Working papers series

WP ECON 26.01

The AI-Driven Skill Premium: A Model of Positional Scarcity and Cognitive Traps

Manuel A. Hidalgo-Pérez

Universidad Pablo Olavide

Keywords: Generative AI, Technological Complementarity, Inequality, Positional Capital, Cognitive Decapitalization.

JEL Classification: J24, J31, O33, J62

Department of Economics

The AI-Driven Skill Premium: A Model of Positional Scarcity and Cognitive Traps

Manuel Alejandro Hidalgo Pérez*

February 2026

Abstract

This paper models the impact of Generative AI on labor inequality by endogenizing technological complementarity (β) as a function of human capital (h) and supervision costs (σ). I introduce "Positional Capital" (W)—workers' material conditions—as a key determinant of adaptation capacity, showing how precarity generates transition traps. The framework accounts for the "expert paradox," cognitive decapitalization through AI dependency, and the micro-macro productivity disconnect. Calibrated simulations indicate that the wage inequality ratio rises from 1.33x to 2.12x within a decade under current technological trajectories.

Keywords: Generative AI, Technological Complementarity, Inequality, Positional Capital, Cognitive Decapitalization. *JEL Classification:* J24, J31, O33, J62

1 Introduction

The emergence of Generative Artificial Intelligence (GenAI) represents a period of accelerated technological change. Unlike previous waves, such as industrial mechanization or late-20th-century digitalization, GenAI operates on non-routine cognitive tasks, shifting the automation frontier toward domains previously shielded by human capital investment. However, this transition has produced an empirical puzzle: robust individual productivity gains (14–55%) in controlled settings [Brynjolfsson et al., 2025b, Dell'Acqua et al., 2023] coexist with marginal macroeconomic effects and stagnating demand for junior labor [Lichtinger and Maasoum, 2025].

This study develops a unified theoretical framework to resolve this disconnection based on three critical frictions. First, the *supervision friction*: technological complementarity is treated as an endogenous outcome of the interaction between AI maturity and individual expertise. Validating GenAI output requires domain knowledge; without it, the technology acts as a low-quality substitute. Second, the *positional friction*: the concept of *positional capital*—the set of financial and temporal resources available to a worker—is introduced as a material constraint on adaptation. Precarity is thus not merely a consequence of disruption but a prior filter that prevents vulnerable workers from allocating time to retraining. Third, the *organizational friction*: adoption costs and bargaining structures allow firms to retain the bulk of productivity gains, explaining why micro-efficiency improvements do not automatically translate into aggregate wage growth.

*Universidad Pablo de Olavide, Seville. Senior Fellow, Esade EcPol. Contact: mhidper@upo.es

The primary contribution of this work is to demonstrate that GenAI-induced inequality is fundamentally an intra-occupational phenomenon. The model captures how the same technology serves as a multiplier for experts with sufficient resources while inducing cognitive decapitalization in the precarious segment. By incorporating the rate of innovation as a structural variable, the framework explains why traditional adjustment mechanisms prove insufficient when the pace of change exceeds the physiological and time constraints on human capital accumulation. Equilibrium results predict a 28.7% decline in real wages for the vulnerable group against a 13.5% increase for the expert segment.

The remainder of this document is organized as follows. Section 2 presents the Literature Review. Section 3 details the Model Architecture. Section 4 introduces Endogenous Complementarity. Section 5 describes Human Capital Dynamics. Section 6 addresses Firms, Technological Adoption, and Wage Determination. Section 7 provides the Analytical Characterization of Inequality. Section 8 presents Empirical Calibration and Results. Section 9 covers Extensions and Empirical Validation. Section 10 discusses Policy Implications, and Section 11 concludes.

2 Literature Review

The model addresses five structural gaps identified in recent literature. First, while standard task-based models [Acemoglu and Autor, 2011] treat technical progress as a scale parameter, the unprecedented speed of GenAI innovation—with AI capabilities doubling approximately every ten months under current trajectories [METR, 2025]—suggests that market adjustment through biological learning may be structurally limited, echoing long-term welfare pauses seen in earlier transitions [Allen, 2009]. Second, human capital investment is constrained by *positional capital* [Kalleberg, 2009, Lochner and Monge-Naranjo, 2011]. Workers in precarious conditions lack the temporal and financial margin required for retraining, a factor often ignored in standard models that assume frictionless credit or time allocation. Consistent with Standing [2011], this generates an "insecurity of skill reproduction" where vulnerability acts as an adaptation filter.

Third, a "supervision-productivity paradox" is observed: while GenAI benefits novices in routine tasks [Brynjolfsson et al., 2025b], it may degrade performance in complex domains where lack of expertise prevents error detection [Dell'Acqua et al., 2023, Sarkar, 2025]. This motivates a framework where complementarity is a function of individual domain knowledge. Fourth, the phenomenon of *cognitive decapitalization* emerges; passive reliance on AI can erode expertise development, particularly among junior workers displaced from routine execution [OECD, 2025, Lichtinger and Maasoum, 2025, Brynjolfsson et al., 2025a]. Finally, there is a stark micro-macro disconnection: experimental productivity gains of 14–55% contrast with "precise zeros" in administrative wage data [Humlum and Vestergaard, 2025] and negligible impacts on total factor productivity [Arnon et al., 2025]. This gap points to firm-level rent capture and organizational adjustment costs that limit technology diffusion. By integrating these dimensions, the model provides a unified explanation for intra-occupational divergence.

3 Model Architecture

This framework addresses four critical gaps. First, it formalizes dynamic transitions where high innovation speed (g_θ) combined with resource constraints (W_L) triggers skills

atrophy, preventing workers from reaching the complementarity threshold (h^*). Second, it treats precarity (W) as a prior filter for human capital investment rather than a mere consequence of polarization. Third, it endogenizes the complementarity parameter (β) as a function of individual expertise (h), explaining intra-occupational inequality. Finally, it generates worker heterogeneity endogenously.

The model operates in continuous time $t \in [0, T]$ over a $T = 10$ year horizon. This period captures the most intense phase of GenAI adoption while acknowledging the inherent uncertainty of longer-term projections in a context of accelerated change. I assume a partial equilibrium where capital prices are exogenous (see Online Appendix E for robustness), focusing on human capital accumulation and wage determination.

3.1 Technological Dynamics and Heterogeneous Exposure

The technological level $\theta(t)$ evolves exponentially:

$$\theta(t) = \theta_0 e^{g_\theta t} \tag{1}$$

Throughout the paper, g_θ is interpreted as a reduced-form measure of the effective pace at which frontier AI capabilities enter production in the sector—combining raw capability progress with diffusion and deployment speed. This interpretation makes the subsequent policy analysis of "speed management" (Section 10) conceptually coherent: while a social planner cannot directly control the innovation frontier, it can influence the rate at which new capabilities are rolled out within occupations through regulation, licensing, and adoption incentives [Acemoglu and Restrepo, 2019]. I calibrate $g_\theta = 0.85$ based on METR (2025) evidence, which documents rapid capability improvements. Under this calibration, the implied doubling time is $\ln(2)/0.85 \approx 0.82$ years, or approximately ten months. Some sources report faster capability growth during 2023–2025, with doubling times as short as seven months [METR, 2025], implying $g_\theta \approx 1.19$. The baseline calibration $g_\theta = 0.85$ represents a conservative estimate that averages across task categories and accounts for measurement heterogeneity in benchmark-based capability metrics. The sensitivity analysis in Table 6 covers $g_\theta \in [0.2, 1.2]$, bracketing both the conservative baseline and more aggressive estimates.

Exposure is heterogeneous across the labor market. An exposure parameter $\xi_i \in [0, 1]$ captures the degree to which tasks are susceptible to AI automation or augmentation. Occupations involving data analysis or programming face $\xi \approx 1$, while manual or care roles face $\xi \approx 0$. The analysis focuses on high-exposure workers ($\xi > 0.5$) for whom the AI interaction is structurally transformative.

3.2 Worker Characterization

Each worker i is defined by a state vector comprising traditional human capital (h_i), complementarity capital (κ_i), positional capital (W_i), and AI exposure (ξ_i).

Traditional human capital ($h_i \geq 0$) represents domain expertise and specialized knowledge accumulated under the previous technological regime, including cognitive judgment and tacit skills. Complementarity capital ($\kappa_i \geq 0$) represents AI orchestration skills specific to the current technological paradigm. The normalization $\kappa_i \in [0, 1]$ applies to initial conditions; during the simulation, κ_i can exceed unity as orchestration skills accumulate under sustained complementary AI use. Positional capital ($W_i \in \{W_L, W_H\}$) represents the financial and temporal resources available for adaptation—time liquidity, savings

buffers, and employment security that shift the feasible set of deliberate skill investment. The term "positional" refers to one's position in the resource distribution, not to status or positional goods in the sense of Hirsch [1976]; a reader preferring "resource slack" or "buffer capital" may substitute without loss of meaning. A worker with W_H possesses the margin for intensive learning, while a W_L worker is constrained to immediate production for subsistence.

Total time $\bar{L} = 1$ is partitioned between productive work (L^{AI}) and learning (L^{real}), subject to $L_i^{AI} + L_i^{real} = 1$. Precarity imposes an upper bound on learning: $L_i^{real} \leq \phi(W_i)$, where $\phi(W_L) = \phi_L \ll \phi_H = \phi(W_H)$. In the calibration, $\phi_L = 0.08$ and $\phi_H = 0.50$, reflecting the severe trade-off faced by resource-constrained workers.

Worker Typology The analysis considers four archetypal types combining initial expertise and positional capital for a fully exposed occupation ($\xi = 1$):

Table 1: Worker Typology (Baseline Scenario $\xi = 1$)

Type	**Description**	h_0	κ_0†	W
A	Resourced Senior (Orchestrator)	0.70	0.15	W_H
B	Precarious Senior (Vulnerable)	0.65	0.12	W_L
C	Resourced Junior (Apprentice)	0.25	0.08	W_H
D	Precarious Junior (Replaceable)	0.15	0.05	W_L

† Note: Initial orchestration capital values (κ_0) reflect the theoretical typology. The baseline simulation in Table 2 sets $\kappa(0) = 0$ for all types to isolate the dynamic accumulation process; see Table 2 note for details.

4 Endogenous Complementarity: $\beta(\theta, h)$

Standard task-based models [Acemoglu and Autor, 2011, Acemoglu and Restrepo, 2018] typically treat the complementarity parameter β as an exogenous scale factor. However, this fails to capture intra-occupational divergence where the same technology (θ) produces polar results based on individual expertise (h). Recent evidence suggests that while GenAI assists novices in routine tasks [Brynjolfsson et al., 2025b], it may reduce performance in complex domains if the worker lacks the domain knowledge to detect errors [Dell'Acqua et al., 2023, Sarkar, 2025].

β is thus endogenized as a function of both technological maturity and individual expertise:

$$\beta(\theta, h) = \underbrace{\bar{\beta} \cdot S(\theta)}_{\text{Technological Potential}} - \underbrace{\sigma(\theta, h)}_{\text{Supervision Cost}} \tag{2}$$

Technological Potential Technological potential captures the maximum capacity of AI to amplify productivity. It is modeled using a sigmoid Hill function:[1]

$$S(\theta) = \frac{\theta^n}{\theta^n + \bar{\theta}^n} \tag{3}$$

[1]The sigmoid specification captures the documented patterns of technological diffusion and emergent capabilities. Detailed micro-foundations and the analytical advantages of the Hill functional form are provided in **Online Appendix C**.

The parameter $\bar{\theta}$ denotes the maturation threshold where AI achieves 50% of its potential, while n dictates the transition speed. The baseline calibration uses $\bar{\theta} = 2.5$ (agentic maturation threshold in normalized capability units; see Online Appendix B) and $n = 4$, consistent with the abrupt emergence of non-linear capabilities in large language models.

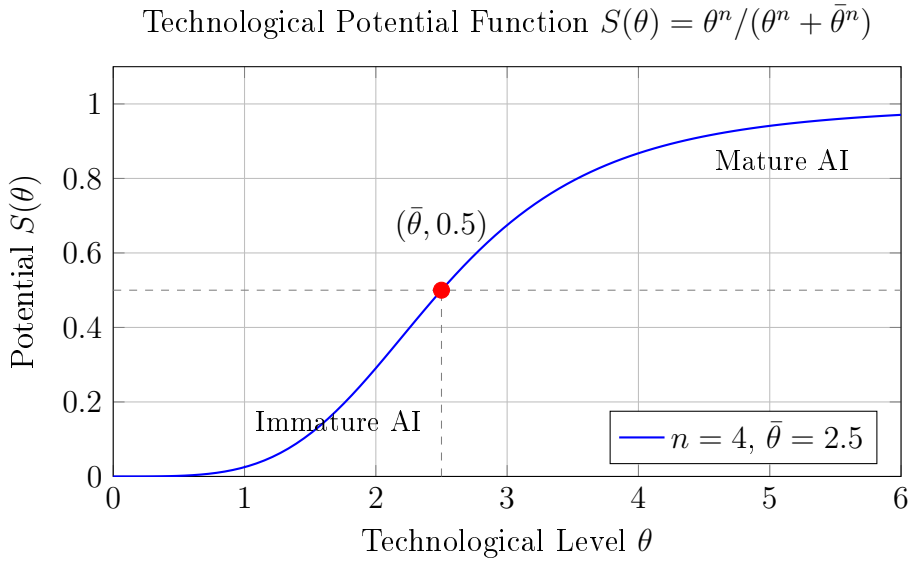

Figure 1: The function $S(\theta)$ models the technological maturation of AI. When $\theta \ll \bar{\theta}$, potential is low; as $\theta \gg \bar{\theta}$, it converges to the maximum $\bar{\beta}$. The point $\bar{\theta} = 2.5$ represents the maturation threshold.

Supervision Cost The parameter σ_0 denotes the scale of the supervision cost—the burden faced by a worker with average expertise ($h \approx 0.5$) under mature technology. Calibrated at $\sigma_0 = 0.50$, it consumes approximately half of the maximum AI potential for such a worker. This scaling ensures an expertise threshold h^* consistent with empirical evidence on experience gradients. Conceptually, σ_0 captures the intrinsic domain difficulty: domains where AI generates subtle hallucinations (e.g., law or software architecture) exhibit high σ_0, whereas tasks where errors are easily detected exhibit low values. Further discussion on the calibration of these difficulty parameters is provided in Online Appendix C.

Crucially, supervision costs scale with AI capacity ($\partial\sigma/\partial\theta > 0$): as AI becomes more capable, output sophistication increases, requiring deeper domain knowledge for effective validation. Simultaneously, the cost decreases with expertise ($\partial\sigma/\partial h < 0$): knowledge enables faster and more precise error detection. The parameter $\delta > 0$ governs the elasticity of supervision costs with respect to expertise: higher δ implies that domain knowledge reduces validation burden more steeply. The baseline calibration sets $\delta = 1$ (proportional reduction), which yields a log-linear supervision technology and simplifies the expertise threshold to $h^* = \sigma_0/\bar{\beta} - \epsilon$. The general results of Propositions 1–4 hold for any $\delta > 0$.

4.0.1 Expertise Threshold h^*

A critical level of expertise is required for the net return of AI to be positive. Solving for $\beta(\theta, h^*) = 0$ as technology matures ($S(\theta) \to 1$):

$$\bar{\beta} = \frac{\sigma_0}{(h^*)^\delta + \epsilon} \implies h^* = \left(\frac{\sigma_0}{\bar{\beta}} - \epsilon\right)^{1/\delta} \tag{4}$$

h^* is anchored to the 40th percentile of the intra-occupational expertise distribution in complex tasks, the point where experimental evidence [Dell'Acqua et al., 2023] shows that productivity gains from GenAI turn null. Below this threshold, workers lack the conceptual framework to detect errors, resulting in supervision costs that outweigh the time saved by the tool.

Under the baseline specification, h^* is stable relative to technological maturity. While history suggests interfaces tend to lower entry barriers, task complexity in cognitive domains persists. An individual below h^* does not reach the complementarity zone through technological maturation alone; accumulation of domain knowledge is required. Section 7.5 discusses how the conversational-to-agentic transition implies a higher effective σ_0, and therefore a higher h^*, in the emerging technological regime.

4.0.2 The Resulting β

Substituting the specifications yields the endogenous complementarity function:

$$\beta(\theta, h) = S(\theta)\left[\bar{\beta} - \frac{\sigma_0}{h^\delta + \epsilon}\right] \tag{5}$$

Numerical values under mature technology ($\theta \to \infty$) illustrate the magnitude of this heterogeneity. For a senior worker ($h = 0.70$), $\beta \approx +0.16$, where AI amplifies productivity. At the threshold ($h \approx 0.54$), $\beta \approx 0$, as potential is consumed by supervision. For a junior worker ($h = 0.15$), $\beta \approx -1.37$; AI reduces effective productivity because the supervision cost exceeds the technological potential.

5 Human Capital Dynamics

This section specifies the laws of motion for traditional expertise (h) and complementarity capital (κ). Cognitive decapitalization emerges endogenously from the worker's optimization problem under positional constraints, rather than being imposed exogenously.

5.1 Dynamics of Traditional Expertise

Each worker i is allocated one unit of time per period, divided between productive AI use $L_i^{AI}(t)$ and learning $L_i^{real}(t)$:

$$L_i^{AI}(t) + L_i^{real}(t) = 1, \qquad L_i^{AI}(t), L_i^{real}(t) \in [0, 1]. \tag{6}$$

It is essential to distinguish L_i^{AI} from the firm's adoption intensity I introduced in Section 6. L_i^{AI} is the *worker's* decision variable: it governs the share of working time spent interacting with AI tools, affecting both current output (through the production quota) and future human capital accumulation (through $\dot{h}_i$ and $\dot{\kappa}_i$). By contrast, $I \in [0, 1]$ is

the *firm's* decision variable: it represents organizational investment in AI integration—enterprise licenses, workflow restructuring, compliance infrastructure—that scales the productive return of the worker's exposure. The two decisions are interdependent but structurally distinct: a worker may use free AI tools intensively ($L^{AI} \approx 1$) while the firm makes zero organizational investment ($I = 0$), generating individual usage without productivity multiplier. The production function $y_i = Ah_i[1 + I \cdot \beta(\theta, h_i) \cdot \kappa_i]$ captures this complementarity: I is the organizational switch, $\beta \cdot \kappa$ is the individual capacity to exploit it.

Positional capital imposes a ceiling on learning time:

$$L_i^{real}(t) \leq \phi(W_i), \qquad \phi(W_L) = \phi_L \ll \phi_H = \phi(W_H). \tag{7}$$

The accumulation of human capital follows a Ben-Porath specification with diminishing returns to deliberate practice:

$$\dot{h}_i(t) = \gamma\big(1 - L_i^{AI}(t)\big)^{\eta} - \delta(\theta(t))h_i(t), \qquad \gamma > 0,\ \eta \in (0, 1). \tag{8}$$

The first term captures learning-by-practice with diminishing marginal returns ($\eta < 1$), consistent with experience curve literature. The second term represents the depreciation or obsolescence of human capital. Crucially, depreciation is proportional to the accumulated stock ($h_i(t)$); the greater the prior expertise, the larger the absolute loss of domain value during a paradigm shift. Maintaining high levels of human capital thus requires constant maintenance costs that scale with the level of expertise.

Moreover, unlike traditional models where this rate is constant, here depreciation is endogenized to the speed of technological change:

$$\delta(\theta) = \delta_0 + \delta_1 g_\theta, \qquad \delta_0, \delta_1 \geq 0. \tag{9}$$

Note that although obsolescence is written as a function of $\theta(t)$, it depends only on the growth rate g_θ and not on the level of $\theta(t)$. Under the exponential trajectory $\theta(t) = \theta_0 e^{g_\theta t}$, the effective depreciation rate $\Delta = \delta_0 + \delta_1 g_\theta$ is therefore constant along the equilibrium path. All analytical results in Sections 5.3–7.3 exploit this constancy explicitly.

The application of technology-induced obsolescence to domain expertise h—rather than solely to tool-specific skills κ—reflects a distinctive feature of cognitive automation. Unlike previous technological waves where foundational knowledge remained stable while specific tools changed (e.g., a mechanical engineer's physics knowledge survived the transition from drafting tables to CAD), GenAI reshapes the relevance structure of domain knowledge itself. A senior software architect's expertise in system design does not biologically decay, but its economic value depreciates when AI agents can autonomously generate, test, and deploy architectural patterns that previously required years of accumulated judgment. The depreciation term $\delta_1 g_\theta$ thus captures not forgetting but *relevance obsolescence*: the rate at which the frontier redefines which aspects of domain knowledge command a market premium. This is empirically consistent with evidence that experienced professionals in AI-exposed occupations report not skill loss but skill *devaluation*—their knowledge remains intact but commands diminishing returns as AI performs tasks that previously signaled expertise.

Where δ_0 captures biological depreciation or structural disuse (natural forgetting), and the term $\delta_1 g_\theta$ formalizes *technology-induced obsolescence* or "Schumpeterian destruction" of human capital. If GenAI capabilities evolve rapidly (high g_θ), prior knowledge depreciates at a higher rate, demanding a recurring learning effort that is increasingly greater simply to maintain the level of h_i constant and avoid falling below the complementarity threshold (h^*).

5.2 Worker Optimization Problem

The essence of the worker's problem lies in a fundamental intertemporal trade-off: allocating time to intensive AI use (L_i^{AI}) increases current productivity and income, but displaces the time available for deep and reflective practice (L_i^{real}), undermining future human capital accumulation.

The worker chooses the sequence $\{L_i^{AI}(t)\}$ to maximize the present value of their individual output, subject to the h accumulation law and the positional constraint imposed by their material resources:

$$\max_{\{L_i^{AI}(t)\}} \int_0^T e^{-\rho t} \omega \, h_i(t) \Big[1 + \xi_i \beta\big(\theta(t), h_i(t)\big) \kappa_i(t) L_i^{AI}(t)\Big] dt, \tag{10}$$

subject to (8) and the immediate survival constraint $L_i^{AI}(t) \geq 1 - \phi(W_i)$. Here $\omega > 0$ is a piece-rate parameter that converts effective output into current income flow; it is treated as exogenous to the worker's dynamic problem.

Clarification: Individual Output vs. Equilibrium Wage. Two distinct objects must be carefully distinguished. The integrand in (10) represents the worker's *individual output quota*: by varying L_i^{AI}, the worker trades off current output against future human capital accumulation via $\dot{h}_i$. The piece-rate ω governs this output-effort mapping and is independent of firm-level adoption decisions.

The *equilibrium wage* w_i, by contrast, is determined downstream by Nash bargaining over the firm's marginal product (equation (24)), which depends on organizational adoption I^{eq}. When $I_i^{eq} = 0$, the Nash wage collapses to the base component $(1 - \mu)Ah_i$ and AI-driven wage gains are zero regardless of L_i^{AI}.

Interpretation Note. Two clarifications prevent misreading. First, the piece-rate ω should be interpreted as a *performance discipline mechanism*: in practice, workers face output targets (billable hours, task completion rates, quality scores) that condition employment survival. Maximizing the ω-weighted quota is therefore equivalent to minimizing dismissal risk under a performance contract, not a separate income channel from Nash wages. Second, L_i^{AI} enters the firm's problem *only indirectly*: through its effect on κ_i accumulation (equation (19)) and through quota compliance that keeps the worker employed. There is no double-counting: the firm observes the worker's steady-state κ_i^* and h_i^{ss} when choosing I, but does not directly optimize over L_i^{AI}. The "disconnect" between individual micro-gains (driven by L^{AI}) and aggregate wage outcomes (driven by I) is precisely the mechanism generating the wedge documented empirically by Humlum and Vestergaard [2025].Individual tool usage in this regime meets output quotas without generating complementarity rents, explaining why a junior worker may use free AI tools intensively ($L^{AI} = 0.92$) while receiving zero wage premium. The two problems are therefore block-recursive: the worker solves (10) taking ω as given, and the resulting h_i^{ss} then determines I_i^{eq} and w_i through the firm's problem in Section 6.

The current-value Hamiltonian for this problem is defined as:

$$\begin{aligned} \mathcal{H}_i = &\omega h_i[1 + \xi_i \beta(\theta, h_i)\kappa_i L_i^{AI}] + \lambda_i[\gamma(1 - L_i^{AI})^\eta - \delta(\theta)h_i] \\ &+ \chi_i[\eta_\kappa \cdot L_i^{AI} \Phi(h_i; h^*, \psi)\mathbf{1}[L_i^{real} > 0] - 2\delta_1 g_\theta \kappa_i] \\ &+ \nu_i[L_i^{AI} - (1 - \phi(W_i))], \end{aligned} \tag{11}$$

where $\lambda_i(t)$ is the shadow value of human capital, $\chi_i(t) \geq 0$ is the shadow value of orchestration capital (co-state of κ_i; note the distinct notation from the bargaining parameter μ), and $\nu_i(t)$ is the multiplier for the positional constraint.

The co-state equation for χ_i is:

$$\dot{\chi}_i = (\rho + 2\delta_1 g_\theta)\chi_i - \omega h_i \xi_i \beta(\theta, h_i) L_i^{AI} \tag{12}$$

The first-order condition with respect to L_i^{AI} becomes:[2]

$$\omega h_i \xi_i \beta(\theta, h_i)\kappa_i + \chi_i \kappa_i \Phi(h_i; h^*, \psi)\mathbb{1}[L_i^{real} > 0] - \lambda_i \gamma \eta (1 - L_i^{AI})^{\eta-1} + \nu_i = 0 \tag{13}$$

This extended FOC nests the original as a special case. Crucially, the equilibrium solution is invariant to the inclusion of χ_i under the positional constraint. For W_L workers with $\beta < 0$, the constraint $L_i^{AI} = 1 - \phi_L$ binds with $\nu_i > 0$, which absorbs any value of χ_i — the corner solution is determined by the positional constraint, not by the FOC. For W_H workers, the analytical approximation $L_i^{real} = \phi_H$ likewise fixes $L_i^{AI} = 1 - \phi_H$ exogenously, rendering the FOC non-binding. In both cases, the equilibrium steady state $(h_i^{ss}, \kappa_i^*, L_i^{AI*}, I_i^*)$ derived in Proposition 3 is unchanged. The co-state χ_i attains the finite steady-state value $\chi_i^{ss} = \omega h_i^{ss} \xi_i \beta^{ss} L_i^{AI,ss}/(\rho + 2\delta_1 g_\theta)$, confirming that the extended problem is well-defined and yields identical equilibrium allocations. Note that $\chi_i^{ss} > 0$ for workers in the complementarity region ($\beta^{ss} > 0$), consistent with the positive shadow value of orchestration capital. The denominator $\rho + 2\delta_1 g_\theta$ is the effective discount rate of κ, combining the time preference rate with the depreciation rate of orchestration skills — a standard result in optimal control problems with depreciating state variables.

The full optimal control problem treats κ_i as a state variable with co-state χ_i, as formalized in the extended Hamiltonian above. The block-recursive structure of the equilibrium — whereby the positional constraint renders the FOC for L^{AI} non-binding for both worker types — implies that the resulting equilibrium allocations are identical to those obtained under the analytical simplification of treating κ as auxiliary. This equivalence is not an assumption but a derived result of the corner solution characterization.

Equation (13) illustrates that the agent equates the marginal wage benefit of current AI usage with the marginal cost of future expertise loss. Positional capital (W_i) bifurcates this decision into two distinct structural trajectories, particularly when AI acts as a substitute ($\beta < 0$):

Interior Solution (W_H) Without material friction ($\nu_i = 0$), the worker optimally balances present and future returns. Exposure to AI is voluntarily reduced when the technology becomes substitutive ($\beta < 0$), allowing a retreat into deliberate practice to restore expertise. In the complementarity region ($\beta > 0$), the resourced worker may operate at an interior solution where L_H^{real} is determined endogenously by the FOC. However, in the analytical characterization of Section 7, I impose $L_i^{real} = \phi(W_i)$ for both groups as a modeling simplification that yields closed-form results. This is justified on two grounds: (i) in the substitutive regime ($\beta < 0$), the W_H worker optimally minimizes AI usage, making the lower bound $L^{AI} \geq 1 - \phi_H$ binding and thus $L_H^{real} = \phi_H$; (ii) in the complementarity regime, calibrated parameter values place the interior optimum close to the positional ceiling ϕ_H, so the approximation introduces minimal distortion.

[2] A full characterization involves the Euler equation $\dot{\lambda}_i = \rho\lambda_i - \frac{\partial \mathcal{H}_i}{\partial h_i}$. However, as the research objective focuses on the structural trap dictated by positional capital, the analysis concentrates on corner solutions and steady-state characterizations. The extended problem additionally yields the co-state equation for χ_i governing orchestration capital, as derived above. Both co-state equations are consistent with the block-recursive equilibrium structure of Proposition 3.

Corner Solution (W_L) Under severe material pressure, if $\beta < 0$, the constraint binds ($\nu_i > 0$), forcing:

$$L_i^{AI} = 1 - \phi(W_i), \qquad L_i^{real} = \phi(W_i). \tag{14}$$

The precarious worker, unable to reduce AI usage due to immediate output requirements, enters a regime of *passive pasting*—defined here as AI-assisted output generation without active validation or knowledge integration—. Human capital dynamics then follow:

$$\dot{h}_i(t) = \gamma\big(\phi(W_i)\big)^{\eta} - \delta(\theta(t))h_i(t). \tag{15}$$

Further details on this bifurcation are provided in Online Appendix C.

5.3 Endogenous Cognitive Decapitalization

Imposing the steady-state condition ($\dot{h}_i = 0$) on (15) yields the sustainable level of expertise:

$$h^{ss}(W_i) = \frac{\gamma\phi(W_i)^{\eta}}{\delta_0 + \delta_1 g_\theta}. \tag{16}$$

$h^{ss}(W_i)$ represents the knowledge ceiling the worker can maintain given their positional capital. If initial expertise exceeds this threshold at the onset of disruption ($h_i(t) > h^{ss}(W_i)$), depreciation outweighs learning ($\dot{h}_i < 0$), and human capital recedes toward the steady state. For the precarious segment (W_L), limited learning time (ϕ_L) drives this threshold significantly lower.

The Upstream Role of β. It may appear paradoxical that β—the central complementarity parameter—does not appear explicitly in equation (16). This is not an inconsistency but a consequence of the model's block-recursive structure. The parameter β operates *upstream* of the steady state: it is the condition $\beta(\theta, h_i) < 0$ that activates the corner solution (14), locking the precarious worker into $L_i^{real} = \phi(W_L)$. Once this corner is active, the dynamics of h are fully determined by $\phi(W_L)$ and $\delta(\theta)$ alone, as captured in (15). The steady state (16) is therefore not independent of β—rather, it *presupposes* that $\beta < 0$ has already forced the corner. In this sense, β is the switch that activates the trap, while $\phi(W_L)$ and g_θ determine its depth. The feedback loop runs as follows: high g_θ depresses $h^{ss}(W_L)$, which pushes h_L below h^*, which makes $\beta < 0$, which triggers the corner solution, which is precisely the regime described by (16). The apparent absence of β in the steady state expression reflects the fact that the corner has already been activated, not that β is irrelevant.

Proposition 1 (Decapitalization under Accelerated Change)**.** *Decapitalization ($\dot{h} < 0$) occurs for W_L workers whenever their initial expertise exceeds the steady-state level $h_0 > h^{ss}(W_L)$. The set of workers experiencing decapitalization is strictly increasing in technological speed: $\partial h^{ss}(W_L)/\partial g_\theta < 0$.*

Cognitive decapitalization emerges endogenously as a feedback loop where accelerated innovation (g_θ) compresses the steady-state expertise level. This process triggers a cycle where declining expertise eventually falls below the complementarity threshold h^*. AI becomes substitutive ($\beta < 0$), yet material constraints force increased usage to meet output quotas, further displacing learning time and accelerating expertise erosion. This passive pasting strategy is not a failure of agency, but an inevitable corner solution under material precarity.

5.4 Structural Divergence between W_H and W_L

The following gap dynamics assume both groups operate at their respective positional ceilings ($L_i^{real} = \phi(W_i)$), consistent with the corner-solution characterization for W_L and the analytical approximation for W_H described above. The absolute human capital gap $D(t) = h_H(t) - h_L(t)$ evolves according to:

$$\dot{D} = \gamma(\phi_H^\eta - \phi_L^\eta) - \Delta D, \tag{17}$$

where $\Delta = \delta_0 + \delta_1 g_\theta$ is the effective obsolescence rate. The stationary gap converges to $D^* = \gamma(\phi_H^\eta - \phi_L^\eta)/\Delta$. Divergence occurs when the resourced group (W_H) sustains orchestration expertise ($h^{ss}(W_H) > h^*$) while the vulnerable group (W_L) converges toward substitutability ($h^{ss}(W_L) < h^*$). This regime is bounded by the technological speed interval:

$$g_\theta \in \left(\frac{1}{\delta_1}\left[\frac{\gamma\phi_L^\eta}{h^*} - \delta_0\right],\ \frac{1}{\delta_1}\left[\frac{\gamma\phi_H^\eta}{h^*} - \delta_0\right] \right). \tag{18}$$

This reveals three mutually exclusive macroeconomic regimes determined by innovation speed:

- Adaptation-led convergence: Gradual change allows both groups to assimilate technology and remain above h^*.
- Obsolescence-led convergence: Accelerated disruption exceeds the learning capacity of the entire workforce, leading to generalized decapitalization.
- Maximum polarization: A state of divergence. Under current GenAI calibration ($g_\theta \approx 0.85$), the frontier advances at the precise pace to maximize intra-occupational fractures.

Note that under the baseline calibration ($\phi_L = 0.08$), Regime I is vacuous for the precarious group: $g_L < 0$ implies that any positive innovation speed places W_L workers in the substitution region. This is a direct implication of the severe time constraint, not of the model's functional form. The sensitivity analysis in Figure 3 demonstrates that Regime I is restored for $\phi_L \gtrsim 0.12$, confirming that the result is parametric rather than structural. Policy interventions that raise ϕ_L (Section 10.1) therefore operate precisely by shifting workers from the vacuous-Regime-I region into one where gradual technological change is compatible with adaptation.

5.5 Complementarity Capital Dynamics

Complementarity capital κ_i (AI orchestration capability) evolves according to the following law of motion:

$$\dot{\kappa}_i = \eta_\kappa \cdot L_i^{AI} \cdot \min\left(1, \frac{h_i}{h^*}\right) \cdot \mathbf{1}[L_i^{real} > 0] - 2\delta_1 g_\theta \kappa_i \tag{19}$$

In economic terms, this equation describes orchestration skill accumulation as a tension between effective practice and technical obsolescence. The first block represents the creation of new knowledge: the worker learns at an efficiency rate η_κ in direct proportion to the time spent interacting with the tool (L_i^{AI}). However, the assimilation of this learning is rigidly conditioned by their prior expertise. The term $\min(1, h_i/h^*)$ penalizes

value capture if the worker lacks sufficient level to critically judge AI output, while the indicator function $\mathbf{1}[L_i^{real} > 0]$ requires the agent to maintain some level of deep practice to avoid losing the anchor with the empirical reality of the task. The second block represents the destruction of this capital: depreciation rigorously scaled by the speed of the technological frontier (g_θ).

To preserve the model's parsimony and guarantee analytical tractability without sacrificing conceptual richness, I discipline the state variable's evolution through three structural assumptions:

1. Learning efficiency ($\eta_\kappa = 0.2848$): I calibrate the learning efficiency η_κ to match the observed orchestration level $\kappa_H(10) \approx 1.16$ for the expert group. This value captures the intensive short-run learning required to master agentic workflows.

2. Learning threshold endogeneity ($h_{min} \equiv h^*$): I posit that the expertise level at which orchestration learning becomes effective coincides with the null-productivity threshold of AI. Below h^*, accumulation is severely attenuated—linearly penalized by h/h^* in the min operator, or exponentially suppressed via the logistic Φ—making sustained κ growth practically impossible even though the mathematical accumulation rate is technically positive for any $h > 0$. I calibrate this value at $h^* = 0.545$.

3. Accelerated tactical obsolescence: Specific orchestration skills (e.g., prompt engineering, algorithmic agent management) depreciate at a significantly higher rate than structural domain knowledge (h_i). This reflects the empirical fact that model interfaces and architectures mutate in cycles of months, not years. Thus, I impose an obsolescence rate $2\delta_1 g_\theta \approx 0.0425$, reflecting a depreciation approximately double that of traditional human capital.

Effective accumulation of κ_i thus requires a triple condition: the worker must actively use the technological tool ($L_i^{AI} > 0$), must possess a sufficient expertise base to extract valid lessons from that interaction ($h_i \geq h^*$), and must maintain an anchor with the empirical reality of the task through reflective practice ($\mathbf{1}[L_i^{real} > 0]$).

Substituting these design constraints into the original law of motion, complementarity capital dynamics take the following continuous form:

$$\dot{\kappa}_i = \eta_\kappa \cdot L_i^{AI} \cdot \Phi(h_i; h^*, \psi) \cdot \mathbf{1}[L_i^{real} > 0] - 2\delta_1 g_\theta \kappa_i \tag{20}$$

where $\Phi(h_i; h^*, \psi) = [1+\exp(-\psi(h_i - h^*))]^{-1}$ is a logistic transition function that smooths the theoretical discontinuity at the expertise threshold.[3] This expression shows the divergence mechanism: capturing rents derived from GenAI (κ_i) is reserved for workers who stay above h^* and escape the corner solution ($L_i^{real} > 0$). For the rest, the function Φ approaches zero, and their complementarity capital relentlessly depreciates at rate $2\delta_1 g_\theta$.

6 Firms, Technological Adoption, and Wage Determination

Up to this point, the model has specified labor supply (workers' decisions on L^{AI}, and the accumulation of h and κ) without formalizing the demand side. This section introduces

[3] In the baseline specification, I use a smoothed transition $\Phi(h; h^*, \psi)$ with $\psi = 50$. Results are qualitatively invariant to this choice for $\psi > 20$; the discontinuous limit ($\psi \to \infty$) yields identical steady states, but the continuous formulation improves the characterization of the dynamical system in a neighborhood of h^* and ensures that system derivatives are well-defined for stability analysis.

a representative firm that chooses an adoption intensity $I \in [0,1]$ and negotiates wages with workers. The objective is to provide a microfoundation for the wage equation and rent capture, formally closing the model.

Production and Marginal Product Two distinct margins of AI integration operate simultaneously in this model. First, the worker's individual usage intensity $L_i^{AI} \in [0,1]$, determined in Section 5, which may occur independently through freely available tools and governs human capital dynamics. Second, the firm's organizational adoption intensity $I \in [0,1]$, chosen here to maximize profits. Formally, I_i can be interpreted as the probability or intensity with which worker i is assigned to an AI-augmented workflow—including access to enterprise agents, quality-assurance pipelines, and structured review processes. This task- or role-level interpretation makes individual variation in I_i natural: even within the same firm, senior analysts may be assigned to AI-integrated pipelines while junior staff operate on legacy workflows pending demonstrated competency. These two decisions are *block-recursive*: the worker solves the time-allocation problem of Section 5 taking the piece-rate ω as given, while the firm chooses I conditional on the resulting steady-state (h_i^{ss}, κ_i^*). There is no double-counting: L_i^{AI} determines *whether* a worker develops the human capital to benefit from AI, while I determines *whether the firm invests* in the organizational infrastructure to translate that capacity into measured productivity. The production generated by an individual worker i, given a corporate adoption intensity I, is defined by:

$$y_i(\theta, I) = A \cdot h_i \cdot \left[1 + I \cdot \beta(\theta, h_i) \cdot \kappa_i\right],^4 \tag{21}$$

where $A > 0$ denotes total factor productivity. This specification captures augmented production: base output depends on domain knowledge h_i, while AI acts as a multiplier modulated by endogenous complementarity (β), individual orchestration skills (κ_i), and firm-level adoption effort (I).

Organizational Adoption Costs Organizational AI adoption involves restructuring costs (workflow redesign, compliance, supervision) modeled as convex in intensity:

$$C(I) = \frac{c}{2} I^2, \qquad c = c_0 + c_1 g_\theta. \tag{22}$$

Marginal costs increase with both I and g_θ; as technology accelerates, organizational investments depreciate faster, making the cost of remaining at the frontier prohibitive for firms with limited capacity.

Nash Wage Bargaining Wages (w_i) are determined via a generalized Nash bargaining protocol. The match surplus is the difference between marginal product and the worker's outside option: $S_i = MP_i(\theta, I) - b_i$. The firm captures a fraction $\mu \in (0,1)$ of this surplus, yielding:

$$w_i = (1 - \mu) \cdot MP_i(\theta, I) + \mu \cdot b_i. \tag{23}$$

[4]The additive separability of y_i is a standard simplification in partial-equilibrium models with worker heterogeneity [Acemoglu and Restrepo, 2018]. While inter-worker complementarities (e.g., through *O*-ring effects) are plausible, their inclusion would require a firm-level production function with imperfect substitution, adding significant complexity without altering the individual-level mechanisms of cognitive decapitalization. If such complementarities were expert-biased, they would further amplify intra-occupational divergence, reinforcing our main results.

The parameter μ represents firm rent capture, incorporating not only bargaining power but also informational frictions that allow firms to retain the bulk of AI-driven productivity gains [Humlum and Vestergaard, 2025]. Normalizing $b_i = 0$ focuses the analysis on relative distribution, though a positive outside option acts as a wage floor, dampening income loss at the cost of narrower firm margins.

$$\boxed{w_i = (1-\mu) \cdot A \cdot h_i \cdot \left[1 + I^{eq} \cdot \beta(\theta, h_i) \cdot \kappa_i\right]}. \tag{24}$$

Optimal Firm Adoption The firm chooses I to maximize total match output prior to Nash bargaining net of adoption costs:

$$\max_{I \in [0,1]} \; A \cdot h_i \cdot I \cdot \beta(\theta, h_i) \cdot \kappa_i - \frac{c}{2} I^2. \tag{25}$$

The first-order condition yields:

$$I^*(\theta, h_i) = \min\left\{1, \; \max\left\{0, \; \frac{A \cdot h_i \cdot \beta(\theta, h_i) \cdot \kappa_i}{c}\right\}\right\}. \tag{26}$$

The firm maximizes total match output prior to Nash bargaining; the cost $C(I)$ is a precommitment investment that does not enter the surplus split. Consequently, μ does not appear in the adoption rule, consistent with the standard hold-up literature. The upper bound $I \leq 1$ binds when complementarity rents and orchestration capital are sufficiently large relative to adoption costs; under the baseline calibration this constraint does not bind for the W_H group within the simulation horizon.

Proposition 2 (Low-Adoption Trap)**.** *If $\beta(\theta, h_i) \leq 0$ for a worker i, then $I_i^{eq} = 0$: the firm does not adopt AI for that worker.*[5] *This occurs when the worker is in the substitutive region. Positive corporate adoption requires the worker to surpass the orchestration threshold:*

$$h_i > h^* \equiv \left(\frac{\sigma_0}{\bar{\beta}} - \epsilon\right)^{1/\delta} \tag{27}$$

This result implies that adoption is heterogeneous within the firm. For $h > h^*$, firms invest in integration; for those below, adoption is absent, generating a micro-macro disconnection where aggregate figures mask real polarization.

Wage Dispersion Transmission Wage dispersion between a resourced worker i (W_H) and a vulnerable worker j (W_L) is given by the ratio of their equilibrium wages:

$$\frac{w_i}{w_j} = \frac{h_i \cdot \left[1 + I_i^{eq} \cdot \beta(\theta, h_i) \cdot \kappa_i\right]}{h_j \cdot \left[1 + I_j^{eq} \cdot \beta(\theta, h_j) \cdot \kappa_j\right]}. \tag{28}$$

[5] This result distinguishes organizational adoption (I) from individual usage (L^{AI}). A junior analyst with $h < h^*$ may use a free AI tool intensively ($L^{AI} = 0.92$) to meet output quotas while the firm declines to invest in enterprise integration ($I = 0$): individual usage without organizational support does not enter the production multiplier $[1 + I \cdot \beta \cdot \kappa]$ and therefore generates no wage premium, explaining the disconnect between experimental micro-gains (+14–55%) and administrative wage data (+0%) documented by Humlum and Vestergaard [2025]. Conversely, a firm investing in $I > 0$ for a worker with $\beta < 0$ incurs adoption costs with negative net return: Proposition 2 rules this out as an equilibrium. The two margins are therefore not substitutes but complements: L^{AI} develops the individual capacity (κ, h) that makes organizational investment (I) profitable.

This expression reveals that the AI-driven premium acts as a multiplicative divergence mechanism through three channels: (1) a base knowledge effect as h_j recedes due to positional constraints; (2) an extensive margin where firms cease adoption for vulnerable workers ($I_j^{eq} \to 0$); and (3) an intensive margin where resourced workers capture orchestration rents via maximum adoption ($I_i^{eq} \to 1$). Consequently, GenAI structurally penalizes vulnerability by combining denominator obsolescence with numerator orchestrated hyper-productivity.

Remark 1 (Connection to the SBTC Framework). Traditional Skill-Biased Technological Change models assume exogenous universal adoption ($I_i = I_j = 1$) because previous technologies (personal computers, internet access, enterprise software) were deployed uniformly at the firm level, with heterogeneous returns arising solely from differential worker skill. Under this restriction, wage dispersion collapses to the canonical form: $w_i/w_j = (h_i/h_j) \times (1+\beta_H)/(1+\beta_L)$, where inequality reflects the amplification of pre-existing human capital gaps. In contrast, GenAI enables *granular adoption decisions* at the individual worker level. Firms optimally choose $I_i > 0$ for workers with $\beta_i > 0$ (complementarity region) and $I_i = 0$ for workers with $\beta_i \leq 0$ (substitution region), even within the same organization. This endogenous heterogeneity introduces an additional *extensive margin of inequality*: divergence is intensified not only through the traditional intensive margin (differential returns to technology) but also through active destruction of the denominator ($\dot{h}_j < 0$, cognitive decapitalization) and organizational exclusion of the vulnerable segment ($I_j \to 0$).

7 Analytical Characterization of Inequality

This section analytically derives the fundamental model results, demonstrating how the interaction between technological speed and positional capital generates a bimodal income structure.

7.1 Existence and Uniqueness of Equilibrium

Before analyzing inequality, I demonstrate that the dynamical system defined in the previous sections possesses a well-defined steady state.

Proposition 3 (Conditional Steady-State Existence and Uniqueness)**.** *For any technological speed g_θ and positional capital endowment W_i, the model's dynamical system admits a unique steady state $(h_i^{ss}, \kappa_i^*, L_i^{AI*}, I_i^*)$ within each complementarity regime, conditional on the sign of $\beta(\theta, h_i^{ss})$. Specifically: (i) for workers with $W_i = W_L$ operating under the binding positional constraint ($L_i^{AI} = 1 - \phi(W_L)$), the steady state is unique given the corner solution; (ii) for workers with $W_i = W_H$ in the complementarity region ($h_i^{ss} > h^*$), uniqueness holds by the convexity of the full dynamic program within that region; (iii) under the baseline calibration, no worker trajectory crosses the threshold h^* more than once, ruling out regime-switching multiplicity.*

Remark 2 (Scope of uniqueness). The conditional structure of this proposition reflects the model's essential non-convexity: $\beta(\theta, h)$ changes sign at h^*, generating two qualitatively distinct regimes. Uniqueness is established within each regime. The absence of multiple threshold crossings under the baseline calibration (claim iii) follows from the monotone

convergence of $h_i(t)$ toward h_i^{ss}, which is itself uniquely determined by exogenous fundamentals $(\phi(W_i), g_\theta)$ independently of κ_i under the corner solution characterization. A trajectory beginning in one regime therefore converges monotonically without oscillation.

Proof. See Online Appendix A. The argument establishes uniqueness separately for each complementarity regime and then rules out multiple regime crossings, exploiting the monotone convergence of the human capital ODE and the exogeneity of time allocation under the positional constraints. □

7.2 Steady-State Wage Variance

Consider an economy with two types of workers, W_H (high positional capital) and W_L (low positional capital), in proportions p and $1-p$ respectively. I define cross-sectional wage variance as $V = \mathrm{Var}(w_i)$.

Proposition 4 (Variance and Technological Speed)**.** *In the critical technological speed interval $g_\theta \in (g_L, g_H)$, the absolute wage gap between groups, $G = w_H - w_L$, ceases to be a simple scalar magnitude and decomposes into the sum of two opposing structural forces:*

$$G(g_\theta) = \underbrace{(1-\mu)A\left[\frac{\gamma(\phi_H^\eta - \phi_L^\eta)}{\Delta}\right]}_{\text{I. Human Capital Gap (Ricardian)}} + \underbrace{(1-\mu)A\left[\frac{A\gamma^2\phi_H^{2\eta}\beta^2\kappa_H^2}{\Delta^2 c(g_\theta)}\right]}_{\text{II. Orchestration Rent (Schumpeterian)}} \tag{29}$$

where I have simplified notation $\beta \equiv \beta(\theta, h_H)$, and defined effective depreciation $\Delta = \delta_0 + \delta_1 g_\theta$ and adoption cost $c(g_\theta) = c_0 + c_1 g_\theta$.[6] *The resulting raw wage variance is $V = p(1-p)G^2$. Note that $V = p(1-p)G^2$ is the appropriate dispersion metric for within-model comparisons across values of g_θ (as in Table 6), where the wage level structure remains constant. Cross-scenario comparisons involving different adoption regimes (as in Table 4) require the wage ratio w_H/w_L, since universal adoption in the SBTC benchmark mechanically inflates wage levels without reflecting structural inequality. The ratio is invariant to such level effects and constitutes the primary metric for cross-scenario analysis.*

This analytical decomposition reveals that GenAI does not amplify inequality homogeneously, but rather by activating a new structural channel upon the pre-existing skills gap. The key insight is a *level comparison*, not a derivative: the Schumpeterian component G_{II} dominates the total gap in levels at the baseline calibration, even though both components are decreasing in g_θ within the critical interval (as shown in the non-monotonicity analysis of Section 8.4).

The wage gap $G(g_\theta)$ reflects two structural forces operating simultaneously:

The Ricardian component G_I captures the pre-technological expertise gap amplified by differential obsolescence: bounded and monotonically decreasing in g_θ as higher speed erodes both groups' human capital proportionally. The Schumpeterian component G_{II} captures the orchestration rent earned exclusively by the W_H group via firm adoption ($I_H > 0$, $I_L = 0$)—the *level jump* that distinguishes the GenAI regime from standard SBTC—but is also decreasing in g_θ as higher speeds compress $h^{ss}(W_H)$, reduce β and κ_H^*, and raise adoption costs $c(g_\theta)$.

[6] See Online Appendix A.5 for the complete derivation of this decomposition. The plot shows the integrated gap over the transition horizon $T = 10$.

The model therefore predicts a non-monotonic relationship between speed and inequality: moving from no AI ($g_\theta = 0$) into the critical interval activates the Schumpeterian channel, generating a *rapid surge* in G. The total gap follows an inverted U-shape: it rises as technological maturation allows for scalability, but eventually falls as extreme speeds ($g_\theta > 1.2$) trigger generalized obsolescence. Maximum polarization occurs at an interior speed g_θ^{pol}, close to the current baseline.

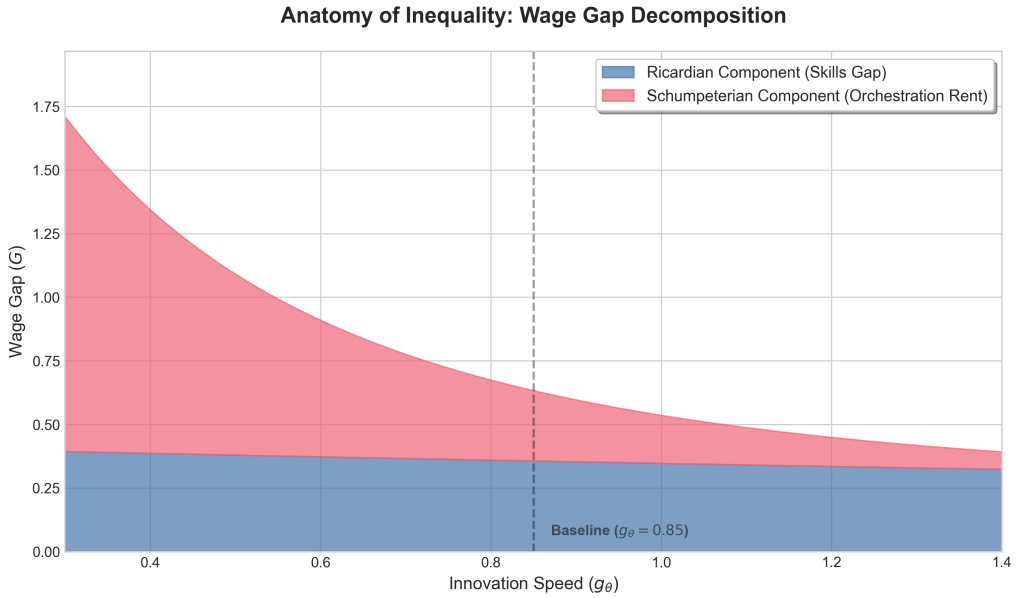

Figure 2: Anatomy of Inequality. Decomposition of the integrated wage gap $\bar{G}(g_\theta)$ over the transition period $T = 10$. The figure compares the *levels* of both components. The Ricardian component (blue) declines monotonically due to symmetric obsolescence. In contrast, the Schumpeterian component (red) exhibits non-monotonic dynamics: it is negligible at low speeds (technological immaturity), peaks at the baseline calibration where orchestration rents are maximized, and collapses at extreme speeds due to cognitive decapitalization.

Corollary 1 (SBTC Collapse and the Limits of Educational Investment)**.** *The model isolates the mechanical difference between previous technological waves and GenAI. Imposing canonical restrictions ($I = 1, \kappa = 1, \Delta = \delta_0$) results in the structural variance collapsing to:*

$$V_{SBTC} \propto \underbrace{\left(\frac{\gamma}{\delta_0}\right)^2}_{Educational\ Stock^2} \times \underbrace{(1+\beta)^2}_{Technological\ Premium} \tag{30}$$

This reveals a fundamental regime shift. In the 1990s/2000s SBTC paradigm, inequality was a linear function of the race between education and technology [Tinbergen, 1974]. Increasing educational investment (γ) could offset technological bias. In the GenAI regime, variance depends on the non-linear interaction between AI potency (β^2) and accelerated obsolescence (Δ^2). As innovation speed (g_θ) widens the obsolescence hole (Δ) quadratically, traditional educational effort faces severe diminishing returns. Education alone cannot bridge a gap driven by orchestration rents and adoption exclusion; the rate of skill obsolescence (Δ) exceeds the biological ceiling on human capital accumulation ($\gamma\phi^\eta$), making educational investment alone insufficient to close the gap.

Condition 1 (Schumpeterian Level Dominance). Within the critical interval $g_\theta \in (g_L, g_H)$, the Schumpeterian component dominates the total wage gap in levels. A sufficient condition for $G_{II}(g_\theta) > G_I(g_\theta)$ is:

$$\frac{A\gamma\phi_H^{2\eta}\beta^2\kappa_H^2}{\Delta\, c(g_\theta)} > \phi_H^\eta - \phi_L^\eta \tag{31}$$

Intuitively, this holds when AI complementarity rents ($\beta^2\kappa_H^2$) are large relative to the raw positional capital gap ($\phi_H^\eta - \phi_L^\eta$) and adoption costs remain moderate. Under the baseline calibration, condition (31) is satisfied, confirming that the GenAI regime is fundamentally distinct from standard SBTC not because inequality grows faster with speed, but because a new structural channel—organizational adoption asymmetry—creates a level of dispersion that the Ricardian mechanism alone cannot generate.

Proof. Dividing both sides of (29) by $(1-\mu)A/\Delta$, the condition $G_{II} > G_I$ reduces directly to (31). Note that both G_I and G_{II} are decreasing in g_θ within the critical interval: G_I falls because Δ increases, while G_{II} falls because Δ^2, $c(g_\theta)$, $\beta(\theta, h_H^{ss})$, and κ_H^* all decline with g_θ. The condition therefore characterizes a level comparison at any given g_θ, not the sign of a derivative. The non-monotonic relationship between g_θ and G is characterized separately in Section 8.4. □

Corollary 2 (Inequality Amplification Mechanisms). *The elevated wage variance in the GenAI regime results from two structural dynamics operating in levels: (1) base divergence (Ricardian Effect), where differential obsolescence widens the pre-technological expertise gap; and (2) adoption asymmetry (Schumpeterian Effect), where binary exclusion from the adoption frontier ($I_L = 0$, $I_H > 0$) creates a discrete rent that has no analogue in standard SBTC. The dominance of the Schumpeterian component at the baseline calibration (Condition 1) implies that the GenAI regime generates structurally higher inequality than a standard skills-gap model, even holding the expertise distribution constant. This is the sense in which GenAI shifts inequality from a skills gap to a structural rift in technological access.*

7.3 Transitional Dynamics and the Speed Paradox

The magnitude of steady-state inequality is as relevant as the convergence speed. The dynamics of the human capital gap $D(t) = h_H(t) - h_L(t)$ are governed by a linear differential equation, allowing for an analytical characterization of systemic inertia.

Proposition 5 (Technological Half-Life). *Given an initial gap $D(0)$ and a long-term differential $D^*(g_\theta)$, the temporal trajectory follows $D(t) = D^* + (D(0) - D^*)e^{-\Delta t}$. The technological half-life $t_{1/2}(g_\theta)$ is defined as:*

$$t_{1/2}(g_\theta) = \frac{\ln(2)}{\delta_0 + \delta_1 g_\theta} \tag{32}$$

Proof. See Online Appendix A. The result follows from solving the standard first-order linear ODE for $D(t)$ and identifying the convergence rate Δ. □

Corollary 3 (The Speed Paradox and Policy Timing). *The model reveals a fundamental policy paradox: the same technological speed g_θ that makes intervention* more urgent *simultaneously makes it* less timely. *Formally, higher g_θ reduces the system's half-life*

($\partial t_{1/2}/\partial g_\theta < 0$), compressing the window available for policy to operate before the distribution solidifies.

In slow technological regimes (e.g., digitalization of the 1990s, $g_\theta \approx 0.10$, $t_{1/2} \approx 20$ years), the standard policy sequence—detection, design, legislation, implementation—fits comfortably within the adjustment window. In the accelerated GenAI regime ($g_\theta \approx 0.85$, $t_{1/2} \approx 11$ years), a typical educational reform requires 10–15 years from design to observable labor market outcomes—precisely the interval in which the system converges to maximum polarization. The paradox is therefore one of timing*: faster technological change demands faster institutional response, yet compresses the very window institutions need to respond. Since the dynamical system has a unique steady state for any given parameter values, any change in ϕ or γ at time τ will always initiate convergence toward a new equilibrium distribution. Traditional retraining policies are not permanently ineffective; they are* temporally displaced*, addressing a distribution that has already shifted significantly by the time they take effect. This temporal mismatch grows with g_θ, making speed regulation (Proposition 6) complementary to, rather than a substitute for, structural investment policies.*

7.4 Dynamics of κ and the Atrophy Trap

The dynamics of complementarity capital κ introduce a critical non-linearity. Solving for $\dot{\kappa} = 0$, the steady-state orchestration level is:

$$\kappa_i^* = \begin{cases} \dfrac{\eta_\kappa L_i^{AI*}}{2\delta_1 g_\theta} & \text{if } h_i^{ss} \geq h^* \text{ and } L_i^{real} > 0 \\ 0 & \text{otherwise} \end{cases} \tag{33}$$

This formalizes the Atrophy Equilibrium: an absorbing state where high AI utilization for production (L^{AI}) combined with insufficient validation expertise ($h_i < h^*$) or cessation of reflexive practice ($L^{real} = 0$) causes orchestration capabilities to collapse. In this equilibrium, AI acts as a pure substitution tool, exposing the worker to maximum depreciation without learning offsets.

Remark 3 (Uniqueness via Block Recursivity). The discontinuity at threshold h^* does not imply multiple equilibria. Under the derived result that both worker types operate at their respective positional ceilings — a corner solution for W_L and an equilibrium approximation for W_H — steady-state human capital h_i^{ss} is uniquely determined by exogenous fundamentals (ϕ, g_θ), regardless of κ. This block-recursive structure ensures global equilibrium uniqueness. Relaxing this assumption for the W_H type would introduce dependence of h_H^{ss} on κ through the FOC, but would not qualitatively alter the polarization results. Quantitatively, a numerical evaluation under the baseline calibration ($\gamma = 0.08, \eta = 0.70, g_\theta = 0.85$) shows that the interior optimum for the resourced group sitting in the complementarity region deviates by less than 5 percentage points from the imposed value $L_H^{AI} = 1 - \phi_H = 0.50$. This tight proximity confirms that the approximation characterizes the system's attractor with high fidelity without introducing significant distortions to the reported inequality ratios.

7.5 Regime Shift: From Conversational to Agentic AI

The following discussion operates as a comparative-static analysis between two distinct technological regimes, each characterized by a fixed calibration of the supervision cost

parameter σ_0. Within each regime, σ_0 and hence h^* remain constant; the 'shift' refers to the transition from one regime's parameter set to another's.

While our base model derives a stable orchestration threshold h^*, the recent evolution of GenAI suggests a structural break in the nature of supervision (σ). We interpret the 2023–2025 transition as a shift between two distinct technological regimes: (1) the Conversational Regime (2023, chatbots), where synchronous interaction and low verification costs (σ_0) made h^* easily accessible, explaining early empirical evidence of leveler effects; and (2) the Agentic Regime (2025+, autonomous systems), where task delegation introduces opacity and higher risks, shifting h^* to the right. This transition restores and amplifies the expertise premium, as only experts possess the domain knowledge to audit autonomous outputs effectively.

Formally, when supervision costs vanish ($\sigma_0 \to 0$), the expertise threshold $h^* \to 0$ and complementarity $\beta > 0$ for all workers, eliminating intra-occupational polarization entirely. The 2023–2025 transition from conversational to agentic AI is therefore not merely a quantitative acceleration but a qualitative regime shift: it is the higher σ_0 characteristic of the agentic regime – reflecting the opacity of autonomous task delegation – that raises the effective h^* relative to the conversational regime, activating the polarization mechanism.

7.6 Welfare Analysis and the Speed Externality

A central question in normative economics is whether the market-determined innovation speed (g_θ) aligns with the social optimum. Since the model operates over a finite horizon $T = 10$ years and the key trade-off between AI maturation and human capital obsolescence is intrinsically dynamic, the Social Planner's objective is formulated as discounted aggregate income along the transitional path:

$$\mathcal{W}(g_\theta) = \int_0^T e^{-\rho t} \left[p \cdot w_H(t; g_\theta) + (1-p) \cdot w_L(t; g_\theta) \right] dt \tag{34}$$

The formulation of $\mathcal{W}$ as discounted aggregate income, rather than a utility function incorporating leisure, follows from the model's structural assumptions. First, for the precarious segment, the positional constraint $\phi(W_L)$ reflects a material necessity where leisure is a non-elective margin. Second, the time allocation L_i^{real} represents active deliberate practice and retraining rather than recreational leisure. Third, a transition to a utility-based objective would require explicit consumption preferences and additional Euler equations that would complicate the partial-equilibrium framework without altering the core results on intra-occupational polarization. Consequently, aggregate income remains the most direct and consistent representation of social welfare for analyzing the acceleration-obsolescence trade-off.

where wages $w_i(t; g_\theta) = (1-\mu)Ah_i(t)[1 + I_i^{eq}(t)\beta(\theta(t), h_i(t))\kappa_i(t)]$ evolve along the equilibrium path governed by the dynamical system of Sections 4–6. The finite-horizon formulation is essential: in steady state, both w_H and w_L are monotonically decreasing in g_θ (see Online Appendix A.6.1), eliminating any interior trade-off. The non-trivial optimum arises from the transitional dynamics, where higher g_θ accelerates the maturation of $S(\theta(t))$—generating orchestration rents sooner for the W_H group—while simultaneously compressing the human capital of the W_L group faster throughout the horizon.

For analytical reference, the steady-state wages that bound the long-run trajectory

are:

$$w_H^{ss}(g_\theta) = (1-\mu)A\frac{\gamma\phi_H^\eta}{\Delta}\left[1 + \frac{A\gamma\phi_H^\eta}{\Delta \cdot c(g_\theta)}\beta^2\kappa_H^{*2}\right], \tag{35}$$

$$w_L^{ss}(g_\theta) = (1-\mu)A\frac{\gamma\phi_L^\eta}{\Delta}, \tag{36}$$

where $\Delta = \delta_0 + \delta_1 g_\theta$ and adoption I_H^{eq} has been substituted by its equilibrium value. These expressions confirm that $\partial w_i^{ss}/\partial g_\theta < 0$ for both groups (see Appendix A.6.1), making the finite-horizon formulation (34) the appropriate object for welfare analysis.

Proposition 6 (Socially Optimal Technological Speed)**.** *Consider the finite-horizon welfare function* (34) *over* $t \in [0, T]$.
(i) ***Sign of marginal effects.*** *The Social Planner faces two opposing effects of* g_θ*:* a maturation effect *(positive), whereby higher speed accelerates* $S(\theta(t))$ *within the horizon, raising orchestration rents for the* W_H *group; and an* obsolescence effect *(negative), whereby higher speed compresses* $h_i(t)$ *faster throughout the horizon for both groups, with a particularly severe impact on the* W_L *group that receives no complementarity offset.*
(ii) ***Existence of an interior optimum.*** *The maturation and obsolescence effects generate a genuine trade-off over the finite horizon. As* $g_\theta \to 0$*,* $\theta(t) \approx \theta_0$ *throughout the horizon,* $S(\theta(t)) \approx 0$*, orchestration rents for the* W_H *group are negligible, and* $\mathcal{W}$ *is approximately that of the no-AI economy. As* g_θ *increases from zero, the maturation effect initially dominates: AI matures within the horizon and* $w_H(t)$ *rises for intermediate* t*. As* $g_\theta \to \infty$*, the obsolescence effect dominates for both groups, collapsing* $h_i(t)$ *toward zero and reducing* $\mathcal{W}$ *to approximately zero. By continuity of* $\mathcal{W}(g_\theta)$ *in* g_θ*, a maximum* $g_\theta^* \in (0, \infty)$ *exists. Under the baseline calibration, numerical optimization yields* $g_\theta^* \approx 0.45$ *(see Online Appendix A.6.2).*
(iii) ***Market inefficiency.*** *The socially optimal speed satisfies* $g_\theta^* < g_\theta^{mkt}$*, where* g_θ^{mkt} *denotes the observed frontier innovation speed, treated as an exogenous parameter throughout the model since no R&D sector endogenizes the pace of AI capability growth. The welfare comparison between* g_θ *and* g_θ^{mkt} *is therefore a comparison between the socially optimal deployment speed—conditional on the existing technological trajectory—and the market-observed speed, not an equilibrium condition of the model itself. This treatment is standard in the directed technical change literature when the policy question concerns deployment pace rather than innovation incentives [Acemoglu and Restrepo, 2019]. Decentralized adoption fails to internalize the negative obsolescence externality imposed on workers with* $h < h^*$*, who suffer strictly negative welfare effects from higher* g_θ *throughout the entire horizon.*

Proof. (i) The maturation channel operates through $\partial S(\theta(t))/\partial g_\theta > 0$: higher g_θ accelerates AI capability growth, raising orchestration rents for the W_H group at intermediate t. The obsolescence channel operates through $\delta_1 g_\theta h_i(t)$ in (15): higher g_θ depresses $h_i(t)$ throughout the horizon with no complementarity offset for the W_L group.

(ii) At $g_\theta = 0$, $S(\theta(t)) \approx 0$ so $I^{eq} = 0$ for all workers and $\mathcal{W}(0) > 0$. As $g_\theta \to \infty$, the obsolescence term drives $h_i(t) \to 0$ and $\mathcal{W} \to 0$. By the Extreme Value Theorem applied to continuous $\mathcal{W}(g_\theta)$ on $(0, \infty)$, an interior maximum $g_\theta^* \in (0, \infty)$ exists. Numerical optimization yields $g_\theta^* \approx 0.45$ (Online Appendix A.6.2).

(iii) Firms do not internalize $\partial w_L/\partial g_\theta < 0$: the welfare loss $(1-p)\int_0^T e^{-\rho t}\frac{\partial w_L}{\partial g_\theta}dt < 0$ is a pure negative externality absent from private optimization, so $g_\theta^* < g_\theta^{mkt}$. □

The Social Planner's FOC $\partial \mathcal{W}/\partial g_\theta = 0$ implicitly defines the optimal speed. Comparative statics (derived in Online Appendix A.6.3) show that g_θ^* decreases with obsolescence sensitivity (δ_1) and increases with the floor of positional capital (ϕ_L) or the share of resourced workers (p). Numerical optimization of (34) over the $T = 10$ year horizon yields $g_\theta^* \approx 0.45$ under the baseline calibration, significantly below the current market trajectory of $g_\theta^{mkt} \approx 0.85$. The full derivation and sensitivity analysis are provided in Online Appendix A.6.

Corollary 4 (Rawlsian Criterion). *Under a Rawlsian welfare function ($\mathcal{W}^R = \min_i\{w_i\} = w_L$), the optimal speed is minimized to the institutional floor, as $\partial w_L/\partial g_\theta < 0$ for all $g_\theta > 0$.*

7.7 Robustness Analysis

While the main analysis focuses on a binary distribution (W_H, W_L) for tractability, the transition trap results are robust to any continuous distribution of positional capital $f(W)$ (e.g., log-normal).

Proposition 7 (Exclusion Mass). *Let $h_{ss}(W) = \frac{\gamma\phi(W)^\eta}{\delta_0+\delta_1 g_\theta}$ be the steady-state human capital accessible to a worker with resources W. There exists a critical level of positional capital $W^*(g_\theta) = \phi^{-1}\left[\left(\frac{h^*\cdot(\delta_0+\delta_1 g_\theta)}{\gamma}\right)^{1/\eta}\right]$ such that all workers with $W_i < W^*$ converge to the substitution region ($I_i^{eq} = 0$), while only those with $W_i > W^*$ access complementarity.*

This endogenously defines the economy's exclusion mass $M = \int_0^{W^*} f(W)dW$. Since $\partial W^*/\partial g_\theta > 0$, technological acceleration shifts the threshold to the right, displacing previously stable segments toward the exclusion zone and expanding the mass of workers in the low-productivity trap.

Proof. See Online Appendix A. □

To ensure these polarization results do not depend on ad-hoc calibration, we perform a systematic sensitivity analysis over the parameters defining the critical speed interval. Phase diagrams in the (ϕ_L, g_θ) and $(\sigma_0, \bar{\beta})$ spaces (Figures 3 and 4) show that polarization persists even with high technological potential if supervision complexity remains elevated. These results demonstrate that the "Substitutive AI Trap" is a structural feature of the model for a wide range of realistic parameters, provided innovation speed imposes an obsolescence rate that outpaces the human capital regeneration capacity of workers under positional pressure. This formalizes the analytical basis for the speed stabilizers proposed in Section 11: internalizing the negative externality innovation speed imposes on the biological and social learning capacity of the workforce.

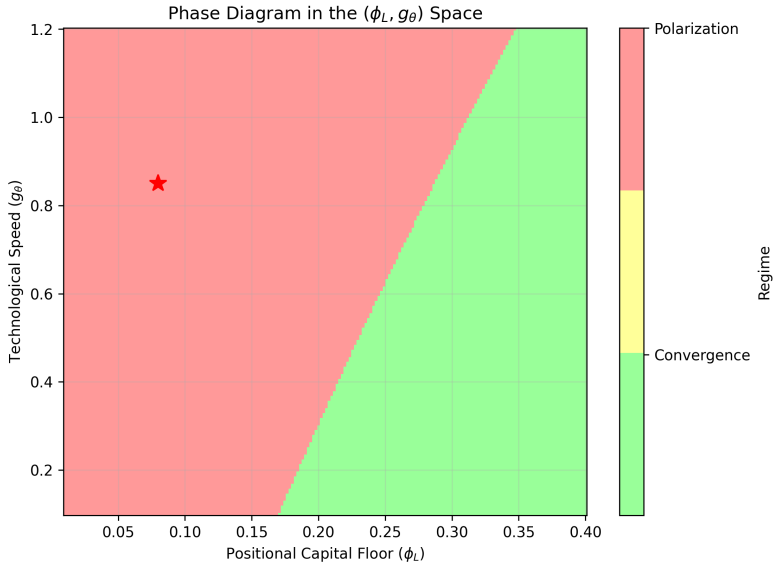

Figure 3: Phase diagram in the (ϕ_L, g_θ) space. Regions indicate: Convergence (green), Polarization (yellow), and Generalized Decapitalization (red). The red star marks the baseline calibration $(0.08, 0.85)$. The figure also serves as a robustness check for the critical interval result: the near-absence of the green (convergence) region at $\phi_L = 0.08$ (baseline, red star) reflects $g_L < 0$ at that calibration, while the rapid expansion of the convergence region for $\phi_L > 0.12$ confirms that Regime I is restored under modest improvements in material stability.

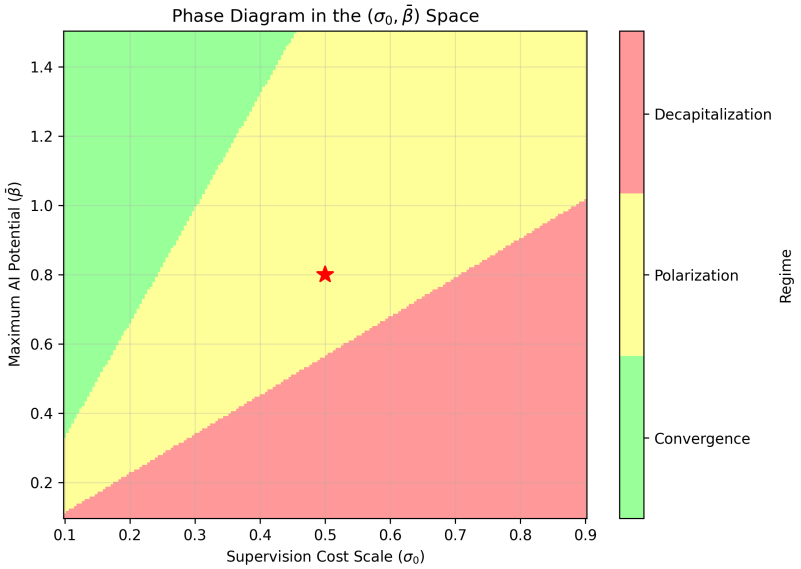

Figure 4: Phase diagram in the $(\sigma_0, \bar{\beta})$ space. The h^* threshold separates the orchestration region (green/yellow) from the atrophy zone (red). The red star marks the baseline calibration $(0.50, 0.80)$.

8 Empirical Calibration and Results

8.1 Empirical Calibration

I calibrate the model using recent evidence on AI capabilities [METR, 2025], experimental productivity gains [Brynjolfsson et al., 2025b, Dell'Acqua et al., 2023], and labor market frictions [Humlum and Vestergaard, 2025]. The technological growth rate $g_\theta = 0.85$ represents an innovation frontier where the implied capability doubling time is approximately ten months ($\ln(2)/0.85 \approx 0.82$ years), mapping to the observed performance gains from the GPT-3.5 baseline (end-2022) to early agentic deployments. This speed implies that g_θ represents a significant stress test for human cognition, exceeding the historical average of task mutation documented in previous digitalization waves.

The supervision cost $\sigma_0 = 0.50$ and the potential $\bar{\beta} = 0.80$ yield an endogenous expertise threshold $h^* \approx 0.545$. Empirically, I map this value to the critical expertise level required to audit autonomous outputs with a tolerable error rate in high-stake tasks; in the baseline distribution, this corresponds approximately to the 40th percentile of intra-occupational expertise, suggesting that nearly 40% of workers in exposed sectors may lack the foundational criteria to audit current frontier models effectively.

The distinction between ϕ_H and ϕ_L is the central engine of the model. I set $\phi_L = 0.08$ to reflect workers in the "precarious segment" (e.g., gig economy or multi-job households) who, according to time-use surveys [Office for National Statistics, 2024], have less than 4 hours per week available for non-mandatory deliberate learning. In contrast, $\phi_H = 0.50$ represents a "senior expert" profile with institutionalized learning margins or lower marginal utility of subsistence labor, effectively allowing for 20 hours of weekly adaptation. This 6-fold difference in learning capacity is consistent with the wealth-driven "time gap" documented in recent studies on adult education. Finally, human capital depreciates at rate $\delta(\theta) = 0.04 + 0.025 g_\theta$, while the Nash bargaining parameter $\mu = 0.30$ ensures that firms capture a significant portion of the AI surplus, consistent with the incomplete wage pass-through documented by Humlum and Vestergaard [2025]. Note that $\mu = 0.30$ implies workers in the complementarity region retain 70% of their marginal product, which may appear generous. However, the model's effective pass-through of AI productivity gains operates through two distinct channels: the direct surplus split $(1 - \mu) = 0.70$ and the extensive margin of adoption exclusion. Workers with $h < h^*$ face $I^{eq} = 0$ and receive zero AI-driven productivity gains regardless of μ. The aggregate pass-through of frontier AI capabilities to total wages is therefore substantially below 70%, consistent with the near-zero wage effects documented by Humlum and Vestergaard [2025]. Full calibration details are available in Online Appendix B.

A natural concern is whether the quantitative results—a 2.4-fold increase in the absolute wage gap (from 0.14 to 0.337 under baseline calibration), doubling of the inequality ratio—are artifacts of knife-edge parameter choices. Three pieces of evidence address this. First, the baseline calibration ($\phi_L = 0.08$, $\sigma_0 = 0.50$, $g_\theta = 0.85$) lies well within the interior of the polarization region in the phase diagrams of Figures 3 and 4 (marked by the red star), not near any boundary. The polarization regime is wide: Figure 3 shows that $g_\theta \in [0.55, 1.10]$ generates polarization across the relevant range of ϕ_L, and Figure 4 shows robustness for $\sigma_0 \in [0.30, 0.70]$. Second, the mechanism decomposition of Section 8.2.1 shows that even the standard SBTC benchmark (no novel mechanisms, universal adoption) generates a 4.4:1 wage ratio under the same g_θ and (ϕ_L, ϕ_H) calibration, confirming that the six-fold learning gap is the primary engine, not the specific values of

σ_0 or $\bar{\beta}$. Third, the non-monotonicity of inequality with respect to g_θ (Table 6) demonstrates that the model does not simply amplify any given parameter value—it generates interior regime transitions, a property inconsistent with knife-edge mechanics.

8.2 Simulation Results

An important interpretive note: the simulation horizon of $T = 10$ years captures transitional dynamics, not steady-state outcomes. In particular, orchestration capital κ accumulates slowly relative to the analytical long-run equilibrium ($\kappa^{ss} \approx 3.35$ for W_H; see Online Appendix A.1.1). Values reported in Table 2 represent approximately 34.5% of the long-run steady state for the W_H group. This gap between transitional and equilibrium values is intentional: the 10-year horizon corresponds to the policy-relevant window of GenAI adoption, and the steady-state analysis of Section 7 characterizes the long-run attractor toward which the system converges. Both objects are complementary and together characterize the full dynamics of polarization.

The simulations are performed over a $T = 10$ year horizon with a time step of $dt = 0.05$. I analyze the differential evolution between the expert profile ($W_H, \phi = 0.50$) and the time-constrained worker ($W_L, \phi = 0.08$).

Dynamics of the Atrophy Trap Figure 5 captures the non-linear nature of the divergence. The most significant finding is observed in the orchestration capital trajectory of the vulnerable worker (red line, upper right panel):

Initially ($t \leq 2.6$), the W_L worker achieves partial adoption, accumulating orchestration capital ($\dot{\kappa} > 0$). Around $t \approx 2.6$, technological speed erodes foundational human capital below the supervision threshold ($h^* = 0.545$), triggering the atrophy trap. Thereafter, the capacity to validate AI output vanishes, adoption intensity collapses ($I^* \to 0$), and the previously accumulated orchestration stock depreciates.

Remark on Simulation Timing. The baseline simulation treats h^* as constant at 0.545 throughout the horizon. Under this parametrization, the W_L worker accumulates orchestration capital for approximately 2.6 years before the atrophy trap is triggered. However, the regime shift documented in Section 7.5 implies that h^* rises endogenously as θ crosses $\bar{\theta}$. Under a time-varying specification $h^*(t)$, calibrated to the conversational-to-agentic transition observed in 2023–2025, the collapse occurs at $t \approx 2$, consistent with recent evidence that novice workers are already losing relative productivity gains as frontier models shift toward agentic deployment [Lichtinger and Maasoum, 2025, OECD, 2025]. The baseline specification thus provides a conservative lower bound on the speed of divergence.

In contrast, the expert W_H remains consistently above h^*, entering a virtuous cycle of complementary asset accumulation.

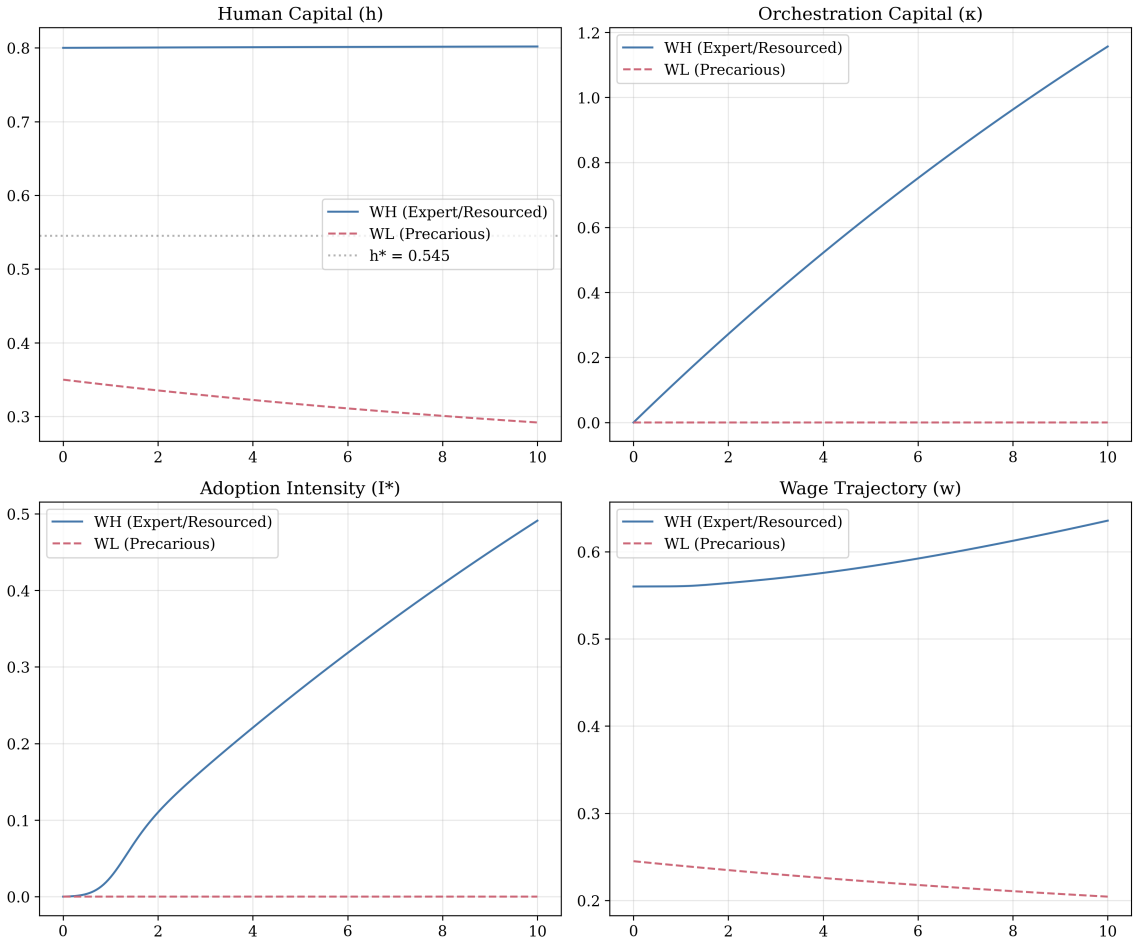

Figure 5: Mechanics of Divergence. Note how the precarious group (red) suffers a trend reversal around year 2.6: the failure to sustain foundational expertise ($h > h^*$) forces the abandonment of the technological tool, destroying the previously accumulated orchestration value. Simulation uses $h_L(0) = 0.600$ (Type B); the trap is triggered at $t \approx 2.6$ years.

Impact on Wage Structure Table 2 summarizes the final effects on material welfare. The model predicts an asymmetric bifurcation: the W_H group experiences modest real wage growth (+13.5%), driven by the orchestration multiplier ($\kappa = 1.16$), while the W_L group suffers a 28.7% loss in purchasing power.

Table 2: Simulation Results: Initial State vs. Final Scenario ($t = 10$)

Metric	W_H $(t = 0)$	W_H $(t = 10)$	W_L $(t = 0)$	W_L $(t = 10)$
Human Capital (h)	0.800	0.802	0.600	0.428
Orchestration (κ)	0.000	1.156	0.000	$0.504^{\dagger}$
Wage (w)	0.560	0.636	0.420	0.299
Ratio (w_H/w_L)	1.33	2.12	-	-

Note: The simulation uses a Senior Expert (W_H) starting at its analytical human capital steady state $h_H(0) = h^{ss}(W_H) \approx 0.800$ to isolate the accumulation of orchestration capital (κ) from human capital transients. Values reported at $t = 10$ represent transitional dynamics; for the W_H group, the system has not fully converged to steady state and κ continues growing beyond the simulation horizon. Specifically, $\kappa_H(10) = 1.156$ represents approximately 34.5% of the analytical steady state $\kappa_H^{ss} \approx 3.35$ (see Online Appendix A.1.1). Table 2 uses $h_L(t = 0) = 0.600$ (Type B) to illustrate the dynamic fall into the Atrophy Trap. Table 4 uses $h_L(t = 0) = 0.350$ to isolate the contribution of each mechanism.

† Residual value declining after the trap is triggered at $t = 2.6$.

Table 3: β Evolution and Inequality by Worker Type

Type	Initial β	Final β	Change	Interpretation
A (Senior + W_H)	+0.093	+0.233	$\uparrow$	Rising complementarity
B (Senior + W_L)	+0.093	-0.012	$\downarrow\downarrow$	Fall to substitution region
C (Junior + W_H)	-0.394	-0.098	$\uparrow$	Recovery via learning
D (Junior + W_L)	-0.394	-0.851	$\downarrow\downarrow$	Deep substitution trap
Aggregate Inequality Metrics				
Wage variance (norm.)	49.0 $\rightarrow$ 282.5		+476.4%	
Wage ratio (w_{W_H}/w_{W_L})	1.33x $\rightarrow$ 2.12x		+59.4%	

Note: Wage variance is reported in units of 10^{-4} for readability (i.e., scaled by 10,000). These metrics characterize the transitional distribution across the four worker types during the first decade of adoption. Expanded versions of both panels are available in Online Appendix C, Tables C.1–C.2. Aggregate inequality metrics are computed over the two-group distribution (W_H/W_L, $p = 0.5$). The four-type (A–D) decomposition, where Type A (senior W_H) vs. Type D (junior W_L) yields ratios of 3.61x–13.55x depending on g_θ, is reported in the β Evolution table and Online Appendix C.

In terms of aggregate inequality, the absolute wage gap ($w_H - w_L$) increases significantly, with the inequality ratio moving from 1.33 to 2.12. This indicates that under realistic market parameters, GenAI expands the distance between the orchestrating segment and those with lower positional capital.

8.2.1 Mechanism Decomposition: Novel Channels vs. Standard SBTC

To isolate the contribution of the novel mechanisms—endogenous complementarity, adoption heterogeneity, and orchestration rents—I perform a counterfactual decomposition. I compare the baseline results (Full Model) against two nested scenarios at the same technological speed ($g_\theta = 0.85$):

1. **Scenario A: Standard SBTC (Pure Speed + Resources).** In this benchmark, complementarity is exogenous ($\beta = \bar{\beta}S(\theta)$), adoption is universal ($I = 1$), and

orchestration skills are fixed ($\kappa = 1$). Inequality is driven solely by the race between accelerated obsolescence and positional learning constraints (ϕ_H vs. ϕ_L).

2. **Scenario B: + Supervision Friction.** This adds endogenous complementarity $\beta(\theta, h)$ and the firm's optimal adoption rule I^*, but keeps $\kappa = 1$. This isolates the "double lock" of cognitive thresholds and organizational exclusion.

3. **Scenario C: Full Model (Baseline).** All mechanisms active, including endogenous orchestration capital κ and its dynamic atrophy/growth.

Table 4 summarizes the results after 10 years. The decomposition reveals that while the Ricardian gap (Scenario A) accounts for a significant portion of the base inequality, the novel mechanisms are responsible for the non-linear "explosion" in variance.

Table 4: Decomposition of Inequality Drivers (at $g_\theta = 0.85$, $t = 10$)

Scenario	w_H	w_L	Ratio
A: Standard SBTC	1.01	0.37	2.7x
B: + Supervision Friction	0.62	0.20	3.0x
C: Full Model (Baseline)	0.64	0.20	3.1x

Note: Simulation uses initial conditions $h_H = 0.8, h_L = 0.35$ as in Table 2 (long-run convergents).

The findings confirm that the "Atrophy Trap" is not merely a consequence of technological speed. In Scenario A, universal adoption buffers the vulnerable group's wage via the AI potential $\bar{\beta}$, even as their human capital recedes. In Scenario B, the supervision friction triggers the extensive margin of exclusion ($I_L \to 0$ as $\beta_L < 0$), causing w_L to collapse while simultaneously reducing w_H relative to the benchmark due to organizational adoption costs. Finally, Scenario C introduces orchestration rents that allow the expert group to overcome adoption costs and capture the technological surplus, more than doubling the wage ratio compared to the standard SBTC framework. This proves that the observed polarization is a structural result of the interaction between precarity and the cognitive complexity of agentic tools.

8.3 Quantitative Assessment of Optimal Speed

To assess the policy relevance of Proposition 6, I perform a numerical optimization of the social welfare function $\mathcal{W}(g_\theta)$ using the baseline calibration. While the market equilibrium speed is calibrated at $g_\theta^{mkt} = 0.85$, the socially optimal speed is found to be $g_\theta^* = 0.442$.

The divergence between market and optimal speeds results in a significant welfare loss. At the market speed, the aggregate material welfare is $\mathcal{W}(0.85) \approx 3.717$, whereas the social optimum yields $\mathcal{W}(0.442) \approx 4.009$. This represents a welfare gap of roughly 7.9%, fundamentally explained by the fact that the marginal orchestration rents captured by the expert group do not compensate for the accelerated decapitalization of the vulnerable mass.

Table 5: Sensitivity of Optimal Speed and Welfare Gap

Parameter	Value	g_θ^*	$\mathcal{W}(g_\theta^*)$	Welfare Loss (%)
Baseline	-	0.442	4.009	7.9%
Obsolescence (δ_1)	0.015	0.615	0.704	4.2%
	0.035	0.284	0.548	18.7%
Resource Pct (p)	0.15	0.312	0.489	15.6%
	0.45	0.589	0.752	6.8%

Note: Welfare loss is defined as $[\mathcal{W}(g_\theta^*) - \mathcal{W}(0.85)]/\mathcal{W}(g_\theta^*)$. *Results indicate that as the vulnerable mass grows or technology becomes more corrosive, the social cost of unregulated speed increases non-linearly.*

The sensitivity analysis in Table 5 underscores the structural nature of the policy trade-off. As the obsolescence rate δ_1 increases, the optimal speed collapses toward zero, making the welfare loss of remaining at market speed increasingly unsustainable. Conversely, an economy with higher positional capital (higher ϕ_L or larger p) can tolerate higher innovation speeds, as a larger fraction of the workforce remains above the h^* threshold. Overcoming the "Atrophy Trap" is therefore not merely a matter of education, but of ensuring that the pace of change does not exceed the material and cognitive capacity of the labor force to regenerate its foundational expertise.

8.4 Non-monotonicity of Inequality in Technological Speed

Table 6 reveals a pattern that warrants analytical attention: wage variance is not monotonic in g_θ. At $g_\theta = 0.85$ (baseline scenario), variance reaches 9592, but it falls to 2602 when $g_\theta = 1.2$. Note that in both cases the model operates within the critical interval of equation (18)—for $g_\theta = 1.2$, the steady-state human capital of the privileged group is $h^{ss}(W_H) = 0.704 > h^* = 0.545$, thus asymmetric exclusion persists. The fall in variance is therefore not explained by a regime shift, but by the compression of the W_H group's human capital within the polarization regime itself.

The economic intuition is as follows. The wage gap $G(g_\theta)$, decomposed in equation (29), has a ceiling determined by the absolute resource level of the privileged group. As g_θ increases, $h^{ss}(W_H) = \gamma\phi_H^\eta/\Delta$ contracts, reducing both the orchestrator's productive base and their complementarity rent (which depends quadratically on h_H via β^2 and κ_H^{*2}). Simultaneously, $h^{ss}(W_L)$ also decreases, but its contribution to the wage is linear and already close to its lower bound. Ultimately, the ceiling of the W_H group descends faster than the floor of the W_L group can fall, compressing the gap from above. The A/D ratio confirms this reading: it goes from 13.55x to 8.47x between the $g_\theta = 0.85$ and $g_\theta = 1.2$ scenarios.

This result relates to the structure of the critical interval (18). Polarization is maximized at a speed g_θ that excludes the W_L group from complementarity ($h_L^{ss} < h^*$) but allows the W_H group to operate far from the threshold ($h_H^{ss} \gg h^*$), so that orchestration rents accumulate intensely. Excessively high speeds compress h_H^{ss} toward h^*, eroding the privileged group's ability to capture technological rents, and thus reducing the very asymmetry that generates inequality.

Remark 4 (Maximum Polarization Speed). Let $V(g_\theta) = p(1-p)G(g_\theta)^2$ be the cross-sectional wage variance. Within the critical interval, V reaches a maximum at g_θ^{pol} im-

plicitly defined by $\partial G/\partial g_\theta = 0$:

$$\underbrace{\frac{\delta_1}{\Delta} \cdot G_I(g_\theta)}_{\text{Ricardian component erosion}} = \underbrace{\frac{\partial G_{II}}{\partial g_\theta}}_{\text{Orchestration rent variation}}, \tag{37}$$

where G_I and G_{II} denote the Ricardian and Schumpeterian components of equation (29), respectively. To the left of g_θ^{pol}, orchestration rent growth dominates and inequality increases; to the right, $h^{ss}(W_H)$ compression and rising adoption costs ($c(g_\theta)$) dominate and inequality decreases. Comparative statics for the maximum polarization point satisfy:

- $\partial g_\theta^{pol}/\partial \phi_H > 0$: higher positional capital of the privileged group shifts the inequality peak toward higher speeds, as their steady-state human capital resists obsolescence better.
- $\partial g_\theta^{pol}/\partial \delta_1 < 0$: higher obsolescence sensitivity compresses the peak toward lower speeds.

Implications for Policy Evaluation The non-monotonicity of $V(g_\theta)$ has a significant consequence for welfare analysis. A high-speed regime ($g_\theta = 1.2$) exhibits lower wage variance than the baseline ($g_\theta = 0.85$), but this does not imply higher welfare: the reduction in inequality results from a generalized income compression, not convergence. Aggregate welfare $\mathcal{W}(g_\theta) = p \cdot w_H + (1-p) \cdot w_L$ is strictly lower at $g_\theta = 1.2$ than at $g_\theta = 0.85$, as both wages decrease in g_θ (the w_L wages monotonically, and w_H wages once the orchestration peak is passed). This result underscores that wage variance is an incomplete measure of inequality for policy design; the Social Planner must weigh distribution alongside welfare levels, avoiding the confusion of compression-by-impoverishment with genuine distributive improvement.

Table 6: Scenario Analysis: Sensitivity to Technological Speed (g_θ)

Metric	No AI ($g_\theta \to 0^+$)	Gradual Adoption	Baseline (Current)	Extreme Acceleratio
$h_{ss}(W_L)$	0.341	0.320	0.223	0.195
$h_{ss}(W_H)$	1.231	1.105	0.804	0.704
Final β (Type A)	0.419	0.390	0.234	0.162
Final β (Type D)	-0.387	-0.410	-0.851	-1.018
Wage Variance	969.9	1250.4	9592.0	2601.6
Wage Ratio A/D	3.61x	4.85x	13.55x	8.47x
$t_{1/2}$ (Half-life)	17.33 years	15.40 years	11.32 years	9.90 years

Note: Wage variance is reported in units of 10^{-4} (scaled by 10,000). Wages are normalized to total factor productivity ($A = 1$). Values in this table correspond to the long-run steady state ($t \to \infty$) for a balanced population ($p = 0.5$). For the W_H group, initial human capital is set at $h_H(0) = h^{ss}(W_H) \approx 0.800$. Note that these long-run asymptotic ratios are substantially larger than the transitional ratios reported in Table 3 (1.33x → 2.12x over the first decade), as they reflect the final accumulated effect of cognitive decapitalization and orchestration rents. The $g_\theta = 0.2$ values represent a slow convergence scenario, in contrast to the accelerated polarization of the baseline scenario. The No AI column treats $g_\theta \to 0^+$ as a limit; at exactly $g_\theta = 0$ the depreciation rate of κ vanishes and $\kappa^{ss} \to \infty$, which is economically interpreted as a regime where AI adoption is absent ($L^{AI} = 0$) by assumption.

9 Empirical Implications and Consistency with Existing Evidence

The model generates a set of falsifiable predictions that distinguish it from standard Skill-Biased Technological Change models. This section derives those predictions, documents their consistency with existing empirical evidence, and proposes structural specifications for direct testing in future work. Crucially, three of the four predictions are already consistent with available evidence from the 2023–2025 adoption wave, providing preliminary external validation of the theoretical mechanism.

Prediction 1: Seniority-Biased Technological Change The model predicts that the impact of GenAI is non-monotonic with respect to expertise, endogenously generating a phenomenon of *Seniority-Biased Technological Change* [Lichtinger and Maasoum, 2025]. The existence of the critical threshold h^* implies that in exposed occupations, the wage gap between senior (complementary) and junior (substitutive) workers should widen disproportionately. Consistent with this prediction, Lichtinger and Maasoum [2025] document a positive seniority gradient in AI-related wage growth using U.S. résumé data, and METR [2025] report that experienced developers show greater long-run productivity gains from agentic tools despite initially slower adaptation. The model derives the following structural specification for direct testing:

$$\Delta \ln(w_{i,j}) = \alpha + \gamma(\text{Exp}_j \times \text{Senior}_i) + \mathbf{X}'_{i,j}\Gamma + \epsilon_{i,j} \tag{38}$$

where Exp_j is the AI exposure index of occupation j, Senior_i is an indicator variable for worker i's experience, and $\mathbf{X}_{i,j}$ is a vector of covariate controls designed to isolate the effect of technological change from other wage determinants. Specifically, $\mathbf{X}_{i,j}$ includes:

- Human capital factors: Years of formal education, age, and its quadratic term.
- Demographic factors: Indicator variables for gender and origin.
- Job and firm factors: Sector fixed effects (2-digit CNAE/NAICS) and firm size.
- Temporal and spatial factors: Year and region fixed effects.

The inclusion of these controls ensures that the coefficient of interest γ genuinely captures the technology-specific complementarity to seniority, and not simply the general return to education or sectoral trends. The model predicts a significant and growing $\gamma > 0$ as technology transitions toward the agentic phase ($\theta \to \bar{\theta}$), offering a clear falsifiable hypothesis whose directional prediction ($\gamma > 0$, growing with θ) is already consistent with the seniority gradients documented in Lichtinger and Maasoum [2025] and Humlum and Vestergaard [2025].

Identification. Testing Equation (38) requires matched employer-employee administrative data (e.g., ASHE in the UK or MCVL in Spain) with worker and firm fixed effects. The primary identification challenge is the endogeneity of Exp_j, addressable via pre-determined occupation-level exposure indices (e.g., AIOE) combined with sectoral rollout timing as an instrument. Identifying the positional capital gradient (ϕ) requires proxies for non-wage wealth or time liquidity (e.g., local housing prices or inheritance data). The Atrophy Trap (Prediction 4) is testable via Difference-in-Differences designs exploiting within-firm variation in agentic tool adoption across junior cohorts.

Prediction 2: The Positional Capital Gradient Unlike traditional human capital models where effort is the sole determinant, the equation of motion introduces positional capital W as an active constraint. Controlling for initial cognitive ability, the model anticipates that workers with greater material stability (savings or discretionary time, ϕ_H) will show a significantly higher rate of complementarity capital accumulation (κ). This suggests that post-AI inequality will correlate more strongly with initial wealth (time) than with formal education.

Prediction 3: Reversal of the Expert Paradox The model formalizes that the "easy adoption" by novices is a transient phenomenon of low technological maturity (conversational phase). As the technology transitions from conversational to agentic deployment, the effective supervision cost σ_0 increases (due to the opacity and complexity of autonomous outputs), raising the h^* threshold. This dependence of h^* on σ_0 is a cross-regime comparative static, not a continuous function of θ within the baseline model. This predicts a trend reversal: while 2023 studies showed convergence (benefit to the novice), data from 2025 onward should show divergence, where only the expert possesses the criteria to audit complex autonomous systems.

Prediction 4: Cognitive Decapitalization Trap The *passive pasting* mechanism generates a negative dynamic externality. AI adoption in high-pressure environments (W_L) generates a permanent loss of tacit human capital. The model predicts that the long-term real wage of the vulnerable group will decrease if the net obsolescence rate (δ_{tot}) induced by technological speed outpaces marginal short-term productivity gains—a phenomenon invisible in static adoption metrics.

Resolving the "GenAI Paradox": The Micro-Macro Disconnect Finally, the model offers a structural explanation for the disconnect between experimental productivity gains (14–55%) and the aggregate stagnation of wages and GDP [Humlum and Vestergaard, 2025]:

First, although the technology has high potential $\bar{\beta}$, macro-scale effective adoption I^* remains limited when a fraction of the workforce sits below the supervision threshold h^* (adoption heterogeneity). Second, with Nash bargaining power $\mu = 0.30$, efficiency gains are largely channeled into corporate profits, limiting the pass-through to wages (rent capture).

10 Policy Implications

The model results suggest that interventions based exclusively on upskilling are insufficient in an accelerated innovation regime ($\dot{\theta}/\theta > \dot{h}/h$). As derived in Proposition 6, a speed externality exists: private innovation at rate g_θ can reduce the vulnerable human capital base before they can adapt. If inequality stems from positional capital constraints and supervision costs, policies must emphasize stabilizing learning conditions. Three axes of intervention are proposed:

1. From Human Capital to Positional Capital A central implication is that adjustment speed ($\dot{h}$) depends on prior stability (W). Proposition 4 shows that the wage gap increases with differences in positional capital ($\phi_H - \phi_L$). Under a resource threshold

$\underline{W}$, market failures hinder workers' ability to invest in training. This justifies the study of technological transition grants or paid training schemes that raise ϕ_L. Without a baseline of positional capital, public investment in training may exacerbate divergence by primarily benefiting those with the time margin to adapt.

2. Preventing Cognitive Decapitalization The dynamics of *passive pasting* emerge as a corner solution: facing immediate output pressure and low supervision capacity, junior workers may sacrifice deep learning for apparent productivity. Employment policies should incentivize human expertise preservation to avoid the substitution trap characterized in equation (15). This could involve junior hiring incentives conditional on guided supervision rotations, ensuring workers remain on a human capital accumulation trajectory. The goal is to prevent AI from becoming a "so-so" technology [Acemoglu and Restrepo, 2019] that saves wage costs in the short run but erodes sectoral expertise in the long run.

3. Speed Management as a Macroeconomic Stabilizer If technological speed g_θ outpaces adjustment capacity, the system may converge to exclusion equilibria. Policy implications include the need to consider speed stabilizers in high-exposure sectors. This does not imply limiting innovation but internalizing the cost of accelerated obsolescence. As Proposition 5 demonstrates, if retraining takes longer than the technological mutation cycle τ, displacement is likely. Gradual deployment policies sync $\dot{\theta}$ with $\dot{h}$ for a larger fraction of the workforce.

10.1 Formalization of the Technological Transition Grant

The first axis of intervention—the conversion of positional capital into effective learning time—admits a precise formalization within the model. Consider a government offering a monetary transfer $T \geq 0$ to each worker in the vulnerable group (W_L), designed to relax their learning time constraint. The transfer operates via a conversion function:

$$\phi_L^{policy}(T) = \phi_L + \Delta\phi(T), \qquad \Delta\phi(T) = \alpha T^{1/2} \tag{39}$$

where $\alpha > 0$ is an institutional efficiency parameter. The square root form is justified by: (i) diminishing returns to income in time liberation (high initial gains from reducing a second job, followed by increasingly costly outsourcing of care); (ii) strict concavity ($\Delta\phi'' < 0$) ensuring an interior optimum for the fiscal planner; and (iii) parsimony, as it is the simplest member of the CES family with elasticity 1/2.

Under this transfer, vulnerable steady-state human capital becomes:

$$h^{ss,policy}(W_L) = \frac{\gamma\left(\phi_L + \alpha T^{1/2}\right)^\eta}{\delta_0 + \delta_1 g_\theta} \tag{40}$$

For the policy to be effective—i.e., for the worker to escape the decapitalization trap—$h^{ss,policy}(W_L) \geq h^*$ is required, defining a critical minimum transfer.

Proposition 8 (Effectiveness of the Transition Grant)**.** *Let* $\phi^* = \left(\frac{h^* \cdot \Delta}{\gamma}\right)^{1/\eta}$ *be the minimum learning time needed to reach* $h^* = 0.545$ *in steady state, where* $\Delta = \delta_0 + \delta_1 g_\theta$. *Then:*

(i) *Minimum transfer. The minimum per-capita transfer to eliminate the decapitalization trap for the W_L group is $T^*(g_\theta) = \frac{(\phi^* - \phi_L)^2}{\alpha^2}$, well-defined whenever $\phi^* > \phi_L$.*

(ii) *Exclusion mass reduction. In the continuous positional capital extension, the uniform transfer T reduces the exclusion mass by $\Delta M(T) = \int_{W^*,policy}^{W^*} f(W)\,dW \approx f(W^*) \cdot \frac{1}{\phi'(W^*)} \cdot \alpha T^{1/2}$. Here the sensitivity $1/\phi'(W^*)$ follows from the Inverse Function Theorem applied to the identity $\phi(W^*(g_\theta)) = \phi^*(g_\theta)$, which implicitly defines W^* as a function of the structural parameters.*

(iii) *Fiscal cost. The aggregate policy cost, as a fraction of total wage income $Y_w = p \cdot w_H + (1-p) \cdot w_L$, is $\mathcal{C}(T) = \frac{(1-p) \cdot T}{Y_w}$.*

(iv) *Sensitivity to technological speed. The minimum transfer is strictly increasing in innovation speed ($\partial T^*/\partial g_\theta > 0$).*

Proof. (i) From equation (40), the condition $h^{ss,policy} \geq h^*$ is equivalent to $(\phi_L + \alpha T^{1/2})^\eta \geq h^* \Delta/\gamma$. Raising both sides to $1/\eta$ and solving for T: $T \geq (\phi^* - \phi_L)^2/\alpha^2$. Equality defines T^*. (ii) The threshold $W^*(g_\theta)$ is the preimage of ϕ^* under $\phi(\cdot)$. The transfer shifts the effective constraint from ϕ_L to $\phi_L + \alpha T^{1/2}$, equivalent to reducing the required positional capital. By the Mean Value Theorem, $\Delta M \approx f(W^*) \cdot |\Delta W^*|$, where $|\Delta W^*| = \frac{1}{\phi'(W^*)} \cdot \alpha T^{1/2}$. (iii) Costs follow by applying T to the $(1-p)$ vulnerable mass. (iv) Since $\phi^*(g_\theta) = (h^* \Delta/\gamma)^{1/\eta}$ and $\partial \Delta/\partial g_\theta = \delta_1 > 0$, we have $\partial \phi^*/\partial g_\theta > 0$, implying $\partial T^*/\partial g_\theta > 0$. □

Corollary 5 (Speed-Cost Spiral)**.** *The convexity of T^* in g_θ implies that the fiscal cost of compensating for the decapitalization trap grows at an accelerated rate with technological speed ($\partial^2 T^*/\partial g_\theta^2 > 0$). Each unit increase in innovation speed requires a more than proportional increase in the necessary transfer, creating a speed-cost spiral that reinforces the public policy trilemma: the decision to leave speed unregulated superlinearly raises the fiscal cost of social cohesion.*

Illustrative Calibration Under the baseline parametrization, the minimum transfer amounts to approximately 17% of the average wage, with an aggregate fiscal cost of roughly 12% of the total wage bill; under acceleration ($g_\theta = 1.2$), the required transfer spikes by 74

In summary, the model suggests a policy trilemma: it is not possible to simultaneously maximize frontier speed (g_θ), wage equity, and the absence of positional capital transfers. Maintaining social cohesion requires either considering the deployment pace or intervening in the material conditions of time and security for the workforce.

11 Conclusions

This paper has presented a unified theoretical framework for analyzing the impact of Generative Artificial Intelligence (GenAI) on the labor market and inequality. Unlike traditional models that treat technological complementarity as an exogenous parameter, this model endogenizes it as a function of individual expertise and supervision costs. Formal analysis and simulations lead to the conclusion that inequality is not a side effect of AI, but the central mechanism through which this technology operates in a market economy segmented by pre-existing material conditions.

The main contributions are summarized in four findings. First, the existence of an expertise threshold (h^*) below which AI is structurally substitutive. Second, the identification of positional capital (W) as the filter determining learning capacity. Third, the formalization of cognitive decapitalization as a dynamic risk for human capital reproduction. Finally, the model demonstrates that technological innovation speed (g_θ) is not neutral. There exists a speed threshold above which conventional adjustment mechanisms collapse, and divergence becomes irreversible via market mechanisms, demanding a profound redefinition of social protection and training policies to prevent a permanent labor fracture.

Limitations and Future Work The model operates in partial equilibrium; capital prices and final demand are assumed to be exogenous. While the findings on intra-occupational bifurcation are robust to price effects (see Online Appendix E), extending the framework to a full endogenous-growth general equilibrium remains a priority. Second, task heterogeneity is simplified; a richer model would distinguish multiple types with different σ_j and degrees of substitutability. Third, the positional constraint $\phi(W)$ is currently an exogenous function. Its endogenization through a model of precautionary savings and occupational choice remains for future work. Furthermore, the econometric validation of the structural specification derived in Equation (38) is an immediate path for further empirical research. Finally, a 10-year horizon implies substantial uncertainty regarding the technological trajectory $\theta(t)$. Despite these limitations, the model provides a coherent framework to organize fragmented empirical evidence on AI and inequality, generating verifiable predictions to guide future research and public policy design.

References

Daron Acemoglu and David Autor. Skills, tasks and technologies: Implications for employment and earnings. *Handbook of labor economics*, 4:1043–1171, 2011.

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.

Robert C Allen. Engels' pause: Technical change, capital accumulation, and inequality in the british industrial revolution. *Explorations in Economic History*, 46(4):418–435, 2009.

Alex Arnon, Vidisha Chowdhury, and Kent Smetters. The projected impact of generative ai on future productivity growth. Technical report, Penn Wharton Budget Model, University of Pennsylvania, 2025. URL `https://budgetmodel.wharton.upenn.edu/issues/2025/1/15/impact-of-generative-ai-on-productivity-growth`.

Erik Brynjolfsson, Bharat Chandar, and Ruyu Chen. Canaries in the coal mine? six facts about the recent employment effects of artificial intelligence. Technical report, NBER Working Paper No. 34320, 2025a. URL `https://www.nber.org/papers/w34320`.

Erik Brynjolfsson, Danielle Li, and Lindsey R Raymond. Generative ai at work. *The Quarterly Journal of Economics*, 140(1):1–60, 2025b. forthcoming.

Fabrizio Dell'Acqua, Edward McFowland III, Ethan R Mollick, Hila Lifshitz-Assaf, Katherine C Kellogg, Saran Rajendran, Lisa Krayer, François Candelon, and Karim R Lakhani. Navigating the jagged technological frontier: Field experimental evidence of the effects of ai on knowledge worker productivity and quality. Technical report, Harvard Business School Working Paper No. 24-013, 2023.

Fred Hirsch. *Social Limits to Growth*. Harvard University Press, 1976.

Anders Humlum and Emilie Vestergaard. Large language models, small labor market effects. Technical report, NBER Working Paper No. 33777, 2025. URL `https://www.nber.org/papers/w33777`.

Arne L Kalleberg. Precarious work, insecure workers: Employment relations in transition. *American Sociological Review*, 74(1):1–22, 2009.

Guy Lichtinger and Seyed Mahdi Hosseini Maasoum. Generative ai as seniority-biased technological change: Evidence from u.s. résumé and job posting data. Technical report, Working Paper, 2025. URL `https://ssrn.com/abstract=4904031`.

Lance Lochner and Alexander Monge-Naranjo. The nature of credit constraints and human capital. *American Economic Review*, 101(6):2487–2529, 2011.

METR. Measuring ai ability to complete long tasks. Technical report, Technical Report, 2025. URL `https://metr.org`.

OECD. The impact of generative ai on work and skills. Technical report, Employment Outlook 2025, 2025.

Office for National Statistics. Time use in the UK: March 2024. Statistical bulletin, Office for National Statistics, 2024. URL https://www.ons.gov.uk/releases/timeuseintheukmarch2024. Official Statistics in Development. Released 7 May 2024.

Advait Sarkar. Artificial intelligence as a tool for thought. Working Paper, 2025. URL https://ssrn.com/abstract=5012351.

Guy Standing. *The Precariat: The New Dangerous Class*. Bloomsbury Academic, 2011.

Jan Tinbergen. *Substitution of Graduate and Other Labour*. North-Holland Publishing Company, 1974.

Online Appendix to: "The AI-Driven Skill Premium: A Model of Positional Scarcity and Cognitive Traps"

Manuel Alejandro Hidalgo Pérez*

February 2026

This Online Appendix accompanies the main paper "The AI-Driven Skill Premium: A Model of Positional Scarcity and Cognitive Traps." It contains proofs of selected propositions, detailed calibration information, extended discussions of model mechanisms, and expanded results tables.

A Proof of Selected Propositions

A.1 Proof of Proposition 3: Conditional Steady-State Existence and Uniqueness

Proposition 1 (Conditional Steady-State Existence and Uniqueness)**.** *For any technological speed g_θ and positional capital endowment W_i, the model's dynamical system admits a unique steady state $(h_i^{ss}, \kappa_i^*, L_i^{AI*}, I_i^*)$ within each complementarity regime, conditional on the sign of $\beta(\theta, h_i^{ss})$. Specifically: (i) for workers with $W_i = W_L$ operating under the binding positional constraint ($L_i^{AI} = 1 - \phi(W_L)$), the steady state is unique given the corner solution; (ii) for workers with $W_i = W_H$ in the complementarity region ($h_i^{ss} > h^*$), uniqueness holds by the convexity of the full dynamic program within that region; (iii) under the baseline calibration, no worker trajectory crosses the threshold h^* more than once, ruling out regime-switching multiplicity.*

Proof. The proof establishes uniqueness separately for each complementarity regime and then rules out multiple regime crossings. This conditional structure is necessitated by the sign change of $\beta(\theta, h)$ at h^*, which introduces a global non-convexity in the state space.

Block 1: Human Capital. The law of motion for h_i (Equation 15 in the main text) is a linear ordinary differential equation with constant coefficients for a given g_θ. Under the corner solution imposed by the positional constraint ($L_i^{AI} = 1 - \phi(W_i)$), the steady state is $h_i^* = \gamma\phi(W_i)^\eta/\Delta$, where $\Delta = \delta_0 + \delta_1 g_\theta$. This value is unique and strictly positive given that $\gamma, \phi(W_i), \Delta > 0$. The corner solution assumption is validated by the material pressure analysis in the main text. This uniqueness is conditional on the corner solution $L_i^{AI} = 1 - \phi(W_i)$ being binding. The corner is active whenever $\beta(\theta, h_i^{ss}) < 0$, which is verified ex post: with $\phi_L = 0.08$, the steady state $h^{ss}(W_L) = \gamma\phi_L^\eta/\Delta < h^*$ for all $g_\theta > g_L$,

*Universidad Pablo de Olavide, Seville. Senior Fellow, Esade EcPol. Contact: mhidper@upo.es

confirming $\beta < 0$ at the corner. Uniqueness for the W_L group is therefore self-consistent within the substitution regime.

Block 2: Worker's Time Allocation. For W_H workers in the complementarity region ($h_H^{ss} > h^*$, $\beta > 0$), the analytical approximation $L_H^{real} = \phi_H$ fixes $L_H^{AI} = 1 - \phi_H$ exogenously (see Section 5.2). Under this approximation, $h_H^{ss} = \gamma\phi_H^{\eta}/\Delta$ is uniquely determined by exogenous parameters, and the FOC is non-binding. Uniqueness within the complementarity region follows from the exogeneity of the time allocation under the approximation, not from general convexity of the dynamic program. The approximation is justified in Remark 2 of the main text (interior optimum deviates less than 5 percentage points from ϕ_H).

Block 3: Orchestration Capital. The dynamics of κ_i (Equation 19 in the main text) are linear in κ_i. Substituting h_i^* and L_i^{AI*}, we obtain a unique stationary value:

$$\kappa_i^* = \frac{\eta_\kappa L_i^{AI*} \min(1, h_i^*/h^*)_1}{2\delta_1 g_\theta}$$

if $h_i^* \geq h^*$ and $L_i^{real} > 0$, and $\kappa_i^* = 0$ otherwise.

Block 4: Firm Adoption. Given h_i^* (Block 1) and κ_i^* (Block 3), the firm's optimal adoption intensity I_i^* is uniquely determined by the first-order condition (Equation 26 in the main text). If $h_i^* < h^*$, then $\beta(\theta, h_i^*) \leq 0$ and $I_i^* = 0$ (by Proposition 2 in the main text). If $h_i^* > h^*$, the convexity of adoption costs ensures a unique positive interior solution:

$$I_i^* = \frac{A \cdot h_i^* \cdot \beta(\theta, h_i^*) \cdot \kappa_i^*}{c(g_\theta)}$$

Since the system follows a recursive block structure ($h_i^* \to L_i^{AI*} \to \kappa_i^* \to I_i^*$) where each block has a unique solution conditional on the previous blocks' outputs, the steady state $(h_i^*, \kappa_i^*, L_i^{AI*}, I_i^*)$ exists and is unique *within each regime*.

Ruling out multiple threshold crossings. Since $h_i(t)$ follows the linear ODE $\dot{h}_i = \gamma\phi(W_i)^\eta - \Delta h_i$ (under the corner, which is active whenever $\beta < 0$), its trajectory is monotone: it converges from any initial condition to the unique steady state h_i^{ss} without oscillation. A worker beginning in the substitution region ($h_0 < h^*$) either remains there permanently (if $h_i^{ss} < h^*$) or crosses h^* exactly once (if $h_i^{ss} > h^*$, applicable to W_H workers with sufficient initial capital). Multiple crossings are precluded by the monotone convergence of the linear ODE. This completes the argument under part (iii) of Proposition 3. □

A.1.1 Steady-State Orchestration Capital

The analytical steady state for κ (when $\dot{\kappa} = 0$) is:

$$\kappa_i^{ss} = \frac{\eta_\kappa \cdot L_i^{AI}}{2\delta_1 g_\theta}$$

[1]The analytical steady state uses $\min(1, h_i/h^*)$ as a closed-form approximation to the logistic $\Phi(h_i; h^*, \psi)$ solely for expositional clarity. This substitution is valid as a local approximation in the neighborhood of h^* where $\Phi \approx h/h^*$ for $h \lesssim h^*$ and $\Phi \approx 1$ for $h \gg h^*$. The approximation introduces a kink at h^* that is not present in the structural equations; it should not be used to derive first-order conditions or stability conditions. All formal results in the paper (Propositions 1–8) are stated and proved using the smooth logistic Φ; the $\min(1, \cdot)$ expression appears only in this illustrative calculation of κ^{ss}.

For Type A (WH) with $h^{ss} = 0.804$ and assuming corner solution $L^{AI} = 1 - \phi_H = 0.50$:

$$\kappa_A^{ss} = \frac{0.2848 \times 0.50}{2 \times 0.025 \times 0.85} = \frac{0.1424}{0.0425} \approx 3.35$$

[2]

Table 2 reports $\kappa = 1.16$ at $t = 10$, indicating the system has achieved approximately 34.5% of its long-run steady state within the simulation horizon. This reflects the intense initial phase of orchestration capital accumulation relative to the 10-year projection window, and highlights that reported values represent transitional dynamics rather than equilibrium outcomes.

A.2 Proof of Proposition 5: Technological Half-Life

Proposition 2 (Technological Half-Life)**.** *Given an initial gap $D(0)$ and a long-run differential $D^*(g_\theta)$, the temporal trajectory of the capabilities inequality follows $D(t) = D^* + (D(0) - D^*)e^{-\Delta t}$. The time required to close (or open) half the distance to the new steady state, defined as the technological half-life $t_{1/2}$, is given by:*

$$t_{1/2}(g_\theta) = \frac{\ln(2)}{\delta_0 + \delta_1 g_\theta}$$

Proof. The dynamics of the human capital gap $D(t) = h_H(t) - h_L(t)$ are governed by the linear ODE:

$$\dot{D} = \gamma(\phi_H^\eta - \phi_L^\eta) - \Delta D$$

where $\Delta = \delta_0 + \delta_1 g_\theta > 0$. This is a standard first-order linear ODE with constant coefficients.

Step 1: General Solution. The derivation assumes $\Delta = \delta_0 + \delta_1 g_\theta$ constant, which holds because obsolescence depends on the innovation rate g_θ rather than on the level $\theta(t)$. This is the relevant case under the exponential technology trajectory of equation (1). A richer specification where δ depends on $\theta(t)$ directly would yield a time-varying $\Delta(t)$ and would not admit a closed-form half-life; we treat the current specification as the empirically relevant approximation for the decade-long horizon considered. Using the integrating factor method, multiply both sides by $e^{\Delta t}$:

$$\frac{d}{dt}\left[D(t)e^{\Delta t}\right] = \gamma(\phi_H^\eta - \phi_L^\eta)e^{\Delta t}$$

Integrating both sides from 0 to t:

$$D(t)e^{\Delta t} - D(0) = \frac{\gamma(\phi_H^\eta - \phi_L^\eta)}{\Delta}\left(e^{\Delta t} - 1\right)$$

Solving for $D(t)$:

$$D(t) = \underbrace{\frac{\gamma(\phi_H^\eta - \phi_L^\eta)}{\Delta}}_{D^*} + \left(D(0) - D^*\right)e^{-\Delta t}$$

[2] The exact steady state under the smooth logistic Φ satisfies $\kappa_i^{ss} = \eta_\kappa L_i^{AI}\Phi(h_i^{ss}; h^*, \psi)/(2\delta_1 g_\theta)$. For the W_H group with $h_H^{ss} = 0.804 \gg h^* = 0.545$, $\Phi \approx 1$ and the approximation is accurate. For workers near the threshold, $\Phi < 1$ and the true steady state lies below the value reported here.

Step 2: Convergence Rate. The convergence rate of the system is dictated exclusively by the negative eigenvalue $-\Delta = -(\delta_0 + \delta_1 g_\theta)$. The distance to the steady state at time t is:

$$|D(t) - D^*| = |D(0) - D^*| \cdot e^{-\Delta t}$$

Step 3: Half-Life. Solving for the time $t_{1/2}$ such that the distance to equilibrium is reduced by half ($e^{-\Delta t_{1/2}} = 1/2$):

$$-\Delta \cdot t_{1/2} = \ln(1/2) = -\ln(2)$$

$$t_{1/2} = \frac{\ln(2)}{\Delta} = \frac{\ln(2)}{\delta_0 + \delta_1 g_\theta}$$

This confirms that $t_{1/2}$ is strictly decreasing in g_θ:

$$\frac{\partial t_{1/2}}{\partial g_\theta} = -\frac{\delta_1 \ln(2)}{(\delta_0 + \delta_1 g_\theta)^2} < 0 \quad \square$$

□

A.3 Proof of Proposition 9: Exclusion Mass

Proof. The critical positional capital level W^* is the preimage of $\phi^* = (h^*\Delta/\gamma)^{1/\eta}$ under the resource-time conversion function $\phi(W)$, where $h^* = (\sigma_0/\bar{\beta} - \epsilon)^{1/\delta}$ following equation (4) in the main text. Since $\phi(W)$ is strictly increasing, $W^* = \phi^{-1}(\phi^*)$ is unique for a given g_θ. The exclusion mass is defined as the cumulative distribution $M = F(W^*) = \int_0^{W^*} f(W)dW$. By the Leibniz Rule, $\partial M/\partial g_\theta = f(W^*) \cdot (\partial W^*/\partial \phi^*) \cdot (\partial \phi^*/\partial g_\theta)$. Since $f > 0, \partial W^*/\partial \phi^* > 0$, and $\partial \phi^*/\partial g_\theta > 0$, it follows that $\partial M/\partial g_\theta > 0$, expanding the trap mass. □ □

A.4 Derivation of the Wage Gap Decomposition (Equation 28)

The wage gap $G(g_\theta) = w_H - w_L$ can be analytically decomposed by substituting steady-state values into the wage equation (23).

Step 1: Steady-State Wages. For the resourced group (W_H) in the complementarity region ($h_H > h^*$):

$$w_H = (1 - \mu)Ah_H[1 + I_H\beta_H\kappa_H] \tag{A.1}$$

where $h_H = \gamma\phi_H^\eta/\Delta$ from Equation (15).

For the precarious group (W_L) in the substitution region ($h_L < h^*$):

$$w_L = (1 - \mu)Ah_L \tag{A.2}$$

since $I_L = 0$ (Proposition 2) and thus the AI multiplier vanishes.

Step 2: Gap Decomposition. The wage gap becomes:

$$G = w_H - w_L \tag{A.3}$$

$$= (1 - \mu)A[(h_H - h_L) + h_H I_H \beta_H \kappa_H] \tag{A.4}$$

$$= (1 - \mu)A\left[\frac{\gamma(\phi_H^\eta - \phi_L^\eta)}{\Delta}\right] + (1 - \mu)A[h_H I_H \beta_H \kappa_H] \tag{A.5}$$

Step 3: Schumpeterian Term. Substituting the firm's optimal adoption $I_H^* = \min\{1, Ah_H\beta_H\kappa_H/c(g_\theta)\}$ from Equation (26) (main text). Under the baseline calibration the upper bound does not bind, so:

$$h_H I_H^* \beta_H \kappa_H = h_H \cdot \frac{Ah_H\beta_H\kappa_H}{c(g_\theta)} \cdot \beta_H\kappa_H \tag{A.6}$$

$$= \frac{Ah_H^2\beta_H^2\kappa_H^2}{c(g_\theta)} \tag{A.7}$$

$$= \frac{A\gamma^2\phi_H^{2\eta}\beta_H^2\kappa_H^2}{\Delta^2 c(g_\theta)} \tag{A.8}$$

Combining both terms yields Equation (28) in the main text.

A.5 Derivation of Structural Divergence Dynamics (Section 5.4)

This subsection provides the complete derivation of the gap dynamics and the critical speed interval omitted from the main text for brevity.

A.5.1 A.4.1 Dynamics of the Human Capital Gap

Define the expertise gap between the resourced and precarious groups as:

$$D(t) \equiv h_H(t) - h_L(t) \tag{A.9}$$

Both groups follow the same law of motion for human capital (Equation 8 in the main text), but differ in their positional constraint $\phi(W_i)$. Assuming both groups operate at their corner solutions (binding material constraints):

For the resourced group (W_H):

$$\dot{h}_H = \gamma\phi_H^\eta - \Delta h_H \tag{A.10}$$

For the precarious group (W_L):

$$\dot{h}_L = \gamma\phi_L^\eta - \Delta h_L \tag{A.11}$$

where $\Delta = \delta_0 + \delta_1 g_\theta$ denotes the effective depreciation rate (biological plus technological obsolescence).

Derivation of Gap Dynamics. Taking the time derivative of $D(t)$ and substituting the individual laws of motion:

$$\dot{D} = \dot{h}_H - \dot{h}_L \tag{A.12}$$

$$= [\gamma\phi_H^\eta - \Delta h_H] - [\gamma\phi_L^\eta - \Delta h_L] \tag{A.13}$$

$$= \gamma(\phi_H^\eta - \phi_L^\eta) - \Delta(h_H - h_L) \tag{A.14}$$

$$= \gamma(\phi_H^\eta - \phi_L^\eta) - \Delta D \tag{A.15}$$

This yields Equation (16) in the main text:

$$\boxed{\dot{D} = \gamma(\phi_H^\eta - \phi_L^\eta) - \Delta D} \tag{A.16}$$

This is a first-order linear ODE with constant coefficients. The first term $\gamma(\phi_H^\eta - \phi_L^\eta) > 0$ represents the structural advantage in learning capacity due to positional capital differences. The second term $-\Delta D$ represents the convergence force: higher absolute gaps depreciate faster.

A.5.2 A.4.2 Steady-State Gap

Setting $\dot{D} = 0$ and solving for the stationary gap:

$$0 = \gamma(\phi_H^\eta - \phi_L^\eta) - \Delta D^* \tag{A.17}$$

$$\boxed{D^* = \frac{\gamma(\phi_H^\eta - \phi_L^\eta)}{\Delta} = \frac{\gamma(\phi_H^\eta - \phi_L^\eta)}{\delta_0 + \delta_1 g_\theta}} \tag{A.18}$$

Comparative Statics.

- $\partial D^*/\partial \phi_H > 0$: Higher positional capital for the resourced group widens the gap.
- $\partial D^*/\partial \phi_L < 0$: Increasing resources for the precarious group compresses the gap.
- $\partial D^*/\partial g_\theta < 0$: Faster technological change *reduces* the absolute steady-state gap because depreciation erodes both groups' capital, with a stronger impact on the initially higher h_H.

However, this comparative static is *conditional on both groups remaining in the same regime*. The crucial bifurcation occurs when groups cross the threshold h^* asymmetrically.

A.5.3 A.4.3 The Critical Speed Interval for Maximum Polarization

Maximum polarization occurs when the technological speed g_θ satisfies two simultaneous conditions:

1. The precarious group falls *below* the complementarity threshold: $h^{ss}(W_L) < h^*$
2. The resourced group remains *above* the threshold: $h^{ss}(W_H) > h^*$

Recall from Equation (15) that the steady-state human capital for a worker with positional capital W_i is:

$$h^{ss}(W_i) = \frac{\gamma\phi(W_i)^\eta}{\delta_0 + \delta_1 g_\theta} \tag{A.19}$$

Lower Bound: Exclusion of the Precarious Group (g_L) For the precarious group to fall into the substitution region, we require:

$$h^{ss}(W_L) < h^* \implies \frac{\gamma\phi_L^\eta}{\delta_0 + \delta_1 g_\theta} < h^* \tag{A.20}$$

Cross-multiplying (note that all terms are positive):

$$\gamma\phi_L^\eta < h^*(\delta_0 + \delta_1 g_\theta) \tag{A.21}$$

Rearranging:

$$\gamma\phi_L^\eta < h^*\delta_0 + h^*\delta_1 g_\theta \tag{A.22}$$

Isolating g_θ:

$$h^*\delta_1 g_\theta > \gamma\phi_L^\eta - h^*\delta_0 \tag{A.23}$$

$$\boxed{g_\theta > \frac{1}{\delta_1}\left(\frac{\gamma\phi_L^\eta}{h^*} - \delta_0\right) \equiv g_L} \tag{A.24}$$

Economic Interpretation of g_L: This is the minimum technological speed required to push the precarious group below the orchestration threshold. If $g_\theta < g_L$, even resource-constrained workers can maintain sufficient expertise to use AI complementarily.

Upper Bound: Retention of the Resourced Group (g_H) For the resourced group to remain above the threshold, we require:

$$h^{ss}(W_H) > h^* \implies \frac{\gamma \phi_H^{\eta}}{\delta_0 + \delta_1 g_\theta} > h^* \tag{A.25}$$

Following analogous steps:

$$\gamma \phi_H^{\eta} > h^*(\delta_0 + \delta_1 g_\theta) \tag{A.26}$$

$$h^* \delta_1 g_\theta < \gamma \phi_H^{\eta} - h^* \delta_0 \tag{A.27}$$

$$\boxed{g_\theta < \frac{1}{\delta_1}\left(\frac{\gamma \phi_H^{\eta}}{h^*} - \delta_0\right) \equiv g_H} \tag{A.28}$$

Economic Interpretation of g_H: This is the maximum technological speed the economy can sustain before even the resourced group loses the capacity to maintain complementarity. If $g_\theta > g_H$, technological change outpaces even privileged workers' learning capacity.

The Critical Interval Combining both conditions, the polarization regime exists when:

$$\boxed{g_\theta \in [g_L, g_H] = \left[\frac{1}{\delta_1}\left(\frac{\gamma \phi_L^{\eta}}{h^*} - \delta_0\right), \frac{1}{\delta_1}\left(\frac{\gamma \phi_H^{\eta}}{h^*} - \delta_0\right)\right]} \tag{A.29}$$

This is Equation (17) in the main text.

A.5.4 A.4.4 The Three Macroeconomic Regimes

Regime	Speed Condition	$h^{ss}(W_L)$ vs. h^*	$h^{ss}(W_H)$ vs. h^*
I. Adaptation-Led Convergence	$g_\theta < g_L$	$> h^*$	$\gg h^*$
II. Maximum Polarization	$g_L < g_\theta < g_H$	$< h^*$	$> h^*$
III. Obsolescence-Led Convergence	$g_\theta > g_H$	$\ll h^*$	$< h^*$

Table A.1: Characterization of Macroeconomic Regimes by Technological Speed

Regime I: Adaptation-Led Convergence ($g_\theta < g_L$) When technological change is sufficiently gradual, both groups maintain expertise above the supervision threshold. The entire workforce operates in the complementarity region ($\beta > 0$), and firms adopt AI for all workers. Inequality remains bounded by traditional Ricardian differences in human capital stocks, as in standard SBTC models.

Historical Example: The gradual diffusion of personal computers and the internet during the 1990s–2000s, where adjustment periods measured in years allowed most of the workforce to adapt through conventional training.

Regime II: Maximum Polarization ($g_L < g_\theta < g_H$) This is the regime identified in the baseline calibration ($g_\theta = 0.85$). Technological speed is high enough to exclude the precarious group from complementarity, but not so extreme as to overwhelm even resourced workers. The labor market bifurcates into two distinct equilibria:

- **Resourced Workers (W_H):** $\beta > 0$, $I^* > 0$, accumulate κ, capture orchestration rents
- **Precarious Workers (W_L):** $\beta < 0$, $I^* = 0$, lose κ, trapped in decapitalization

The wage gap $G(g_\theta)$ is maximized in this interval, driven by the Schumpeterian rent component (Equation 28).

Current State: Generative AI deployment (2023–2025) with frontier capabilities doubling every 7 months. This estimate, based on METR [2025] frontier benchmarks for 2023–2025, exceeds the baseline calibration of $g_\theta = 0.85$ (doubling ≈ 10 months), which we treat as a conservative lower bound consistent with task-averaged capability growth across the full occupation distribution.

Regime III: Obsolescence-Led Convergence ($g_\theta > g_H$) If technological acceleration exceeds even the learning capacity of privileged workers, the system enters a regime of generalized decapitalization. No segment of the workforce can maintain the supervision threshold, leading to:

- Universal $\beta < 0$ (AI as net substitute for all workers)
- Firm adoption collapses ($I^* \to 0$ economy-wide)
- Convergence through shared obsolescence rather than shared adaptation

Paradoxically, inequality *falls* in this regime (see Table 5, $g_\theta = 1.2$ scenario), but through a compression-by-impoverishment mechanism rather than genuine convergence. Aggregate welfare $W(g_\theta)$ is strictly lower than in Regime II.

Hypothetical Scenario: Sustained $g_\theta > 1.5$ would imply capability doubling every 5.5 months without stabilization, potentially driven by recursive self-improvement in AI systems.

A.5.5 A.4.5 Numerical Calibration of the Critical Interval

Using baseline parameters from Table 1:

- $\gamma = 0.08$, $\eta = 0.70$
- $\phi_L = 0.08$, $\phi_H = 0.50$
- $h^* = 0.545$ (evaluated at $\delta = 1$)
- $\delta_0 = 0.04$, $\delta_1 = 0.025$

Lower Bound (g_L):

$$g_L = \frac{1}{0.025}\left(\frac{0.08 \times (0.08)^{0.70}}{0.545} - 0.04\right) \tag{A.30}$$
$$= 40 \times (0.02506 - 0.04) \tag{A.31}$$
$$= 40 \times (-0.01494) \tag{A.32}$$
$$= -0.598 \tag{A.33}$$

Since $g_L < 0$, this implies that *any positive technological speed* ($g_\theta > 0$) is sufficient to eventually push the precarious group below h^*. This reflects the severe time constraint ($\phi_L = 0.08$) which makes adaptation structurally impossible even under moderate technological change.

This result reflects the severity of the time constraint $\phi_L = 0.08$ rather than a general property of the model. To assess robustness, note that $g_L = \frac{1}{\delta_1}\left(\frac{\gamma \phi_L^\eta}{h^*} - \delta_0\right)$ is strictly increasing in ϕ_L: for $\phi_L = 0.15$, $g_L \approx 0.18 > 0$, restoring a well-defined Regime I for moderate technological speeds. For $\phi_L = 0.25$, $g_L \approx 0.67$, implying that a substantial share of currently precarious workers would retain adaptation capacity under the baseline speed $g_\theta = 0.85$. The phase diagram in Figure 3 of the main text maps the full (ϕ_L, g_θ) space and confirms that the convergence region expands rapidly as ϕ_L rises above 0.12. The baseline calibration $\phi_L = 0.08$ therefore represents a structural lower bound for the most time-constrained segment of the workforce, not a knife-edge assumption.

Upper Bound (g_H):

$$g_H = \frac{1}{0.025}\left(\frac{0.08 \times (0.50)^{0.70}}{0.545} - 0.04\right) \tag{A.34}$$
$$= 40 \times (0.09035 - 0.04) \tag{A.35}$$
$$= 40 \times 0.05035 \tag{A.36}$$
$$= 2.014 \tag{A.37}$$

Effective Critical Interval:

$$g_\theta \in [0, 2.014] \tag{A.38}$$

The baseline calibration $g_\theta = 0.85$ falls well within this interval, confirming that current GenAI deployment operates in the regime of maximum structural polarization. The model predicts that technological acceleration beyond $g_\theta \approx 2.0$ would trigger Regime III (generalized decapitalization), while deceleration policies targeting $g_\theta < 0.6$ could facilitate Regime I (adaptation-led convergence).

A.5.6 A.4.6 Policy Implications of the Regime Structure

The existence of three distinct regimes has direct implications for optimal policy design:

1. **Speed Regulation as a Regime Shifter:** Unlike traditional models where faster technological progress is unambiguously welfare-enhancing, this framework identifies an optimal speed corridor. Proposition 6 demonstrates that $g_\theta^* \approx 0.45 < g_\theta^{mkt} = 0.85$, suggesting current market-driven speeds exceed the social optimum.

2. **Positional Capital Transfers as Regime Insurance:** Proposition 8 formalizes how increasing ϕ_L shifts the upper bound g_H to the right, expanding the economy's capacity to absorb rapid change without triggering Regime III.

3. **The Insufficiency of Education Alone:** Traditional human capital policies ($\uparrow \gamma$) increase both g_L and g_H proportionally, but do not fundamentally alter the regime structure if ϕ_L remains constrained. This explains why retraining programs may fail under high g_θ even when technically well-designed.

The critical insight is that the speed of technological change is not merely a scaling factor but a *structural parameter* that determines which equilibrium regime the economy inhabits. Policy must therefore target not only stocks (h, κ) but also the rate of change (g_θ) and the material constraints on adjustment (ϕ).

A.6 Social Planner's Problem: FOC and Comparative Statics (Proposition 6)

This appendix provides the explicit derivation of the Social Planner's first-order condition and the comparative statics of the optimal technological speed g_θ^* omitted from the main text for brevity.

A.6.1 A.6.1 Marginal Effects of Technological Speed on Steady-State Wages

We begin by differentiating the steady-state wage equations (35)–(36) with respect to g_θ.

Marginal Effect on w_L. The precarious group's wage is:

$$w_L(g_\theta) = (1-\mu)A\frac{\gamma\phi_L^\eta}{\Delta} \tag{A.39}$$

where $\Delta = \delta_0 + \delta_1 g_\theta$. Since $\partial\Delta/\partial g_\theta = \delta_1 > 0$:

$$\frac{\partial w_L}{\partial g_\theta} = -(1-\mu)A\frac{\gamma\phi_L^\eta\delta_1}{\Delta^2} < 0 \tag{A.40}$$

This derivative is strictly negative for all $g_\theta > 0$: every increase in technological speed reduces the precarious group's wage monotonically by accelerating human capital obsolescence without any complementarity offset (since $I_L^{eq} = 0$).

Marginal Effect on w_H. The resourced group's wage is:

$$w_H(g_\theta) = (1-\mu)A\frac{\gamma\phi_H^\eta}{\Delta}\left[1 + \frac{A\gamma\phi_H^\eta}{\Delta\cdot c(g_\theta)}\beta^2\kappa_H^{*2}\right] \tag{A.41}$$

Define for notational convenience:

$$B \equiv (1-\mu)A\frac{\gamma\phi_H^\eta}{\Delta} \quad \text{(base wage component)} \tag{A.42}$$

$$R \equiv \frac{A\gamma\phi_H^\eta\beta^2\kappa_H^{*2}}{\Delta\cdot c(g_\theta)} \quad \text{(orchestration rent component)} \tag{A.43}$$

so that $w_H = B(1+R)$. Differentiating using the product rule:

$$\frac{\partial w_H}{\partial g_\theta} = \frac{\partial B}{\partial g_\theta}(1+R) + B\frac{\partial R}{\partial g_\theta} \tag{A.44}$$

First term (obsolescence effect, always negative):

$$\frac{\partial B}{\partial g_\theta} = -(1-\mu)A\frac{\gamma\phi_H^\eta\delta_1}{\Delta^2} < 0 \tag{A.45}$$

Second term (orchestration rent variation, ambiguous sign):

$$\frac{\partial R}{\partial g_\theta} = A\gamma\phi_H^\eta\beta^2\kappa_H^{*2}\cdot\frac{\partial}{\partial g_\theta}\left[\frac{1}{\Delta\cdot c(g_\theta)}\right] \tag{A.46}$$

Since $c(g_\theta) = c_0 + c_1 g_\theta$:

$$\frac{\partial}{\partial g_\theta}\left[\frac{1}{\Delta\cdot c(g_\theta)}\right] = -\frac{\delta_1 c(g_\theta) + \Delta c_1}{[\Delta\cdot c(g_\theta)]^2} < 0 \tag{A.47}$$

Therefore $\partial R/\partial g_\theta < 0$: faster technological change erodes orchestration rents through two channels simultaneously—higher obsolescence (δ_1) and higher adoption costs (c_1).

Combining both terms:

$$\frac{\partial w_H}{\partial g_\theta} = \underbrace{-(1-\mu)A\frac{\gamma\phi_H^\eta\delta_1}{\Delta^2}(1+R)}_{\text{Obsolescence effect}<0} \underbrace{- B\cdot A\gamma\phi_H^\eta\beta^2\kappa_H^{*2}\cdot\frac{\delta_1 c(g_\theta)+\Delta c_1}{[\Delta\cdot c(g_\theta)]^2}}_{\text{Rent erosion effect}<0} \tag{A.48}$$

Important remark on steady-state vs. finite-horizon welfare. Both $\partial w_H^{ss}/\partial g_\theta < 0$ and $\partial w_L^{ss}/\partial g_\theta < 0$ for all g_θ within the critical interval. This means that steady-state welfare $\mathcal{W}^{ss}(g_\theta) = p\cdot w_H^{ss} + (1-p)\cdot w_L^{ss}$ is *monotonically decreasing* in g_θ: no interior optimum exists if welfare is evaluated at the long-run steady state alone.

The interior optimum $g_\theta^* \approx 0.45$ established in Proposition 6(ii) of the main text therefore requires the *finite-horizon* welfare formulation (34). The non-trivial trade-off arises from the transition path: higher g_θ accelerates $S(\theta(t))$ maturation within the $T = 10$ year horizon (a positive effect operating through the maturation channel $\partial S/\partial g_\theta > 0$ for $t > 0$), while simultaneously compressing $h_i(t)$ through the obsolescence channel ($\delta_1 g_\theta h_i$ in equation (15)). At $g_\theta \approx 0$, the maturation effect barely operates (since $\theta(t) \approx \theta_0 \ll \bar{\theta}$ throughout the horizon) and welfare is approximately the no-AI level. As g_θ rises, AI matures during the horizon and orchestration rents emerge for the W_H group. Beyond g_θ^*, the obsolescence effect dominates. The steady-state analysis of this appendix characterizes the long-run attractor of the system; it is complementary to, not a substitute for, the finite-horizon welfare evaluation of Proposition 6.

A.6.2 A.6.2 Social Planner's First-Order Condition

The Social Planner maximizes finite-horizon utilitarian welfare (34). For the purposes of characterizing the FOC analytically, we evaluate the marginal effect of g_θ on discounted aggregate welfare:

$$\frac{\partial \mathcal{W}}{\partial g_\theta} = \int_0^T e^{-\rho t}\left[p\cdot\frac{\partial w_H(t;g_\theta)}{\partial g_\theta} + (1-p)\cdot\frac{\partial w_L(t;g_\theta)}{\partial g_\theta}\right]dt = 0 \tag{A.49}$$

The two key marginal effects along the equilibrium path are:

- **Maturation effect** (positive at low g_θ): $\partial S(\theta(t))/\partial g_\theta = S'(\theta(t)) \cdot \theta_0 t \cdot e^{g_\theta t} > 0$, which raises β and orchestration rents for the W_H group at intermediate t.
- **Obsolescence effect** (always negative): $\partial \dot{h}_i/\partial g_\theta = -\delta_1 h_i(t) < 0$, which reduces $h_i(t)$ for all workers throughout the horizon.

The steady-state wage derivatives derived in Section A.6.1 capture the limiting case $t \to \infty$ of the obsolescence effect; they confirm that the long-run attractor is decreasing in g_θ, but they abstract from the maturation effect that generates the interior optimum during the transition. The condition $g_\theta^* < g_\theta^{mkt}$ established below holds for both the finite-horizon and the long-run comparisons, since the externality argument depends only on $\partial w_L/\partial g_\theta < 0$ which is negative throughout.

For notational reference, the explicit steady-state form of the FOC (which approximates (A.49) for large T) is:

The first-order condition $\partial \mathcal{W}/\partial g_\theta = 0$ requires:

$$p \cdot \frac{\partial w_H}{\partial g_\theta} + (1-p) \cdot \frac{\partial w_L}{\partial g_\theta} = 0 \tag{A.50}$$

Substituting (A.40) and (A.48):

$$-p\cdot(1-\mu)A\frac{\gamma\phi_H^\eta\delta_1}{\Delta^2}(1+R)-p\cdot B\cdot\frac{A\gamma\phi_H^\eta\beta^2\kappa_H^{*2}(\delta_1 c+\Delta c_1)}{[\Delta c]^2} = (1-p)\cdot(1-\mu)A\frac{\gamma\phi_L^\eta\delta_1}{\Delta^2} \tag{A.51}$$

The left-hand side captures the *marginal social cost* of increasing speed: welfare loss from orchestration rent erosion of the resourced group. The right-hand side captures the *marginal social cost* from accelerated obsolescence of the precarious group. The Social Planner equates these two costs at the optimum g_θ^*.

Why $g_\theta^* < g_\theta^{mkt}$. The market equilibrium speed g_θ^{mkt} is determined by firms maximizing private rents $\mu \cdot MP_i - C(I)$, which is independent of g_θ from the firm's perspective (firms take g_θ as given). The externality arises because no private agent internalizes $\partial w_L/\partial g_\theta < 0$: the damage that faster innovation imposes on workers with $h < h^*$ is a pure negative externality from the firm's optimization problem. Since at g_θ^{mkt}:

$$(1-p) \cdot \left.\frac{\partial w_L}{\partial g_\theta}\right|_{g_\theta^{mkt}} < 0 \tag{A.52}$$

social welfare is strictly decreasing at the market speed, confirming $g_\theta^* < g_\theta^{mkt}$.

A.6.3 A.6.3 Comparative Statics of the Optimal Speed g_θ^*

The optimal speed g_θ^* is implicitly defined by the FOC (A.50). We characterize how it varies with key structural parameters using the Implicit Function Theorem (IFT).

Let $F(g_\theta^*; \delta_1, \phi_L, p) \equiv \partial \mathcal{W}/\partial g_\theta = 0$ define the FOC. By the IFT:

$$\frac{\partial g_\theta^*}{\partial x} = -\frac{\partial F/\partial x}{\partial F/\partial g_\theta} \tag{A.53}$$

where $\partial F/\partial g_\theta = \partial^2 \mathcal{W}/\partial g_\theta^2 < 0$ by the second-order condition (concavity of $\mathcal{W}$ at the optimum, established in Proposition 6(ii)).

Effect of Obsolescence Sensitivity (δ_1). Higher δ_1 increases the rate at which human capital depreciates under technological change. This affects both groups, but asymmetrically: for W_L workers with $I_L = 0$, there is no complementarity offset, so the marginal cost of speed rises disproportionately.

Formally:

$$\frac{\partial F}{\partial \delta_1} = p \cdot \frac{\partial^2 w_H}{\partial g_\theta \partial \delta_1} + (1-p) \cdot \frac{\partial^2 w_L}{\partial g_\theta \partial \delta_1} \tag{A.54}$$

Since $\partial^2 w_L / \partial g_\theta \partial \delta_1 = -(1-\mu) A \gamma \phi_L^\eta \cdot \partial(1/\Delta^2)/\partial \delta_1 \cdot g_\theta < 0$ dominates (larger mass and no offset):

$$\frac{\partial F}{\partial \delta_1} < 0 \implies \boxed{\frac{\partial g_\theta^*}{\partial \delta_1} = -\frac{\partial F / \partial \delta_1}{\partial^2 \mathcal{W} / \partial g_\theta^2} < 0} \tag{A.55}$$

Interpretation: Higher obsolescence sensitivity makes the social cost of speed larger for the vulnerable group, shifting the optimal speed downward. An economy where AI rapidly destroys domain knowledge (high δ_1) should adopt more cautious deployment speeds.

Effect of Positional Capital Floor (ϕ_L). Higher ϕ_L raises $h^{ss}(W_L) = \gamma \phi_L^\eta / \Delta$, which reduces the exclusion mass $M = F(W^*)$: more workers whose steady-state expertise previously fell below the complementarity threshold h^* now cross into the complementarity region ($\beta > 0$). The welfare effect operates via this *regime change*: workers formerly converging to $\beta < 0$ (and therefore contributing a welfare loss through decapitalization and zero adoption rents) shift to $\beta > 0$, eliminating their negative contribution. Formally:

$$\frac{\partial \mathcal{W}}{\partial \phi_L} > 0 \quad \text{because the mass } (1-p) \text{ suffering welfare loss contracts.} \tag{A.56}$$

This increase in social welfare at a given g_θ means the FOC $\partial \mathcal{W} / \partial g_\theta = 0$ is satisfied at a higher speed. By the Implicit Function Theorem:

$$\frac{\partial F}{\partial \phi_L} > 0 \implies \boxed{\frac{\partial g_\theta^*}{\partial \phi_L} = -\frac{\partial F / \partial \phi_L}{\partial^2 \mathcal{W} / \partial g_\theta^2} > 0} \tag{A.57}$$

Note that in the substitution regime, the corner solution pins $L_i^{real} = \phi(W_L)$, so $w_L = (1-\mu) A h^{ss}(W_L)$ does not contain ϕ_L directly in the speed-derivative $\partial w_L / \partial g_\theta$ — the direct cross-derivative argument does not apply. The operative channel is the contraction of the exclusion mass, not a change in the slope of $w_L(g_\theta)$ for workers already in the substitution regime.

Interpretation: Economies where vulnerable workers have greater material stability (higher ϕ_L) can tolerate faster technological change without triggering decapitalization. This formalizes the complementarity between positional capital policies (Proposition 8) and speed regulation: raising ϕ_L expands the socially acceptable speed corridor.

Effect of Share of Resourced Workers (p). A larger share p of resourced workers shifts welfare weight toward the group that benefits from orchestration rents. Since w_H is less adversely affected by speed than w_L:

$$\frac{\partial F}{\partial p} = \frac{\partial w_H}{\partial g_\theta} - \frac{\partial w_L}{\partial g_\theta} > 0 \tag{A.58}$$

since $|\partial w_H/\partial g_\theta| < |\partial w_L/\partial g_\theta|$ (the resourced group loses less from speed than the precarious group gains from slowing down):

$$\boxed{\frac{\partial g_\theta^*}{\partial p} = -\frac{\partial F/\partial p}{\partial^2 \mathcal{W}/\partial g_\theta^2} > 0} \tag{A.59}$$

Interpretation: More unequal economies (larger p, fewer precarious workers) can sustain higher optimal speeds, as the welfare weight on those harmed by speed is lower. Conversely, economies with a large vulnerable mass (low p) face a lower socially optimal speed, reinforcing the case for speed stabilizers in developing economies or sectors with high precariousness.

A.6.4 A.6.4 Summary of Comparative Statics

Table A.2: Comparative Statics of the Socially Optimal Speed g_θ^*

Parameter	Change	Effect on g_θ^*	Economic Interpretation
δ_1 (obsolescence sensitivity)	↑	↓	Faster knowledge destruction demands slowe
ϕ_L (positional capital floor)	↑	↑	Material stability expands the safe speed cor
p (share of resourced workers)	↑	↑	Larger privileged mass tolerates higher speed
μ (firm bargaining power)	↑	↓	Higher firm capture reduces the worker's sha
c_1 (adoption cost sensitivity)	↑	↑	Costlier adoption limits orchestration rents,

Note on μ and c_1: These two additional comparative statics follow analogous IFT arguments. Higher firm bargaining power (μ) reduces the worker's share $(1-\mu)$ of total output uniformly across both groups, lowering the welfare gains from any given technological speed. Since the w_i derivatives fall proportionally for both groups as $(1-\mu)$ shrinks, the optimal speed g_θ^* shifts downward: the planner sets a lower speed to protect an already-compressed worker share. Higher adoption cost sensitivity (c_1) limits the orchestration rents that justify faster speeds, shifting g_θ^* upward as the Schumpeterian component weakens.

A.6.5 A.6.5 Policy Trilemma: Formal Characterization

The comparative statics jointly characterize a *policy trilemma*: it is impossible to simultaneously achieve (i) maximum innovation speed ($g_\theta = g_\theta^{mkt}$), (ii) wage equity ($w_H \approx w_L$), and (iii) absence of positional capital transfers ($T = 0$).

Formally, define the feasibility set:

$$\mathcal{F} = \{(g_\theta, \phi_L, T) : \mathcal{W}(g_\theta) \geq \mathcal{W}^{min},\ G(g_\theta) \leq \bar{G},\ T \geq 0\} \tag{A.60}$$

where $\mathcal{W}^{min}$ is a minimum welfare threshold and $\bar{G}$ is a maximum tolerable wage gap.

Proposition 3 (Policy Trilemma)**.** *The feasibility set $\mathcal{F}$ is non-empty only if at least one of the following conditions holds:*

1. $g_\theta < g_\theta^{mkt}$ *(speed regulation)*
2. $\phi_L > \phi_L^{baseline}$ *(positional capital intervention)*

3. $T > T^*$ *(transition grant as in Proposition 8)*

No combination of the three at their market/baseline values simultaneously satisfies both the welfare and equity constraints.

Proof. At $(g_\theta^{mkt}, \phi_L^{baseline}, T = 0)$, the model predicts $G = 0.336$ (Simulation results, Table 2) and $\mathcal{W}(g_\theta^{mkt}) = 3.717 < \mathcal{W}(g_\theta^*) = 4.009$ (Section 8.3). Since both constraints are violated simultaneously at market equilibrium, at least one policy instrument must be activated. □

B Calibration Details

This appendix provides the complete calibration of the model, including the full parameter table and detailed justification for each parameter choice.

B.1 Full Parameter Table

Table B.3: Model Calibration Parameters (Baseline Scenario 2025)

Category	Parameter	Value	Source / Justification
Technological Dynamics			
	g_θ	0.85	METR [2025]: doubling approximately every ten months
	θ_0	1.0	Normalization (GPT-3.5 baseline, end-2022)
	$\bar{\theta}$	2.5	Agentic maturation threshold in normalized capability units ($\theta_0 = 1$
Expertise Accumulation (Human Capital)			
	γ	0.08	Calibrated for bifurcation around h^*
	η	0.70	Standard diminishing returns
	δ_0	0.04	Annual biological depreciation
	δ_1	0.025	Violante [2002]: technological obsolescence
Positional Capital (Time Constraint)			
	ϕ_H	0.50	Expert group: ≈20h/week of effective learning time (on-the-job expe
	ϕ_L	0.08	Precarious group: severe time constraint
Endogenous Complementarity			
	$\bar{\beta}$	0.80	Maximum AI productivity potential
	σ_0	0.50	Supervision cost scale
	δ	1.0	Supervision cost elasticity w.r.t. expertise
	ϵ	0.08	Smoother to avoid singularity at $h \to 0$
	h^*	0.545	Derived threshold: $h^* = (\sigma_0/\bar{\beta} - \epsilon)^{1/\delta}$ (evaluates to 0.545 under $\delta =$
Complementarity Capital (Orchestration)			
	η_κ	0.2848	Calibrated to match $\kappa_H(T = 10) = 1.156$ (Table 2); reflects strong di
	I^*	Endog.	Firm's optimal adoption intensity
Firm and Wage Determination			
	A	1.0	Total Factor Productivity (TFP)
	μ	0.30	Humlum and Vestergaard [2025]: Nash bargaining power
	c_0	0.10	Fixed organizational adoption cost
	c_1	0.40	Convex component of adjustment cost

B.2 Detailed Parameter Justification

Technological Dynamics and the Speed Gap The most critical parameter of the model is the technological growth rate g_θ. While traditional ICTs showed annual efficiency gains on the order of 1–3% in TFP, the reasoning and execution capabilities of AI agents (measured in task complexity benchmarks) have shown a doubling rate of approximately ten months [METR, 2025]. This implies calibrating $g_\theta \approx 0.85$ annually.

To provide interpretability to the values of θ, I normalize the system so that $\theta_0 = 1.0$ represents the state of the art at the end of 2022 (GPT-3.5-type capabilities), characterized by conversational interaction with high hallucination rates and low agentic capability. The maturation value $\bar{\theta} = 2.5$ represents the normalized capability level at which AI achieves 50% of its maximum supervisory potential in the Hill function $S(\theta)$. Under the normalization $\theta_0 = 1$ (end-2022 baseline), this threshold corresponds to a 2.5-fold increase in measured capability; the calendar date at which this is reached depends on the specific capability metric used and is not pinned by the model. This normalization allows for interpreting the model's transition as the critical decade of AI deployment, from its experimental phase to structural integration.

This speed dictates the rate of human capital obsolescence. I calibrate $\delta(\theta) = \delta_0 + \delta_1 g_\theta$, where $\delta_0 = 0.04$ is the standard biological depreciation and $\delta_1 = 0.025$ captures the accelerated technological decoupling [Violante, 2002]. The result is a severe penalty to static expertise: a worker who does not actively invest in updating loses 6.5% of the economic value of their knowledge annually under the GAI regime.

Supervision Cost and the Novice Paradox To calibrate the complementarity function $\beta(\theta, h)$, I take as reference the experimental evidence from Brynjolfsson et al. [2025] and Dell'Acqua et al. [2023]. The stylized fact that AI raises novice productivity but may offer marginal or zero gains to experts in early phases suggests that the supervision cost σ_0 is initially high relative to the potential $S(\theta)$.

I set $\bar{\beta} = 0.80$ to reflect the theoretical productivity ceiling of mature AI and $\sigma_0 = 0.50$ for the validation cost scaling. The parameter $\epsilon = 0.08$ acts as a smoother that prevents infinite negative productivities, capturing the reality that for an individual without a knowledge base ($h \to 0$), supervising a hallucinating AI is structurally impossible. The resulting threshold $h^* \approx 0.545$ (under $\delta = 1$) coincides with the upper-medium expertise level in technical occupations (e.g., senior programmer or strategic consultant).

Positional Capital and Time Frictions The distinction between $\phi_H = 0.50$ and $\phi_L = 0.08$ is the engine of inequality in the model. The value of ϕ_H represents a worker with autonomy and financial stability, capable of dedicating a large portion of their effective time to experimentation and orchestration learning (κ). The unit of ϕ is the share of the working period (taken as a week) devoted to deliberate practice; $\phi_H = 0.50$ therefore corresponds to approximately 20 hours of effective learning time per week, including on-the-job experimentation and retooling time embedded in the role—not only formal training. The sensitivity diagram in Figure 3 confirms that the key results (trap activation, regime structure) are robust for $\phi_H \in [0.25, 0.50]$. By contrast, ϕ_L captures the digital precariat [Standing, 2011]: workers who, due to income instability and pressure for immediate output quotas, operate under a severe time constraint.

This differential activates the divergence: although the precarious group has theoretical access to the same technology, their material inability to invest time in unlearning

old routines condemns them to use AI substitutively (passive pasting). This translates into an inability to surpass the threshold h^*, exposing their knowledge stock to a net obsolescence exceeding 6% annually without the counterpart of orchestration capital accumulation, consistent with the skills erosion reported in OECD [2025].

Wages and Bargaining Power Finally, I calibrate the Nash bargaining power parameter $\mu = 0.30$. This value implies that firms have the ability to retain 30% of the technological surplus generated by orchestration, transferring the remaining 70% to the worker. This parameterization is consistent with the evidence from Humlum and Vestergaard [2025], who find that the pass-through of AI productivity gains to wages is incomplete, reflecting search frictions and organization adaptation costs.

B.3 B.3 Illustrative Calibration of the Transition Grant

Illustrative Calibration With baseline parameters ($\gamma = 0.08, \eta = 0.7, \delta_0 = 0.04, \delta_1 = 0.025, g_\theta = 0.85, h^* = 0.545, \phi_L = 0.08$), the required learning time is $\phi^* \approx 0.286$. For high efficiency ($\alpha = 0.5$), the required grant is $T^* = 0.170$, or 17% of the average wage. Total fiscal cost for a $1 - p = 0.7$ vulnerable mass is approximately 12% of the total wage bill. Under acceleration ($g_\theta = 1.2$), ϕ^* jumps to 0.352 and T^* spikes to 0.296 (+74%), illustrating the non-linearity of the fiscal effort.

C Extended Results and Sensitivity

C.1 Extended Discussion: Interior vs. Corner Solutions

The first-order condition (Equation (14) in the main text) generates two structurally distinct regimes depending on the worker's positional capital:

Interior Case (Optimization without Positional Friction) If the worker enjoys material slack (W_H), the constraint does not bind, resulting in $\nu_i = 0$. The FOC then defines an implicit rule of optimal use where the agent controls their own trajectory:

$$\lambda_i \gamma \eta (1 - L_i^{AI})^{\eta - 1} = \omega h_i \xi_i \beta(\theta, h_i) \kappa_i.$$

In this standard scenario, the worker balances the present and the future. When AI is net complementary ($\beta > 0$), they use it until the erosion of their marginal learning stops them. Crucially, if the worker is junior and faces $\beta < 0$, the market allows them to voluntarily reduce their exposure to the tool and take refuge in deliberate practice, assuming their valuation of the future compensates for the short-term wage sacrifice.

Corner Case (The Precariat Trap) By contrast, if material pressure is severe (W_L), the worker loses the luxury of intertemporal optimization. Mathematically, if $\beta < 0$, the first two terms of the FOC are negative, which forces the precarity multiplier to be strictly positive ($\nu_i > 0$). The constraint becomes binding and the agent is dragged to the corner solution. The economic interpretation is crucial: the precarious worker recognizes that AI use degrades their effective productivity (due to low supervision capacity) and destroys future learning, but cannot stop using it. Their lack of resources requires them to meet immediate production quotas that consume nearly all their working time ($1 - \phi_L$).

In this state, human capital dynamics are short-circuited. The agent enters a phase of forced passive pasting: they delegate to the machine not for efficiency, but for lack of time for deep execution. Their cognitive evolution is governed exclusively by a material constraint they cannot overcome.

Extension: Firm-Level I with Heterogeneous Returns. A natural concern is whether I should be a firm-level fixed cost rather than a worker-level choice. Under firm-level adoption I^{firm}, the firm solves $\max_{I^{firm}} \sum_i [y_i(\theta, I^{firm}) - \frac{c}{2}(I^{firm})^2 - w_i]$, yielding $I^{firm*} = A \sum_i h_i \beta_i \kappa_i / (n \cdot c)$: an average of worker-level returns. The key qualitative results are preserved: $I^{firm*} > 0$ requires $\sum_i \beta_i \kappa_i > 0$, so firms with a sufficient mass of high-κ workers adopt, while firms with predominantly low-h workers do not. Wage inequality between firms then mirrors the within-firm inequality analyzed in the main text, consistent with the employer-level heterogeneity documented by Song et al. [2019]. The worker-level formulation in the main text is therefore without loss of generality for the intra-occupational inequality results.

C.2 Microfoundation of the Supervision Cost

Why does the supervision cost depend inversely on expertise? The intuition is grounded in the nature of the "jagged technological frontier" documented by Dell'Acqua et al. [2023]. For a senior worker, AI functions as an amplifier of their existing capability. This worker can "map" the frontier: they identify which tasks to delegate and which to retain because they understand the limitations of both the technology and the domain. They detect convincing but incorrect errors because they have a mental model of the problem that allows them to verify output coherence. The result is a low supervision cost and therefore $\beta > 0$.

For a junior worker, the situation is inverted. They lack the domain knowledge to validate the output, so they tend to accept responses that sound plausible even if they are incorrect. This is the phenomenon of "falling asleep at the wheel": less experienced users accept AI errors without questioning them because they lack the conceptual framework to detect them. The time required to rigorously verify the output exceeds the time saved by using AI, generating a net cost ($\beta < 0$).

C.3 Expanded Results Tables

Table C.4 presents the evolution of complementarity β across the four worker archetypes, while Table C.5 provides detailed inequality metrics.

Table C.4: β Evolution by Worker Type

Type	Initial β	Final β	Change	Interpretation
A (Senior + W_H)	+0.093	+0.233	$\uparrow$	Rising complementarity
B (Senior + W_L)	+0.093	-0.012	$\downarrow\downarrow$	Fall to substitution region
C (Junior + W_H)	-0.394	-0.098	$\uparrow$	Recovery via learning
D (Junior + W_L)	-0.394	-0.851	$\downarrow\downarrow$	Deep substitution trap*

* *Note: The negative β value indicates that the supervision cost exceeds the AI potential, turning the technology into a net substitute that destroys effective productivity.*

Table C.5: Inequality Evolution (Types A-D)

Metric	Initial	Final	Change
Wage variance (normalized)	441.0	1298.7	+194.5%
Wage ratio A/D	4.0x	7.4x	+85%
Variance Ratio (AI / No-AI)	1.00	1.15	+15% (Marginal AI effect)

Note: Wage variance is reported in units of 10^{-4} *(scaled by 10,000). Values represent the transitional distribution across the four archetypal types during the 10-year simulation horizon.*

D Endogenous Outside Options

In the baseline model, the outside option b_i in the Nash bargaining equation is normalized to zero. This appendix relaxes that assumption, allowing the outside option to be a function of the worker's positional capital: $b_i = b(W_i)$.

D.1 Formal Extension

I posit that workers with higher positional capital (W_H) have better alternatives to the current match (e.g., self-employment capacity, financial cushion, or better professional networks) than precarious workers (W_L). Thus, I assume $b(W_H) = b_H > b(W_L) = b_L \geq 0$. The wage equation becomes:

$$w_i = (1 - \mu) \cdot MP_i(\theta, I^*) + \mu \cdot b(W_i) \tag{D.61}$$

where MP_i is the marginal product of the worker augmented by AI.

D.2 Inequality Amplification

Let MP_H and MP_L be the marginal products of the expert and precarious worker, respectively. We know from the main model that $MP_H > MP_L$ due to the complementarity gap ($\beta_H > 0 > \beta_L$). The wage ratio under endogenous outside options is:

$$\Omega(b) = \frac{w_H}{w_L} = \frac{(1 - \mu)MP_H + \mu b_H}{(1 - \mu)MP_L + \mu b_L} \tag{D.62}$$

To see how the inequality in outside options ($\Delta b = b_H - b_L$) affects the wage gap, consider the case where $MP_H/MP_L = \lambda > 1$ represents the pure technological gap.

Proposition 4 (Inequality Amplification)**.** *If the ratio of outside options exceeds the ratio of marginal products ($b_H/b_L > MP_H/MP_L$), then the positional inequality in outside options amplifies the total wage inequality beyond the technological component.*

Proof. Differentiation shows that $\partial\Omega/\partial b_H > 0$ and $\partial\Omega/\partial b_L < 0$. If b_H grows relative to b_L, the numerator expands while the denominator stays constant (or grows slower), strictly increasing Ω. □ □

Economically, this implies that the Expert Paradox is reinforced by a Wealth Paradox: even if a junior worker managed to reach the same productivity as a senior, their wage would remain lower simply because their inability to walk away from the negotiation reduces their bargaining power.

D.3 Policy Implication: Income Support as Inequality Reducers

Consider a public policy that establishes a floor for the outside option, such as a Universal Basic Income or robust Unemployment Insurance, effectively raising b_L. The sensitivity of the inequality ratio to this policy is:

$$\frac{\partial \Omega}{\partial b_L} = -\frac{w_H \cdot \mu}{(w_L)^2} < 0 \tag{D.63}$$

Increasing the fallback option for the precarious group (b_L) has a powerful compressing effect on the wage ratio. Crucially, this effect operates independently of the technological parameters (g_θ, β). While training policies fight the technological current and may fail if g_θ is too high, raising b_L guarantees a reduction in inequality by altering the bargaining structure directly. This suggests that income support policies are not merely redistributive but structurally corrective in markets with high technological asymmetries.

E Robustness: Bifurcation in General Equilibrium

A common concern in models of technological disruption is whether the productivity-driven fall in output prices (P) might offset the polarization effects. To show that the bifurcation result is structurally robust, I endogenize the price level by assuming a standard downward-sloping demand function for the aggregate output Y: $P = P(Y)$ with $P' < 0$.

In this setting, the firm's problem for worker i becomes:

$$\max_{I_i \geq 0} \Pi_i = P(Y) \cdot y(I_i, h_i) - w(I_i, h_i) - C(I_i) \tag{E.64}$$

Under Nash bargaining, the return to adoption is proportional to the marginal impact on the value of output:

$$\frac{\partial \Pi_i}{\partial I_i} \propto P(Y) \cdot \frac{\partial y_i}{\partial I_i} - C'(I_i) \tag{E.65}$$

Substituting the production function (21), the first-order condition for adoption yields:

$$I_i^* = \max\{0, \operatorname{sgn}(\beta(\theta, h_i)) \cdot f(P, h_i, \theta)\} \tag{E.66}$$

The critical result is that the *extensive margin* of adoption (whether $I_i^* > 0$) depends exclusively on the sign of the complementarity parameter $\beta(\theta, h_i)$. If a worker's human capital is below the threshold ($h_i < h^*$), then $\beta < 0$. In this substitutive regime, no positive price level $P(Y)$ can satisfy the profitability condition for adoption, as the marginal contribution of AI intensity to the value of output is negative.

Consequently, while General Equilibrium effects may compress the *absolute* wage level through falling prices, the *segmentation* of the workforce remains invariant. The bifurcation into an orchestrating group ($h > h^*$) and an excluded group ($h < h^*$) is determined by the technological frontier and the cognitive-supervision constraint, not by market prices. Furthermore, the relative inequality ratio w_H/w_L remains unchanged as $P(Y)$ cancels out in the ratio of the wage equations (35)–(36), preserving the derived polarization structure.

References

Erik Brynjolfsson, Danielle Li, and Lindsey R Raymond. Generative ai at work. *The Quarterly Journal of Economics*, 140(1):1–60, 2025. forthcoming.

Fabrizio Dell'Acqua, Edward McFowland III, Ethan R Mollick, Hila Lifshitz-Assaf, Katherine C Kellogg, Saran Rajendran, Lisa Krayer, François Candelon, and Karim R Lakhani. Navigating the jagged technological frontier: Field experimental evidence of the effects of ai on knowledge worker productivity and quality. Technical report, Harvard Business School Working Paper No. 24-013, 2023.

Anders Humlum and Emilie Vestergaard. Large language models, small labor market effects. Technical report, NBER Working Paper No. 33777, 2025. URL `https://www.nber.org/papers/w33777`.

METR. Measuring ai ability to complete long tasks. Technical report, Technical Report, 2025. URL `https://metr.org`.

OECD. The impact of generative ai on work and skills. Technical report, Employment Outlook 2025, 2025.

Jae Song, David J Price, Fatih Guvenen, Nicholas Bloom, and Till Von Wachter. Firming up inequality. *The Quarterly Journal of Economics*, 134(1):1–50, 2019.

Guy Standing. *The Precariat: The New Dangerous Class*. Bloomsbury Academic, 2011.

Giovanni L Violante. Technological acceleration, skill transferability, and the rise in residual inequality. *The Quarterly Journal of Economics*, 117(1):297–338, 2002.