. The equilibrium is jointly characterized by two conditions. The **Cournot condition** gives:

$$p_{AI} = \frac{mc \cdot N \eta^{\text{agg}}(p_{AI})}{N \eta^{\text{agg}}(p_{AI}) - 1} \quad (38)$$

and the **free-entry condition** gives $N = (p_{AI} - mc) X(p_{AI}) / F$. Eliminating N between the two yields the reduced-form equilibrium equation:

$$\tilde{G}(p_{AI}) \equiv \frac{\eta^{\text{agg}}(p_{AI}) \cdot (p_{AI} - mc)^2 \cdot X(p_{AI})}{p_{AI}} = F \quad (39)$$

$X(p_{AI})$ is strictly decreasing for $\theta > 1$. From the adoption condition (16), let $h(t) = f_s \phi_s t^{\rho-1} / [(1 - f_s) t^\rho + f_s^\rho]$.

Differentiating with respect to t :

$$h'(t) \propto t^{\rho-2} [-(1 - f_s) t^\rho + (\rho - 1) f_s^\rho]$$

For $\rho \in (0, 1)$: both terms are negative, so $h'(t) < 0$. Since $h(t) = p_{AI} / w_s$, a higher p_{AI} implies lower t_s and lower $r_s = t_s / \phi_s$ in every sector. Therefore $X(p_{AI})$ is strictly decreasing.

$\tilde{G}(p_{AI})$ is strictly increasing. Taking the logarithmic derivative:

$$\frac{d \ln \tilde{G}}{d \ln p_{AI}} = \underbrace{\frac{d \ln \eta^{\text{agg}}}{d \ln p_{AI}}}_{\geq 0} + \underbrace{\frac{p_{AI} + mc}{p_{AI} - mc}}_{\equiv A} - \eta^{\text{agg}}$$

The first term is non-negative because η^{agg} is increasing in p_{AI} (established in Section 3.4). For the remaining terms, $A > 1$, so a sufficient condition for $d \ln \tilde{G} / d \ln p_{AI} > 0$ is:

$$p_{AI} < \frac{mc(1 + \eta^{\text{agg}})}{\eta^{\text{agg}} - 1}$$

Since $mc(1 + \eta)/(\eta - 1)$ is decreasing in η and $\eta^{\text{agg}} < \theta$, the right-hand side exceeds its flat-CES counterpart $mc(1 + \theta)/(\theta - 1)$. At the baseline calibration ($mc = \$0.26$, $\theta = 2.079$): the threshold is $\$0.26 \times 3.079/1.079 = \0.741 , well above the observed $p_{AI}^* = \$0.34$. The condition is verified at the calibrated baseline and throughout the scenario range considered below, so $\tilde{G}^{-1}$ exists and $p_{AI}^*(F) = \tilde{G}^{-1}(F)$ is strictly increasing in F .

Parts (i)–(iv) now follow. For (i): p_{AI}^* increases in F (monotonicity of $\tilde{G}$); lower p_{AI} raises r_s in every sector, so X rises and from free entry $N^* = (p_{AI}^* - mc)X^*/F$ falls. For (ii): $\mathcal{L} = 1 - mc/p_{AI}^*$ is strictly increasing in p_{AI}^* , hence in F . For (iii): $\tilde{G}(p_{AI}) \rightarrow 0$ requires $(p_{AI} - mc)/p_{AI} \rightarrow 0$, i.e., $p_{AI} \rightarrow mc$ and $\mathcal{L} \rightarrow 0$; then $N^* = 1/(\mathcal{L}\eta^{\text{agg}}) \rightarrow \infty$. For (iv): in an economy with all L_s scaled by $\lambda > 1$, total demand scales to $\lambda X(p_{AI})$ at any given p_{AI} , so $\tilde{G}_\lambda(p_{AI}) = \lambda \tilde{G}(p_{AI})$. The equilibrium condition $\tilde{G}_\lambda(p_{AI}^{*,\lambda}) = F$ gives $\tilde{G}(p_{AI}^{*,\lambda}) = F/\lambda < F$. By monotonicity of $\tilde{G}$: $p_{AI}^{*,\lambda} < p_{AI}^*$, so $\mathcal{L}^{*,\lambda} = 1 - mc/p_{AI}^{*,\lambda} < \mathcal{L}^*$ and $N^{*,\lambda}$ is larger. The formula $N^* = p_{AI}^*/[\eta^{\text{agg}}(p_{AI}^* - mc)]$ (from equating the two equilibrium conditions) is used throughout. $\square$ $\square$

Proof of Proposition 3 (Equilibrium Wage Path)

Proof. Part (i): By Lemma 2(i), $p_{AI,t}^*$ falls monotonically. Lower p_{AI} raises r_s from (16), raising $\ell_{s,t}$. By Proposition 1, real wages rise at rate $1 - (1 - \alpha)/\beta - \rho \kappa_s^\rho / \Gamma_s(r_s)$; the expression $1 - (1 - \alpha)/\beta - \rho$ is the low-adoption approximation valid when $\kappa_s^\rho / \Gamma_s(r_s) \approx 1$.

Part (ii): Rising $r_{s,n}$ lowers $\kappa_{s,n}$ and $\Gamma(r_{s,n})$, adding the task-bundle compression of Proposition 1(ii).

Part (iii): Two cases arise.

Case A: $F_n \rightarrow 0$ with $mc > 0$. By Lemma 2, $N_n^* \rightarrow \infty$, $p_{AI,n}^* \rightarrow mc$, and $\Pi_{AI,n} \rightarrow 0$. Since $mc > 0$, the adoption condition (16) pins down a finite limiting ratio $t_{s,\infty}$ and hence finite limiting adoption $r_{s,\infty} = t_{s,\infty}/\phi_s$. Therefore $\ell_{s,n} \rightarrow \ell_{s,\infty} < \infty$, and real wages converge to a finite competitive-AI ceiling. The wage floor $\Gamma(r_{s,\infty}) \geq 1 - f_s > 0$ is preserved.

Case B: $F_n \rightarrow 0$ and additionally $mc_n \rightarrow 0$. Then $p_{AI,n}^* \rightarrow 0$. For any $f_s < 1$, the high-adoption limit of condition (16) gives $r_s \approx f_s \phi_s w_s / [(1 - f_s)p_{AI}]$, so $r_{s,n} \rightarrow \infty$ as $p_{AI,n} \rightarrow 0$. Therefore $\ell_{s,n} \rightarrow \infty$, GDP grows without bound, and $\tilde{w}_{s,n} \rightarrow \infty$. Meanwhile $\Gamma(r_{s,n}) \rightarrow (1 - f_{s,n}) > 0$, so the wage floor is maintained throughout.

Part (iv): Follows directly from Proposition 2 at each n . $\square$ $\square$

Proof of Price-Regulation Remark

Proof. Part (i): Lower $\bar{p}_{AI}$ raises adoption r_s in every sector, raising ℓ_s , Y_s , and GDP (via Proposition 1(i)). AI profits fall since the markup per unit shrinks. The CES price index P falls. Worker welfare $V^W = E^W/P$ rises in the calibrated scenario because the CPI decline and the reduction in AI markup rents dominate any increase in real AI resource costs. Formally, worker income $E^W = \alpha \text{GDP}^G - \Pi_{AI} - mcX - NF$ rises when the labor-share gain from higher output exceeds the net increase in markup and resource costs, and the CES price index P falls as output rises.

Part (ii): By continuity of Ω in μ and Part (i), the welfare gradient is negative for all $\mu \geq \bar{\mu}$, making the Ramsey price the constrained optimum.

Part (iii): Under symmetric Cournot with $N > 1$, total AI revenue satisfies $\partial(p_{AI}X)/\partial p_{AI}|_{p_{AI}^*} = X^*(N-1)/N > 0$ at the unconstrained equilibrium (since the Cournot price exceeds the revenue-maximizing price). Hence a cap strictly reduces AI profits. Capitalist welfare V^C depends on both capital income $(1-\alpha)\text{GDP}$ and AI profits $\Pi_{AI}/(1-\lambda)$; the sign of $\partial V^C/\partial \bar{p}_{AI}$ is therefore ambiguous.

Part (iv) follows from Part (i) by continuity in μ . □ □

Proof of Proposition 4 (Barrier Policy)

Proof. Part (i): Lower τ_s reduces the effective price $p_{\text{eff},s} = \tau_s p_{AI}$, relaxing the adoption condition (16) and raising r_s . Total AI quantity $X = \sum_{s'} L_{s'} r_{s'}$ rises while p_{AI} is unchanged, so $\Pi_{AI} = (p_{AI} - mc)X - NF$ rises strictly.

Part (ii): Higher r_s raises ℓ_s and hence Y_s and GDP. Capital income $rK = (1-\alpha)\text{GDP}$ rises.

Part (iii): From Lemma 1, worker nominal income equals $\alpha \text{GDP} - \Pi_{AI}$. Both terms move in the same direction under barrier reduction (both rise), so the net effect is $\alpha \cdot \Delta \text{GDP} - \Delta \Pi_{AI}$, positive when the stated condition holds. The CES price index also falls as output rises, amplifying the real-wage gain independently of the nominal income effect. At the calibration, barrier removal is the dominant source of productivity gain (S3 CPI falls 21.3%), ensuring the condition holds with a large margin. □ □

B Wage Scenarios: Full Sector Table

Table 7 reports nominal and real wage changes for all 22 SOC-2 sectors under the six scenarios defined in Section 7.2.

Table 7: Wage, employment, and welfare changes from no-AI baseline across five scenarios.

Sector	f_s	S1	S2	S3	S4	S5
<i>Panel A: Nominal wage % change from no-AI baseline</i>						
Management	0.227	-2%	-4%	+1%	-2%	-2%
Business & Fin.	0.320	-3%	-5%	+2%	-3%	-2%
Computer & Math	0.513	+3%	-6%	+3%	+5%	-9%
Arch. & Eng.	0.159	-4%	-1%	-1%	-0%	-3%
Life/Phys. Sci.	0.248	+1%	+1%	+1%	+2%	-2%
Community & Social	0.293	-3%	-5%	-1%	-3%	-4%
Legal	0.254	-3%	-5%	+0%	-3%	-3%
Education	0.391	-3%	-6%	+2%	-3%	-4%
Arts & Media	0.330	+1%	-0%	+1%	+2%	-3%
Healthcare Pract.	0.165	-3%	-4%	-2%	-3%	-4%
Healthcare Supp.	0.103	-2%	-4%	-3%	-3%	-5%
Protective Svc.	0.060	-2%	-4%	-4%	-3%	-6%
Food Prep.	0.138	-3%	-4%	-3%	-3%	-5%
Building & Grounds	0.055	-2%	-4%	-4%	-3%	-6%
Personal Care	0.121	-2%	-4%	-3%	-3%	-5%
Sales	0.386	-3%	-6%	+1%	-4%	-5%
Office & Admin.	0.250	-2%	-4%	-1%	-2%	-4%
Farming	0.019	-2%	-3%	-5%	-2%	-7%
Construction	0.044	-2%	-4%	-4%	-3%	-6%
Install./Repair	0.057	-2%	-4%	-4%	-3%	-6%
Production	0.026	-2%	-3%	-5%	-2%	-6%
Transport	0.045	-2%	-4%	-4%	-3%	-6%
Agg. (wtd.)	—	-2%	-4%	-1%	-2%	-5%
<i>Panel B: Real wage % change (nominal / CPI)</i>						
Management	0.227	+8%	+13%	+28%	+10%	+37%
Business & Fin.	0.320	+7%	+11%	+30%	+9%	+36%
Computer & Math	0.513	+14%	+10%	+30%	+18%	+26%
Arch. & Eng.	0.159	+10%	+16%	+26%	+12%	+35%
Life/Phys. Sci.	0.248	+12%	+18%	+28%	+15%	+36%
Community & Social	0.293	+7%	+11%	+26%	+9%	+33%
Legal	0.254	+7%	+12%	+27%	+9%	+35%
Education	0.391	+7%	+10%	+30%	+9%	+33%
Arts & Media	0.330	+11%	+17%	+28%	+14%	+34%
Healthcare Pract.	0.165	+8%	+12%	+25%	+9%	+33%
Healthcare Supp.	0.103	+8%	+13%	+23%	+9%	+31%
Protective Svc.	0.060	+8%	+13%	+22%	+9%	+31%
Food Prep.	0.138	+8%	+12%	+24%	+9%	+32%
Building & Grounds	0.055	+8%	+13%	+22%	+9%	+31%
Personal Care	0.121	+8%	+12%	+23%	+9%	+32%
Sales	0.386	+7%	+10%	+28%	+8%	+32%
Office & Admin.	0.250	+8%	+12%	+26%	+10%	+34%
Farming	0.019	+8%	+13%	+21%	+10%	+29%
Construction	0.044	+8%	+13%	+22%	+9%	+30%
Install./Repair	0.057	+8%	+13%	+22%	+9%	+30%
Production	0.026	+8%	+13%	+21%	+10%	+30%
Transport	0.045	+8%	+13%	+21%	+9%	+30%
Agg. (wtd.)	—	+8%	+12%	+25%	+10%	+32%
CPI change	—	-9%	-15%	-21%	-11%	-28%
<i>Panel C: Employment % change and worker welfare</i>						
Management	0.227	-2%	-3%	-6%	-2%	-7%
Business & Fin.	0.320	-2%	-3%	-6%	-2%	-7%
Computer & Math	0.513	-3%	-2%	-6%	-4%	-6%
Arch. & Eng.	0.159	-2%	-4%	-6%	-3%	-7%
Life/Phys. Sci.	0.248	-3%	-4%	-6%	-3%	-7%
Community & Social	0.293	-2%	-3%	-6%	-2%	-7%
Legal	0.254	-2%	-3%	-6%	-2%	-7%
Education	0.391	-2%	-2%	-6%	-2%	-7%
Arts & Media	0.330	-3%	-4%	-6%	-3%	-7%
Healthcare Pract.	0.165	-2%	-3%	-5%	-2%	-7%
Healthcare Supp.	0.103	-2%	-3%	-5%	-2%	-7%
Protective Svc.	0.060	-2%	-3%	-5%	-2%	-6%
Food Prep.	0.138	-2%	-3%	-5%	-2%	-7%
Building & Grounds	0.055	-2%	-3%	-5%	-2%	-6%
Personal Care	0.121	-2%	-3%	-5%	-2%	-7%
Sales	0.386	-2%	-2%	-6%	-2%	-7%
Office & Admin.	0.250	-2%	-3%	-6%	-2%	-7%
Farming	0.019	-2%	-3%	-5%	-2%	-6%
Construction	0.044	-2%	-3%	-5%	-2%	-6%
Install./Repair	0.057	-2%	-3%	-5%	-2%	-6%
Production	0.026	-2%	-3%	-5%	-2%	-6%
Transport	0.045	-2%	-3%	-5%	-2%	-6%
Agg. employ. (wtd.)	—	-2%	-3%	-5%	-2%	-7%
Agg. worker welfare	—	+6%	+9%	+18%	+7%	+23%
Agg. GDP change	—	+6%	+9%	+19%	+7%	+23%

Notes: Baseline is pre-AI adoption with wages equal to BLS May 2019 medians. S1: current task susceptibility and calibrated informational barriers (R5 effective price). S2: task susceptibility expanded by 50 percent, R5 price. S3: adoption barriers removed, firms face the market AI price directly. S4: competitive AI pricing (price set to marginal cost), barriers unchanged. S5: all three – S2, S3, and S4 – combined. Real wage equals nominal wage divided by the CPI. Employment change uses the separable CRRA Marshallian elasticity $\varepsilon_M = (1 - \gamma)/(\eta + \gamma) = -0.25$ ($\gamma = 2$, $\eta = 2$); employment falls when real wages rise because the income effect dominates the substitution effect. Worker welfare is the CRRA welfare index $\mathcal{W}_s \propto (w_s/P)^{(1+\eta)/(\eta+\gamma)}$ with $\gamma = 2$, $\eta = 2$, exponent = 0.75; it rises whenever w_s/P rises for any $\gamma > 0$. Sector weights are BLS 2019 employment shares.

C Derivation of Marginal Products

Derivation of $\partial \ell_s / \partial L_s$ and $\partial \ell_s / \partial x_s$. Let $A \equiv (1 - f_s)L_s$ and $B \equiv \ell_s^{\text{aug}}$, so $\ell_s = A^{1-f_s} B^{f_s}$. Since L_s enters both A and B :

$$\frac{\partial \ell_s}{\partial L_s} = \underbrace{(1 - f_s) \cdot A^{-f_s} B^{f_s} \cdot (1 - f_s)}_{\text{via } A} + \underbrace{f_s \cdot A^{1-f_s} B^{f_s-1} \cdot \frac{\partial B}{\partial L_s}}_{\text{via } B}$$

Using $\partial B / \partial L_s = (f_s L_s / B)^{\rho-1} \cdot f_s$ (standard CES derivative), and the fact that $A^{-f_s} B^{f_s} = \ell_s / A$ and $A^{1-f_s} B^{f_s} = \ell_s$:

$$\begin{aligned} \text{Via } A: \quad (1 - f_s)^2 \frac{\ell_s}{A} &= (1 - f_s) \frac{\ell_s}{L_s} \\ \text{Via } B: \quad f_s^2 \ell_s \frac{(f_s L_s)^{\rho-1}}{B^\rho} &= f_s \kappa_s^\rho \frac{\ell_s}{L_s} \end{aligned}$$

Summing gives equation (9). Equation (10) follows analogously, using $\partial B / \partial x_s = (\phi_s x_s / B)^{\rho-1} \phi_s$.

D Robustness: Full Sector Table at 95% CI Bounds

Table 8 repeats the sector-level analysis of Appendix B at $\theta \in \{1.219, 2.940\}$ (95% CI lower and upper bounds). Adoption barriers are re-calibrated at each alternative elasticity; all other parameters are held fixed.

Table 8: Robustness: sector-level real wages across 95% CI for σ_{LA} .

Sector	f_s	S1			S2			S3			S4			S5		
		Base	Lo	Hi	Base	Lo	Hi	Base	Lo	Hi	Base	Lo	Hi	Base	Lo	Hi
Panel A: Nominal wage % change																
Management	0.227	-1.8%	+9.4%	-2.5%	-3.0%	+13.7%	-4.0%	+0.7%	+16.5%	-1.4%	-2.0%	+10.4%	-2.8%	-2%	+23%	-5%
Business & Fin.	0.320	-2.2%	+14.9%	-3.0%	-3.8%	+22.4%	-4.8%	+2.1%	+28.7%	-1.2%	-2.5%	+16.5%	-3.5%	-2%	+40%	-7%
Computer & Math	0.513	-2.8%	+41.8%	-6.8%	-7.6%	+67.6%	-13.8%	+2.5%	+56.2%	-3.7%	-2.5%	+46.5%	-7.1%	-9%	+78%	-20%
Arch. & Eng.	0.159	-1.2%	+6.4%	-2.1%	-2.1%	+9.0%	-3.4%	-0.9%	+8.1%	-2.2%	-1.2%	+7.0%	-2.2%	-3%	+11%	-5%
Life/Phys. Sci.	0.248	-0.2%	+17.3%	-2.4%	-1.0%	+26.4%	-4.4%	+0.7%	+18.8%	-1.7%	+0.3%	+19.0%	-2.2%	-2%	+26%	-6%
Community & Social	0.293	-2.0%	+15.3%	-3.2%	-3.5%	+23.1%	-5.3%	-0.8%	+22.0%	-3.8%	-2.1%	+16.9%	-3.6%	-4%	+30%	-9%
Legal	0.254	-2.0%	+9.2%	-2.3%	-3.2%	+13.4%	-3.6%	+0.1%	+18.9%	-2.3%	-2.2%	+10.2%	-2.7%	-3%	+26%	-6%
Education	0.391	-2.5%	+23.5%	-4.2%	-4.7%	+36.4%	-7.3%	+1.9%	+37.4%	-2.4%	-2.6%	+26.0%	-4.7%	-4%	+52%	-11%
Arts & Media	0.330	-0.7%	+25.1%	-3.7%	-2.3%	+39.4%	-7.0%	+1.0%	+28.6%	-2.5%	-0.1%	+27.7%	-3.5%	-3%	+40%	-9%
Healthcare Pract.	0.165	-1.7%	+3.9%	-1.8%	-2.7%	+5.2%	-2.8%	-1.9%	+7.6%	-3.3%	-1.9%	+4.3%	-2.1%	-4%	+10%	-6%
Healthcare Supp.	0.103	-1.5%	-1.4%	-1.1%	-2.3%	-2.8%	-1.7%	-3.4%	+0.4%	-4.1%	-1.8%	-1.5%	-1.4%	-5%	-0%	-6%
Protective Svc.	0.060	-1.4%	-2.6%	-1.1%	-2.3%	-4.6%	-1.8%	-4.0%	-3.8%	-4.2%	-1.7%	-2.9%	-1.4%	-6%	-6%	-6%
Food Prep.	0.138	-1.5%	-0.4%	-1.1%	-2.4%	-1.3%	-1.7%	-2.6%	+4.4%	-3.7%	-1.8%	-0.4%	-1.4%	-5%	+5%	-6%
Building & Grounds	0.055	-1.4%	-3.6%	-1.0%	-2.2%	-6.0%	-1.6%	-4.1%	-4.4%	-4.3%	-1.7%	-3.9%	-1.2%	-6%	-7%	-6%
Personal Care	0.121	-1.6%	+1.1%	-1.5%	-2.5%	+0.9%	-2.4%	-3.0%	+2.4%	-3.9%	-1.8%	+1.2%	-1.8%	-5%	+2%	-6%
Sales	0.386	-2.2%	+13.3%	-2.0%	-3.5%	+19.8%	-3.2%	+0.9%	+35.4%	-3.3%	-2.6%	+14.9%	-2.5%	-5%	+49%	-11%
Office & Admin.	0.250	-2.0%	+8.6%	-2.1%	-3.2%	+12.4%	-3.4%	-0.7%	+17.4%	-3.2%	-2.2%	+9.5%	-2.6%	-4%	+24%	-7%
Farming	0.019	-0.9%	-4.0%	-0.7%	-1.6%	-6.8%	-1.3%	-5.0%	-8.1%	-4.8%	-1.2%	-4.6%	-0.9%	-7%	-12%	-6%
Construction	0.044	-1.4%	-3.7%	-1.0%	-2.2%	-6.2%	-1.6%	-4.3%	-5.4%	-4.4%	-1.6%	-4.0%	-1.3%	-6%	-8%	-6%
Install./Repair	0.057	-1.4%	-2.9%	-1.1%	-2.2%	-5.0%	-1.8%	-4.3%	-4.4%	-4.5%	-1.7%	-3.2%	-1.3%	-6%	-7%	-6%
Production	0.026	-1.3%	-4.3%	-1.0%	-2.1%	-7.0%	-1.6%	-4.6%	-7.2%	-4.4%	-1.6%	-4.7%	-1.2%	-6%	-11%	-6%
Transport	0.045	-1.4%	-4.2%	-0.9%	-2.1%	-6.8%	-1.5%	-4.5%	-5.6%	-4.5%	-1.6%	-4.5%	-1.1%	-6%	-8%	-6%
Agg. (wtd.)	—	-1.7%	+6.5%	-2.0%	-3.0%	+9.6%	-3.3%	-1.5%	+12.8%	-3.4%	-2.0%	+7.2%	-2.3%	-5%	+17%	-8%
Panel B: Real wage % change																
Management	0.227	+3.8%	+42.8%	+1.2%	+6.0%	+74.5%	+1.9%	+28.0%	+83.3%	+23.1%	+4.9%	+47.6%	+1.7%	+37%	+141%	+28%
Business & Fin.	0.320	+3.4%	+49.9%	+0.7%	+5.2%	+87.9%	+1.1%	+29.7%	+102.5%	+23.3%	+4.4%	+55.8%	+1.0%	+36%	+175%	+25%
Computer & Math	0.513	+2.8%	+85.0%	-3.3%	+1.1%	+157.2%	-8.5%	+30.3%	+145.6%	+20.1%	+4.4%	+95.9%	-2.8%	+26%	+249%	+8%
Arch. & Eng.	0.159	+4.5%	+38.8%	+1.6%	+7.0%	+67.3%	+2.6%	+25.9%	+70.0%	+22.1%	+5.7%	+43.1%	+2.3%	+35%	+117%	+28%
Life/Phys. Sci.	0.248	+5.6%	+53.1%	+1.3%	+8.3%	+94.0%	+1.5%	+28.0%	+86.9%	+22.7%	+7.3%	+59.1%	+2.3%	+36%	+147%	+27%
Community & Social	0.293	+3.6%	+50.4%	+0.5%	+5.5%	+88.9%	+0.6%	+26.1%	+91.9%	+20.0%	+4.8%	+56.3%	+0.9%	+33%	+155%	+23%
Legal	0.254	+3.7%	+42.5%	+1.5%	+5.8%	+74.0%	+2.4%	+27.3%	+86.9%	+21.9%	+4.6%	+47.4%	+1.9%	+35%	+147%	+26%
Education	0.391	+3.1%	+61.1%	-0.6%	+4.2%	+109.4%	-1.5%	+29.5%	+116.1%	+21.8%	+4.3%	+68.5%	-0.3%	+33%	+199%	+20%
Arts & Media	0.330	+5.0%	+63.3%	-0.0%	+6.9%	+113.9%	-1.2%	+28.4%	+102.2%	+21.7%	+6.8%	+70.7%	+1.0%	+34%	+174%	+23%
Healthcare Pract.	0.165	+4.0%	+35.6%	+1.9%	+6.4%	+61.4%	+3.2%	+24.6%	+69.2%	+20.6%	+4.9%	+39.5%	+2.4%	+33%	+115%	+26%
Healthcare Supp.	0.103	+4.2%	+28.6%	+2.7%	+6.8%	+49.1%	+4.4%	+22.8%	+57.9%	+19.6%	+5.1%	+31.7%	+3.2%	+31%	+96%	+26%
Protective Svc.	0.060	+4.3%	+27.0%	+2.6%	+6.9%	+46.4%	+4.3%	+22.1%	+51.4%	+19.6%	+5.2%	+29.9%	+3.2%	+31%	+84%	+26%
Food Prep.	0.138	+4.2%	+29.9%	+2.7%	+6.8%	+51.5%	+4.4%	+23.8%	+64.1%	+20.2%	+5.1%	+33.2%	+3.2%	+32%	+107%	+26%
Building & Grounds	0.055	+4.3%	+25.8%	+2.8%	+7.0%	+44.3%	+4.5%	+21.9%	+50.4%	+19.5%	+5.2%	+28.6%	+3.4%	+31%	+83%	+26%
Personal Care	0.121	+4.1%	+31.9%	+2.2%	+6.6%	+54.9%	+3.7%	+23.3%	+61.1%	+19.9%	+5.1%	+35.4%	+2.8%	+32%	+101%	+26%
Sales	0.386	+3.5%	+47.9%	+1.7%	+5.6%	+83.8%	+2.8%	+28.3%	+113.0%	+20.6%	+4.2%	+53.6%	+2.0%	+32%	+193%	+19%
Office & Admin.	0.250	+3.7%	+41.7%	+1.6%	+5.9%	+72.5%	+2.6%	+26.2%	+84.7%	+20.8%	+4.6%	+46.4%	+2.0%	+34%	+143%	+25%
Farming	0.019	+4.8%	+25.2%	+3.1%	+7.6%	+43.1%	+4.9%	+20.8%	+44.6%	+18.8%	+5.8%	+27.6%	+3.8%	+29%	+73%	+26%
Construction	0.044	+4.3%	+25.6%	+2.7%	+7.0%	+44.0%	+4.5%	+21.6%	+48.8%	+19.3%	+5.2%	+28.3%	+3.3%	+30%	+80%	+26%
Install./Repair	0.057	+4.3%	+26.7%	+2.7%	+6.9%	+45.8%	+4.3%	+21.6%	+50.3%	+19.1%	+5.2%	+29.5%	+3.2%	+30%	+83%	+26%
Production	0.026	+4.4%	+24.8%	+2.7%	+7.1%	+42.7%	+4.5%	+21.3%	+46.0%	+19.2%	+5.3%	+27.5%	+3.4%	+30%	+75%	+26%
Transport	0.045	+4.3%	+25.0%	+2.8%	+7.0%	+43.0%	+4.6%	+21.4%	+48.6%	+19.1%	+5.3%	+27.7%	+3.5%	+30%	+80%	+26%
Agg. (wtd.)	—	+3.9%	+38.9%	+1.7%	+6.1%	+68.1%	+2.7%	+25.2%	+77.5%	+20.6%	+4.9%	+43.4%	+2.2%	+32%	+130%	+24%
CPI change	—	-5.5%	-23.3%	-3.7%	-8.6%	-34.8%	-5.9%	-21.3%	-36.4%	-19.9%	-6.5%	-25.2%	-4.4%	-28%	-49%	-26%

Notes: Each column set uses the indicated θ with τ_s re-calibrated via FOC inversion (AEI R3–R5 convos/worker), so the effective AI price $p_{\text{eff},s}$ reflects the elasticity assumption throughout. The baseline $\sigma_{LA} = 2.079$ is the OLS pre-trend estimate (BLS 2015–2025, AEI R3); $\sigma_{LA} = 1.219$ and 2.94 are the 95% CI bounds (Section 6.3). Scenario definitions: S1 current f_s , R5 τ ; S2 $f_s \times 1.5$; S3 $\tau_s \rightarrow 1$; S4 $p_{AI} \rightarrow mc$; S5 all combined. Implied Lerner indices: 0.241 (baseline), 0.410 (CI lower), 0.170 (CI upper). Real wage = nominal/CPI; ω_s = BLS 2019 employment shares.

E Physical Capacity Ceilings by Sector

Table 9 reports the physical ceiling on AI adoption for all 22 SOC-2 sectors. The ceiling $r_{\text{max},s} = f_s \cdot (8 \times 60/d_s) \times 260$ is the maximum conversations per worker per year consistent with an 8-hour workday and the sector-average conversation duration d_s . $T_{\text{max},s} = \phi_s \cdot r_{\text{max},s}$ converts this to the worker-year units used in the model. The final column shows the ratio of the physical ceiling to observed AEI adoption at R5, quantifying how much adoption could increase before physical

constraints rather than organizational barriers become binding.

Table 9: Physical capacity ceilings and observed AI adoption by sector.

Sector	f_s	d_s (min)	$r_{\max,s}$ (conv/yr)	$T_{\max,s}$ (wkr-yr)	r_{obs} (conv/yr)	Ratio
Management	0.227	15.1	1,876	3.19	15.2	123×
Business & Fin.	0.320	14.4	2,770	4.36	20.3	137×
Computer & Math	0.513	15.9	4,032	6.97	185.6	22×
Arch. & Eng.	0.159	14.7	1,350	2.18	26.3	51×
Life/Phys. Sci.	0.248	15.8	1,958	3.15	189.9	10×
Community & Social	0.293	18.0	2,031	2.23	62.3	33×
Legal	0.254	15.7	2,011	2.74	15.1	133×
Education	0.391	15.2	3,212	4.63	58.7	55×
Arts & Media	0.330	16.3	2,531	3.56	184.2	14×
Healthcare Pract.	0.165	12.8	1,605	1.44	14.1	114×
Healthcare Supp.	0.103	11.2	1,150	0.72	1.9	600×
Protective Svc.	0.060	7.5	998	0.53	4.2	240×
Food Prep.	0.138	8.5	2,016	1.11	2.7	742×
Building & Grounds	0.055	7.5	906	0.49	2.2	413×
Personal Care	0.121	8.5	1,780	0.91	14.0	127×
Sales	0.386	7.9	6,085	3.86	15.2	400×
Office & Admin.	0.250	11.5	2,719	2.17	19.6	139×
Farming	0.019	8.9	270	0.13	62.2	4×
Construction	0.044	12.3	449	0.40	0.9	513×
Install./Repair	0.057	13.3	530	0.35	3.6	147×
Production	0.026	13.6	236	0.34	2.2	109×
Transport	0.045	9.5	587	0.31	0.6	907×

f_s = task susceptibility (fraction of O*NET tasks with AEI coverage). d_s = conversation-weighted average AI-assisted duration (AEI R5, minutes). $r_{\max,s} = f_s \cdot (8 \times 60 / d_s) \cdot 260$: physical ceiling assuming 8 working hours per day, 260 days per year, and f_s share of augmentable tasks. $T_{\max,s} = \phi_s \cdot r_{\max,s}$: worker-year equivalent used as the adoption cap in Scenario S3. r_{obs} : AEI-scaled observed conversations per worker at R5 (February 2026). Ratio = $r_{\max,s} / r_{\text{obs}}$: factor by which physical capacity exceeds current adoption.