

The Skill Content of AI: First-Personal Experience and Returns to Human Capital

Michael B. Wong*

May 2026

Abstract

AI's reach into human work has a structural limit: its training corpus contains no first-personal experience. Existing AI exposure measures, by contrast, load primarily on cognitive complexity and are conditionally near-orthogonal to the relational dimension. We study optimal human capital investment in a model in which humans and AI are symmetric learning systems differing only in compute and data endowment, with each task combining distributional and first-personal skills. The permanent human domain — complement-type tasks with positive first-personal weight — is invariant to the compute price. As AI cheapens, an aggregate Baumol channel raises the wage of first-personal skill, which persistently overtakes ordinary cognitive skill if human-exclusive tasks are sufficiently first-personal-skill intensive. Optimal investment shifts toward first-personal skill, with distributional-skill investment non-monotone — falling early as routine cognitive work is reassigned to AI, then recovering.

Keywords: artificial intelligence, automation, comparative advantage, task content, human capital.

JEL codes: J24, O33, D23.

*HKU Business School. Email: mbwong@hku.hk. I thank Jie Gong, John Klopfer, Alan Kwan, and Jin Li for helpful discussion.

1. Introduction

In 2021, IKEA introduced the AI chatbot Billie, which within two years was resolving roughly 47% of customer inquiries. Rather than cut headcount, IKEA retrained 8,500 call-centre workers as remote interior design consultants, generating roughly EUR 1.3 billion in revenue from the new remote design channel (Reid, 2023; Ingka Group, 2022). Both jobs are non-routine cognitive work; on conventional complexity measures, the design consultation is, if anything, the easier. AI absorbed one and the workers were redeployed to the other. Why?

The answer, we argue, is architectural. A learning system whose relevant data is *first-personal* — registered in the body of someone living the situation, and absent from any training corpus — cannot be substituted by one whose data is text. First-personal inputs include the affective attunement of a teacher reading the moment a student has shut down, the co-constructed back-and-forth that lets a designer build a room with a client, the experiential authority of a clinician who has accompanied many patients through difficult news. Distributional skill, by contrast, is corpus-derived: codified knowledge and analytical patterns that scale with compute on external records. AI capability scales with compute on data that exists in external corpora; the capacities produced by first-personal experience do not, because no external corpus contains them. The fraction of tasks AI can profitably serve converges to a ceiling strictly below one as compute prices fall, and the complement — the permanent human domain — is invariant to the price of compute. In the transition, an aggregate Baumol channel raises the wage of first-personal skill, which persistently overtakes the wage of ordinary cognitive skill if human-exclusive tasks are sufficiently first-personal-skill intensive.

A recent field experiment makes the point at a single task. Dell'Acqua et al. (2023) gave consultants at BCG access to GPT-4 and measured performance across 18 tasks inside the AI frontier plus one designed to lie outside it. On the inside tasks, AI produced large gains. On the outside task — integrating an interview with an executive with internal data to form a view about organizational dynamics — AI users were 19 percentage points less likely to produce correct

answers than consultants working alone. Same professionals, same AI system, similar apparent difficulty. The outside task required judgment drawing on context that a text-corpus model cannot reconstruct.

The cross-section makes the point at scale. We build an O*NET-based index of relational content for 746 occupations and compare it against three publicly available AI exposure measures (Eloundou et al., 2023; Felten et al., 2021). The relational index correlates *positively* with each measure ($r = 0.40\text{--}0.54$), the opposite of what a one-dimensional substitution model predicts. The cause is structural: existing exposure measures load primarily on cognitive complexity, and once cognitive complexity is partialled out, the residual correlation collapses sharply — to near-zero for general measures and to a small positive residual for language-model-specific measures. The relational dimension is largely invisible to the measures that frame the policy debate. The high-exposure stratum splits sharply on the relational axis: registered nurses, elementary school teachers, and managers on one side; office clerks, secretaries, and paralegals on the other.

To formalize the argument, we model humans and AI as symmetric learning systems differing only in compute and data endowment. Each task lives in a two-dimensional space (α_{iH}, σ_i) , where α_{iH} is the task’s first-personal weight and σ_i is the within-task elasticity of substitution between first-personal and distributional inputs. Tasks are produced one-agent-per-task and assigned to the lower-cost producer; humans allocate finite lifetime investment between distributional and first-personal skill bundles. Comparative advantage between humans and AI is *derived* from the task’s skill requirements rather than imposed by an exogenous task-content rubric.

Three structural results follow. The set of complement-type tasks with positive first-personal weight ($\sigma_i < 1, \alpha_{iH} > 0$) forms a permanent floor invariant to the compute price: the within-task Leontief corner forces AI output to zero regardless of how cheap compute becomes. As AI absorbs the contestable complement of this floor, an aggregate Baumol channel raises w_H — expenditure shifts toward the floor while human supply of first-personal skill is bounded, bidding the price up. The channel extends to the factor level a documented sectoral pattern: relative prices of healthcare, education, and personal services have grown 0.6–1.3 percentage points per year

faster than consumer goods over the post-1950 record (Baumol, 2012; Helland and Tabarrok, 2019), and the framework predicts that divergence will widen rather than close as AI diffuses.

The return to distributional skill is non-monotone over the transition: w_D dips briefly as routine cognitive work is reassigned to AI, then recovers via within-human complementarity on complement-type tasks. At fixed pre-AI skill stocks the same mechanism produces the senior-vs-junior incidence pattern in the early empirical record — Hampole et al. (2025) and Humlum and Vestergaard (2025) document essentially zero earnings effects on senior workers in AI-exposed occupations through 2024 alongside contracting early-career hiring, with within-firm declines concentrated where workers’ portfolios skew toward AI-exposed task content. With skill reallocation, w_H overtakes w_D in the long run if human-exclusive tasks are sufficiently first-personal-skill intensive. The level of the long-run gap depends on the size of the permanently-human domain — a magnitude the framework’s robustness exercises bracket rather than predict.

The framework draws on multiple strands of the labor-AI literature. The task-content tradition — Autor, Levy, and Murnane (2003), Acemoglu and Restrepo (2018, 2019), Acemoglu (2025), Autor and Thompson (2025) — supplies the unit of analysis: changes occur task by task, and within-occupation expertise composition matters; but none of these papers offers a structural reason for the automation boundary to stop. AI-era theoretical extensions sharpen the question without resolving it: Ide and Talamas (2025) embed AI in a hierarchical knowledge economy à la Garicano (2000), distinguishing autonomous AI agents from non-autonomous copilots; Garicano and Rayo (2025) introduce AI into a dynamic apprenticeship model and derive a viability threshold for the apprenticeship pipeline. Our contribution grounds the automation bottleneck in the architecture of training data, making the boundary permanent rather than contingent on engineering progress. We add a second axis to the task space — the within-task elasticity σ_i alongside the H -bundle weight α_{iH} — generating both the structural floor and the transition-path dynamics that one-axis frameworks cannot deliver. And we treat human capital investment as a Ben-Porath problem against an AI agent on the same skill-formation technology, a margin most of this literature treats as predetermined.

The closest related framework is Garicano, Li, and Wu (2026), who also study whether AI removes the human from the job. They build a two-task model in which AI either assists inside a bundle or supplies the task autonomously: in weak-bundle occupations where tasks are independently reallocable, AI unbundles and displaces workers; in strong-bundle occupations where tasks are not reallocable, AI improves performance inside the bundle without removing the human. The mechanisms are complementary but distinct. Their bundle strength turns on the cost of separating tasks from the job — a job-design margin — whereas σ_i here is a within-task production elasticity that operates inside a single task’s CES. Both arguments converge on the prediction that human-exclusive content is protected; this paper grounds the protection in training-data architecture, parameterizes the boundary as a 2D task space, and yields aggregate transition dynamics — the Baumol channel and the U-shape — that a two-task model does not deliver.

A separate empirical literature constructs AI exposure measures. Eloundou et al. (2023) score LLM exposure from human ratings; Felten et al. (2021) weight O*NET ability requirements by AI’s overlap with each ability; Frey and Osborne (2017) score probability of computerisation, assigning social-intelligence variables explicit bottleneck weights; Loaiza and Rigobon (2024) introduce the EPOCH framework — five capability groups LLM-scored at the task level — naming where AI struggles. The measurement contribution here is the conditional orthogonality of the relational dimension to leading exposure measures once cognitive complexity is partialled out, complementary to Loaiza and Rigobon (2024)’s construction of a new occupation-level metric. The framework’s wage predictions also extend a documented pre-AI pattern: Deming (2017) documents faster employment and wage growth in social-skill-intensive occupations from 1980 to 2012, and Weinberger (2014) shows that the cognitive–social complementarity rose between NLSY79 and NLSY97 cohorts.

The paper proceeds as follows. Section 2 develops the mind architecture and the empirical decomposition of AI exposure. Section 3 sets up the learning-systems model. Section 4 derives the static division of labor — the permanent floor, the jagged frontier, and the common boundary between automation and measurement. Section 5 develops the analytical transition path: the Baumol

channel, the non-monotone return to D -skill, the short-run wage incidence, and the within- D portfolio margin. Section 6 traces the transition quantitatively. Section 7 draws out the implications for human capital formation.

2. Mind Architecture and AI Exposure

2.1. The architecture of AI and human minds

A deep-learning AI system is a high-dimensional function approximator trained by gradient descent on a large external corpus, almost entirely textual. Two properties of this architecture matter for what follows.

Property 1: AI knowledge is distributional and corpus-derived. The weight matrix encodes the statistical structure of the training distribution; capability is a relationship between the system and the corpus it was trained on, not a general property. The system performs well on inputs resembling the corpus and degrades on inputs that differ substantially from it. Because the corpus is constructed primarily from text — transcripts, documents, code, captioned images — the model’s knowledge is a model of how the world is described in language and other recorded signals, not of the world itself. For tasks where the recorded signal is a faithful surface over the relevant reality, this is adequate; for tasks whose relevant reality is generated outside any external record, the gap is structural rather than incidental.

Property 2: AI has no first-personal experience. Human cognition is generated by an embodied organism with stakes in outcomes. This generates a class of capacities whose exercise produces external records (notes, recordings, narratives) but whose generative signal exists only inside the experiencing organism. Examples: noticing what a patient’s posture is communicating before words are exchanged; reading when a student has shut down; the experiential authority of a clinician who has accompanied many patients through difficult news. AI systems have no access to the generative signal because no external record preserves it. This is a statement about training

data, not about consciousness.¹ The asymmetry between AI and human capability is therefore a property of what data exists in the world, not a property of current model capability.

The two properties together generate a structural prediction. AI capability scales with the data, compute, and modelling effort applied to the distributional signal in external corpora. The capacities produced by first-personal experience do not have an external corpus, so they do not scale with compute, no matter how cheap. The natural formal counterpart is a model in which the data endowment of AI on first-personally generated capacities is zero, and humans are the only source of that input.

2.2. The relational dimension is invisible to leading AI exposure measures

If first-personal skill is what AI cannot acquire, the natural prior is that occupations heavy in relational content are exactly the occupations existing AI exposure measures should rate as least exposed. Relational content is the observable subset of first-personal content: interpersonal interaction picked up by occupational data, conservative because it does not capture solo experiential judgment (a surgeon’s haptic feedback, a craftsman’s tacit feel for materials). We test this prior against the data and find it is not borne out by the raw cross-section.

We build an occupation-level relational index α_H as the simple average of importance ratings on ten O*NET items capturing interpersonal care, coordination, and influence (five Work Activities and five Skills; full list and validation in Appendix A). The first principal component over the ten items explains 60.1% of variance, with all loadings positive in $[0.24, 0.37]$; the ten items measure overlapping facets of a single coherent dimension. We compare α_H with three publicly available occupation-level AI exposure measures: the Eloundou et al. (2023) human-rated exposure score $E_2 = \beta + 0.5\gamma$, the Felten et al. (2021) general AI Occupational Exposure (AIOE) index, and the same authors’ language-modelling-specific AIOE.² Each correlates *positively* with α_H in the raw

¹The argument is independent of any philosophical position on the nature of subjective experience: a corpus contains records of what such workers produced; it does not contain the experiential signal that produced those records.

² β and γ are the respective shares of an occupation’s tasks where LLMs alone or LLMs with additional tooling can reduce time-to-task by at least 50%; the Felten indices weight O*NET ability requirements by AI’s overlap with each ability. The three exposure measures are highly correlated with each other ($r > 0.87$ pairwise).

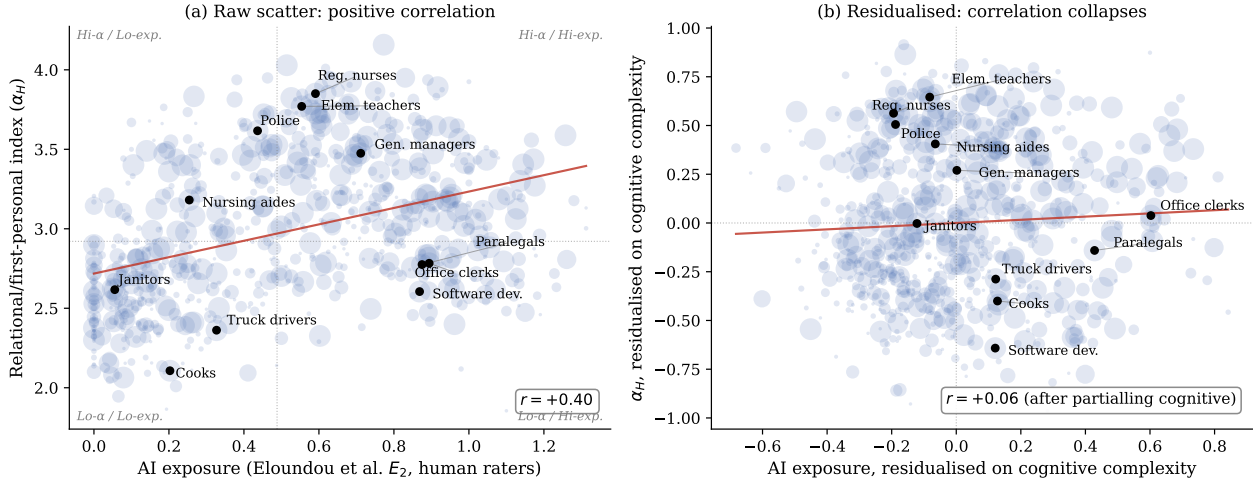


Figure 1: The relational dimension and AI exposure. Panel (a): raw scatter of α_H against the Eloundou et al. E_2 exposure score across $N = 746$ SOC-6 occupations. Bubbles sized by employment (BLS OEWS, May 2021); selected occupations labelled. Raw correlation $r = +0.40$. Panel (b): both axes residualised on a cognitive-complexity index built from fourteen O*NET items distinct from the relational items used in α_H . Conditional correlation $r = +0.06$. The positive raw correlation is fully accounted for by the joint loading of α_H and existing exposure measures on cognitive complexity.

cross-section ($r = 0.40, 0.48, 0.54$). Figure 1(a) shows the scatter against the Eloundou measure across 746 SOC-6 occupations.

The negative slope a one-dimensional substitution reading would predict is decisively contradicted by the data, robustly so across three independent exposure measures and two index constructions. The fact has a structural explanation rather than a contingent one. Existing exposure measures are constructed primarily from O*NET cognitive scales — Eloundou’s task ratings load on linguistic and analytical content; Felten’s AIOE on analytical and reasoning abilities — so they are best read as estimates of an occupation’s *cognitive intensity*. The natural hypothesis is that they correlate with α_H because professional occupations are typically dense in both cognitive and relational content. To test it, we construct a cognitive-complexity index from fourteen O*NET items capturing information processing, analysis, and abstract reasoning, distinct from the relational items in α_H (full list in Appendix A). It correlates 0.78 with log mean wage — it is the dimension that wage usually proxies for — and 0.65–0.74 with the three AI exposure measures, consistent with the interpretation that existing exposure measures are primarily capturing cognitive complexity. The cognitive-complexity index also correlates 0.55 with α_H , because professional oc-

cupations are both cognitively dense and relationally dense.

After partialling cognitive complexity out of both axes, the conditional association between α_H and AI exposure collapses sharply but unevenly across the three measures. The standardised conditional coefficient on α_H is +0.05 (Eloundou; not significant, $p = 0.12$), +0.10 (Felten general AIOE; $p = 0.002$), and +0.22 (Felten LM AIOE; $p < 0.001$); see Figure 1(b) and Table 1 for full inference. Both Felten residuals retain statistical significance — the LM-specific in particular reflects that some relational items (Persuasion, Negotiation, Instructing) are themselves language-mediated. The substantive claim is that the bulk of the raw positive correlation is accounted for by joint loading on cognitive complexity: R^2 rises from 0.16/0.23/0.30 in the univariate α_H specification to 0.43/0.56/0.53 when cognitive complexity is added. The relational dimension is therefore best described as conditionally near-orthogonal to general AI exposure measures and substantially attenuated for language-model-specific measures, rather than uniformly invisible.³

We also tested the Frey and Osborne (2017) probability-of-computerisation score against α_H and found a qualitatively different pattern: a *negative* raw correlation ($r = -0.58$) that does not collapse when cognitive complexity is partialled out (standardised conditional coefficient -0.32 , partial $r = -0.36$). This reflects the construction of the score: Frey and Osborne explicitly weighted social-intelligence variables (social perceptiveness, persuasion, negotiation, assisting and caring for others) as bottlenecks to computerisation. The post-2020 LLM-era exposure measures behave qualitatively differently from the earlier rubric-based score in which the relational dimension was incorporated by design.

The four-quadrant decomposition of Figure 1(a) makes the within-exposure heterogeneity visible. The high-exposure column splits sharply on α_H : registered nurses, elementary school teachers, general managers, and customer service representatives sit on the upper side; office clerks, secre-

³Two robustness facts qualify this. First, the Eloundou null is partly sample-dependent: on the common $N = 664$ Felten sample the conditional coefficient rises to +0.10 ($t = 2.95$, $p = 0.003$), so the +0.05 headline relies on the 82 narrow-SOC occupations the Felten files drop. Second, the residual is sensitive to the breadth of the cognitive index: a parsimonious 9-item cognitive index (excluding language-mediated items like Reading Comprehension and Writing) raises the Eloundou residual to $\approx +0.15$. We adopt the broad 14-item cognitive index because Eloundou’s task ratings themselves load on linguistic content; consistency between the cognitive control and the exposure measure justifies the inclusion. Robustness sweeps reported in Appendix A.

Table 1: α_H and AI exposure: raw and conditional on cognitive complexity

Specification	Outcome: AI exposure measure (z-scored)		
	Eloundou E_2	Felten AIOE	Felten LM AIOE
Raw	+0.395* (0.034)	+0.481* (0.034)	+0.543* (0.033)
+ log(wage)	+0.262* (0.034)	+0.324* (0.032)	+0.413* (0.032)
+ cognitive	+0.052 (0.033)	+0.097* (0.031)	+0.219* (0.032)
+ log(wage)+ cog.	+0.039 (0.033)	+0.089* (0.032)	+0.210* (0.033)
R^2 at base spec	0.156	0.231	0.295
R^2 +log(wage)	0.270	0.405	0.415
R^2 +cognitive	0.428	0.557	0.527
R^2 +log(wage)+cog.	0.435	0.558	0.528
N	746	664	664

Notes: Each cell reports the standardised coefficient on α_H in an OLS regression of the column outcome on α_H and the indicated controls; all variables z-scored before estimation. Standard errors in parentheses. * indicates $|t| > 1.96$. Sample is the intersection of O*NET 30.1, the indicated AI exposure measure, and BLS OEWS May 2021 national wages, at the SOC-6 occupation level. α_H is the simple average of importance ratings on ten O*NET items (five Work Activities: Assisting and Caring for Others; Establishing/Maintaining Interpersonal Relationships; Coaching and Developing Others; Resolving Conflicts and Negotiating; Performing for/Working Directly with the Public; five Skills: Social Perceptiveness, Service Orientation, Instructing, Persuasion, Negotiation). Cognitive complexity index is the simple average of importance ratings on fourteen O*NET items (Critical Thinking, Complex Problem Solving, Mathematics, Reading Comprehension, Writing, Active Learning, Judgment and Decision Making, Systems Analysis, Systems Evaluation, Analyzing Data or Information, Processing Information, Making Decisions and Solving Problems, Updating and Using Relevant Knowledge, Thinking Creatively). Eloundou $E_2 = \beta + 0.5\gamma$ from human ratings in Eloundou et al. (2023). AIOE and LM AIOE from Felten et al. (2021) and the language-modeling extension. Sample sizes differ across outcomes because the Felten files exclude some narrow SOC categories.

taries, bookkeepers, and paralegals sit on the lower side. Existing one-dimensional measures place the two groups at the same point on the exposure axis.

3. Model Setup: Production and Skill Investment

What then are the implications of AI for human capital investment? We model humans and AI as learning agents with the same skill-production technology; they differ only in what experience they can draw on and the amount of compute. Humans have a life and accumulate human-bundle (H) skill from it; AI has a corpus and cannot. Tasks combine distributional (D) and human-bundle

(H) skill inputs in proportions and substitutabilities that vary across tasks. Production runs one agent per task, with each task going to the lower-cost producer. Humans allocate finite lifetime investment between the two skill bundles.

3.1. Learning agents: the symmetric skill technology

We use the following notation. For each agent $a \in \{AI, H\}$ and each bundle $j \in \{D, H\}$, X_j^a is the agent's cumulative investment in (or data endowment on) the j -bundle, and s_j^a is the resulting skill level. The superscript indexes the agent; the subscript indexes the bundle.

AI's skill on the D-bundle depends on training data and compute $C = S/p_s$, where $S > 0$ is a scale constant and $p_s > 0$ is the price of compute:

$$s_D^{AI} = \bar{s}_D^{AI} \cdot \left(1 - e^{-\gamma_D^{AI} X_D^{AI} C}\right). \quad (1)$$

The parameter $\bar{s}_D^{AI} > 0$ is the asymptotic skill ceiling; $\gamma_D^{AI} > 0$ governs the elasticity to data and compute jointly; X_D^{AI} is large and taken as a fixed endowment. As $p_s \rightarrow 0$, $C \rightarrow \infty$ and $s_D^{AI} \rightarrow \bar{s}_D^{AI}$.

AI's skill on the H-bundle is identically zero by the structural zero:

$$s_H^{AI} \equiv 0 \quad \text{for all } C, \text{ because } X_H^{AI} \equiv 0. \quad (2)$$

This is the only primitive asymmetry between humans and AI. The claim embedded in $X_H^{AI} \equiv 0$ is that AI's H-bundle data is not small but absent: a corpus can contain records of what first-personal capacities produce, but not the experiential signal that produces them.

Human skill depends only on accumulated investment; biological compute is fixed and absorbed into the parameters:

$$s_D^H = \bar{s}_D^H \cdot \left(1 - e^{-\gamma_D^H X_D^H}\right), \quad s_H^H = \bar{s}_H^H \cdot \left(1 - e^{-\gamma_H^H X_H^H}\right). \quad (3)$$

Both are concave increasing with $s_j^H(0) = 0$ and $s_j^H \rightarrow \bar{s}_j^H$ as $X_j^H \rightarrow \infty$.

The Ben-Porath budget (Ben-Porath, 1967). Humans face

$$X_D^H + X_H^H \leq \bar{X}(T), \quad (4)$$

where T is the lifetime investment horizon. Both components are investment choices of the representative agent; there is no inherited constitutive stock in the baseline model.

The functional form in (1)–(3) is motivated by, but not a literal parameterization of, the neural scaling laws literature (Kaplan et al., 2020; Hoffmann et al., 2022): it captures the key properties that performance improves with data, that returns are diminishing, and that a finite investment horizon bounds human skill.

3.2. Assignment: task indivisibility and one-agent-per-task

Each task i is produced entirely by a single agent $a \in \{H, AI\}$, not jointly. We take this indivisibility as primitive: an AI that drafts notes for the physician and a physician who sees the patient are doing two tasks, not collaborating on one. Cleanly separable subtasks are modelled as multiple tasks rather than as coproduction within one task.

Given the indivisibility, assignment is cost-efficient: tasks go to whichever agent produces them more cheaply at the current wages and prices. This is the task-content tradition of Autor, Levy, and Murnane (2003) and Acemoglu and Restrepo (2018, 2019). What differs here is the structural zero, which on \mathcal{C} -mixed tasks (defined below) makes AI’s output identically zero regardless of cost, so no cost comparison is ever close. The assignment rule operates over the complement \mathcal{H}^{*c} .

AI per-task cost. We model AI as a competitive provider with zero marginal cost at the task level. A trained model serves an unbounded number of simultaneous tasks at compute cost $p_s \cdot c_i$ per task, where $c_i > 0$ is the per-task compute requirement. This is the standard platform assumption of fixed training cost and near-zero marginal serving cost; it is consistent with the observed deployment economics of frontier deep-learning systems and implies that p_s — the price

of compute — is the macro-parameter that varies as AI capability expands.

Human per-task cost. The human supplies both bundles on a task through the same embodied effort; the associated per-task cost is $w(i) = w_D s_D^H + w_H s_H^H$, where (w_D, w_H) are the equilibrium skill prices developed in Section 3.4.

Given these costs, assignment on contestable tasks is cost-efficient: the agent who delivers output per dollar at a higher rate wins the task. We write the formal assignment condition in Section 3.3 once the task production function is introduced.

3.3. Complementarities: within-task production and the structural floor

A task i is characterized by two parameters that do separate work. The first, $\alpha_{iH} \in [0, 1]$, is the H -bundle weight — how much the task requires human-exclusive first-personal skill. A routine legal summary has $\alpha_{iH} \approx 0$; end-of-life accompaniment has $\alpha_{iH} \approx 1$. The second, σ_i , is the elasticity of substitution between D - and H -bundle inputs within the task.

When $\sigma_i < 1$, the two inputs are complements within the task: more distributional skill on the task raises the marginal value of first-personal skill applied to the same task. A physician delivering a difficult cancer diagnosis needs both medical knowledge — to be precise about what is happening — and the first-personal capacity to read the patient and respond; either one without the other is far weaker than both together. When $\sigma_i > 1$, the two inputs are substitutes within the task: more distributional skill reduces the marginal value of first-personal skill on that task. A customer-service agent who can answer the question directly does not need much empathy to satisfy the customer; the first-personal layer is dispensable when the informational answer is good enough. The two parameters are orthogonal: a task can have any combination of H -weight and within-task elasticity.

Formally, we partition the task space $[0, 1]$ into two types based on σ_i :

$$\mathcal{C} = \{i : \sigma_i = \sigma^C < 1\}, \quad \mathcal{S} = \{i : \sigma_i = \sigma^S > 1\}, \quad (5)$$

with $\mu_C \equiv \mu(\mathcal{C}) \in (0, 1)$ and $\mu_S = 1 - \mu_C$.

Task production. Task output combines D -bundle and H -bundle skill via a within-task CES with elasticity σ_i :

$$q_i = \left[\alpha_{iD} \cdot s_D^{\frac{\sigma_i-1}{\sigma_i}} + \alpha_{iH} \cdot s_H^{\frac{\sigma_i-1}{\sigma_i}} \right]^{\frac{\sigma_i}{\sigma_i-1}}, \quad \alpha_{iD} + \alpha_{iH} = 1, \quad (6)$$

where s_D and s_H are the D - and H -bundle skills supplied to task i by whichever agent is assigned to it. The exponent $1/\sigma_i$ on the within-agent marginal products

$$\frac{\partial q_i}{\partial s_D} = \alpha_{iD} \cdot s_D^{-1/\sigma_i} \cdot q_i^{1/\sigma_i}, \quad \frac{\partial q_i}{\partial s_H} = \alpha_{iH} \cdot s_H^{-1/\sigma_i} \cdot q_i^{1/\sigma_i} \quad (7)$$

governs the within-agent cross-partial: when s_D rises on a task, the H -skill marginal product rises if $\sigma_i < 1$ (within-task complementarity) and falls if $\sigma_i > 1$ (within-task substitutability). The sign of $\sigma_i - 1$ is the source of the non-monotone investment result of Section 5.1.

The structural zero, plugged in. Substituting $s_H^{AI} = 0$ into (6) and applying the partition (5) yields three cases:

- **Type- \mathcal{C} mixed task** ($\sigma^C < 1$, $\alpha_{iH} \in (0, 1)$): $q_i^{AI} = 0$. AI cannot produce these tasks at any compute price — the within-task Leontief structure requires both inputs to be positive.
- **Type- \mathcal{S} mixed task** ($\sigma^S > 1$, $\alpha_{iH} \in (0, 1)$): $q_i^{AI} = \alpha_{iD}^{\sigma^S/(\sigma^S-1)} \cdot s_D^{AI}$. AI can produce at a bundle-compressed level, and wins against humans at sufficiently low p_s .
- **Pure \mathcal{D} -task** ($\alpha_{iH} = 0$): $q_i^{AI} = s_D^{AI}$; standard cost-efficiency applies.

We develop only these three cases; the corner $\alpha_{iH} = 1$ (purely first-personal work) is treated as a limit of the \mathcal{C} -mixed case. Illustrative examples of purely first-personal tasks — grief accompaniment, end-of-life presence — should be read as \mathcal{C} -mixed tasks with α_{iH} near the upper boundary. Analytically, the $\alpha_{iH} = 1$ corner adds no qualitative content that the interior \mathcal{C} -mixed case does not already supply: $q_i^{AI} = 0$ in both, for the same Leontief-corner reason.

The *structural floor* is therefore

$$\mathcal{H}^* \equiv \{i : \sigma_i = \sigma^C, \alpha_{iH} \in (0, 1)\}, \quad (8)$$

the set of \mathcal{C} -mixed tasks, on which AI cannot produce at any p_s . Its measure is invariant to compute prices. The complement \mathcal{H}^{*c} comprises pure-D tasks and \mathcal{S} -type mixed tasks, both of which AI can produce. The distinctive feature of this floor relative to the baseline $\{i : \alpha_{iH} > 0\}$ of simpler models is that \mathcal{S} -type mixed tasks have positive H-weight but are *not* permanently human: AI's positive-exponent CES lets it compress the missing H-bundle into a bundle-weighted form of the D-bundle.

Cost-efficient assignment on contestable tasks. On $i \in \mathcal{H}^{*c}$, agents are ranked by output per dollar:

$$i \in \mathcal{A}(p_s) \iff \frac{q_i^{AI}(p_s)}{p_s \cdot c_i} > \frac{q_i^H}{w(i)}, \quad (9)$$

where $w(i) = w_D s_D^H + w_H s_H^H$ evaluated at the human's equilibrium skill vector. As p_s falls, the left-hand side rises weakly, so $\mathcal{A}(p_s)$ expands monotonically over \mathcal{H}^{*c} . We establish in Proposition C.1 that $\mathcal{A}(p_s) \rightarrow \mathcal{H}^{*c}$ as $p_s \rightarrow 0$: AI eventually takes every contestable task, and only \mathcal{H}^* remains human.

Within-human complementarity. A consequence of one-agent-per-task is that on \mathcal{C} -mixed tasks — which are permanently human — the $\sigma^C < 1$ elasticity binds *within* the same human agent's own s_D^H and s_H^H , not across agents. More s_H^H raises the marginal product of same-agent s_D^H , and vice versa. This within-human complementarity is the engine of the non-monotone aggregate-investment result in Section 5.1.

3.4. Workers, wages, and the allocation problem

The representative agent's labor income is the bilinear aggregate

$$W(X_D^H, X_H^H) = w_D \cdot s_D^H(X_D^H) + w_H \cdot s_H^H(X_H^H), \quad (10)$$

where w_D and w_H are the equilibrium prices of one unit of D -skill and one unit of H -skill, each integrated across the task distribution through the task price schedule $\{p_i\}$.

The allocation problem. Because wages enter the objective linearly in skills (10) and skills are concave functions of investment (3), the lifetime investment problem collapses to a static allocation of the Ben-Porath budget across bundles: the agent maximizes (10) over (X_D^H, X_H^H) subject to (4). The interior first-order condition equalizes marginal returns across bundles,

$$w_D \cdot (s_D^H)'(X_D^{H*}) = w_H \cdot (s_H^H)'(X_H^{H*}), \quad (11)$$

so that X_D^{H*} moves with w_D and X_H^{H*} moves with w_H : investment mirrors wages. Under stationary wages the intratemporal FOC (11) coincides with the interior Euler condition of the full dynamic Ben-Porath problem; we work throughout with the static FOC, and none of the results below depend on the intertemporal structure.

Equilibrium and comparative statics in p_s . The structural results in Section 4 hold pointwise at any equilibrium wage vector (w_D, w_H) and compute price p_s . The transition results in Section 5 characterize how the equilibrium responds as p_s falls: $\{p_i\}$ are determined by the aggregate CES aggregator (16), and (w_D, w_H) solve labor-market clearing at each p_s . The equilibrium is stated formally as a fixed point in (w_D, w_H, Y) in Appendix B; its existence is established via Kakutani's theorem in Appendix C. Regularity conditions on the task distribution, functional forms, market structure, and equilibrium selection — standard assumptions A1–A5 invoked throughout the proofs — are collected in Appendix B.

4. The Human–AI Division of Labor

4.1. The performance ceiling: where AI cannot reach

AI faces a hard performance ceiling on any task requiring the H-bundle. More compute does not raise this ceiling because the bottleneck is missing data, not insufficient processing.

Proposition 1 (Ceiling structure). (i) *Component-level zero.* AI's H -skill is zero at every compute level.

(ii) *Task-level ceiling.* AI output is zero on C -mixed tasks (the Leontief corner), bundle-compressed on S -mixed tasks, and equal to full D -skill on pure- D tasks. Falling p_s raises the second two ceilings without bound but leaves the first untouched.

Part (i) is the structural zero: $X_H^{AI} \equiv 0$ plugged into the skill function (1) gives $s_H^{AI} = 0$ for any C . For part (ii), substituting $s_H^{AI} = 0$ into the within-task CES (6) yields

$$q_i^{AI} = \begin{cases} 0 & \text{on } C\text{-mixed tasks (Leontief corner),} \\ \alpha_{iD}^{\sigma^S/(\sigma^S-1)} \cdot s_D^{AI} & \text{on } S\text{-mixed tasks (bundle-compressed } D), \\ s_D^{AI} & \text{on pure-}D \text{ tasks.} \end{cases}$$

Compute acts only on s_D^{AI} , so the second two ceilings rise with compute and the first does not. The structural floor \mathcal{H}^* is a set of tasks placed permanently outside AI's reach by structure, not by cost.

4.2. Comparative advantage: who does what, and why

Under one-agent-per-task assignment, every contestable task goes to whichever agent produces it most cost-efficiently. Combining the cost-efficiency criterion with the ceiling structure gives a three-case sorting.

Proposition 2 (Comparative advantage). *Under one-agent-per-task assignment:*

- (i) *Pure- D tasks go to AI once p_s is small enough.*
- (ii) *S -mixed tasks also go to AI, but more slowly: bundle compression delays entry without preventing it.*
- (iii) *C -mixed tasks remain with humans at every p_s , because AI's ceiling is zero there.*

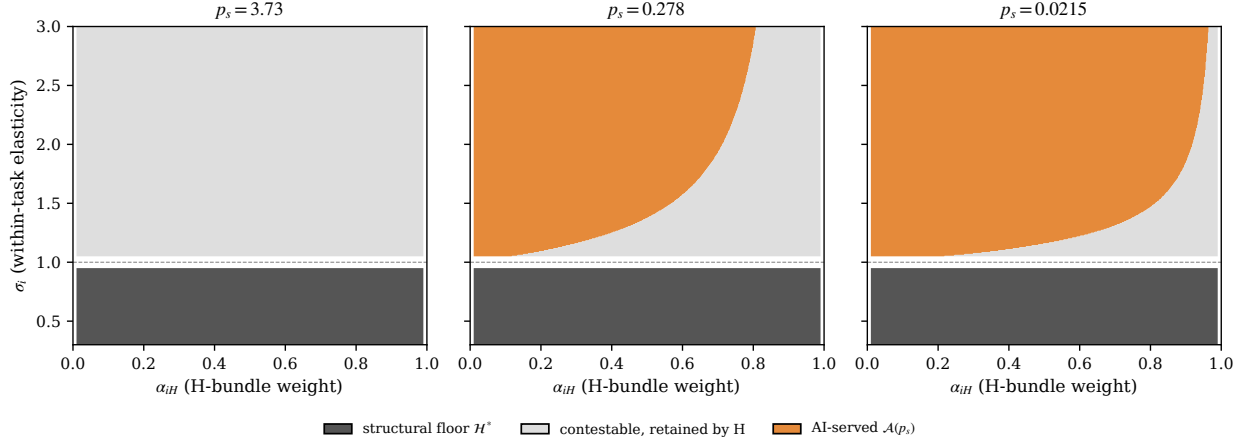


Figure 2: Task assignment across the (α_{iH}, σ_i) space at three levels of the compute price p_s . Dark shading: structural floor \mathcal{H}^* (human, permanent). Light shading: contestable tasks retained by humans at current p_s . Orange: AI-served tasks $\mathcal{A}(p_s)$. The boundary curve shifts right as p_s falls, but the $\sigma_i < 1$ band with $\alpha_{iH} > 0$ is invariant. Numerical GE: at each p_s , the equilibrium wages and skills $(w_D, w_H, s_D^H, s_H^H, s_D^{AI})$ are solved from the full model, then a dense (α, σ) grid is shaded by cost-efficiency (9). Dashed line marks the excluded $\sigma = 1$ locus. Parameters as in Figure 3.

Cases (i) and (ii) comprise the contestable set \mathcal{H}^{*c} that AI fully absorbs as $p_s \rightarrow 0$ (Proposition C.1); case (iii) is the structural floor \mathcal{H}^* .

4.3. The jagged frontier

Corollary 1 (The jagged frontier). *Define task difficulty operationally as the expected completion time or error rate for a trained non-expert human. So defined, difficulty is an unreliable predictor of AI performance. The reliable predictor is α_{iH} : tasks with $\alpha_{iH} \approx 0$ have high AI ceilings regardless of difficulty; tasks with $\alpha_{iH} > 0$ face a binding ceiling below human performance regardless of apparent simplicity.*

The frontier is jagged because α_{iH} is orthogonal to difficulty. Complex legal research has high α_{iD} and $\alpha_{iH} \approx 0$: hard but fully within the frontier. Sensing that a patient’s “I’m fine” is not fine has low apparent complexity but high α_{iH} : simple but outside the frontier. The jaggedness tracks the α_{iH} -weight, not difficulty.

Figure 2 plots the assignment regions across (α_{iH}, σ_i) at three values of p_s .

4.4. The automation boundary is also a measurement boundary

Remark 1 (Common boundary). The tasks most resistant to AI substitution — those with $\alpha_{iH} > 0$ — are also the tasks least accessible to standard benchmark measurement: the automation frontier and the measurement frontier share the same boundary.

A benchmark requires a pre-specifiable correct output evaluable without the evaluator participating in the interaction. For high- α_{iH} tasks this is violated for the same structural reason that $X_H^{AI} \equiv 0$: the H -bundle is constituted from inside the interaction. Existing AI productivity studies — Dell’Acqua et al. (2023), Noy and Zhang (2023), Brynjolfsson, Li, and Raymond (2023) — by design measure performance on tasks with pre-specifiable outputs and therefore sample from inside the AI frontier; the 15–40% gains they report are best read as gains on the D -bundle layer of professional work, not on professional work as a whole.

5. The Transition Path

Section 5.1 derives the long-run non-monotone return to D -skill; Section 5.2 derives the unambiguous rise in w_H via an aggregate Baumol channel; Section 5.3 closes the wage analysis with the short-run partial equilibrium at fixed pre-AI skill stocks — the comparative statics that map to the early empirical record; Section 5.4 records a separate within- D specialization margin governed by a distinct elasticity ρ .

5.1. Non-monotone D-investment

We work throughout this subsection in the main regime $\sigma^C < \eta$, in which AI cheapening generates a non-monotone D -wage; the complementary regime $\sigma^C \geq \eta$ is documented in Remark 2. As AI improves, the return to a worker’s own D -skill is the sum of two opposing forces, formalized below.

Within-task complementarity on \mathcal{C} -tasks. Corollary 2 (proved below) gives a strict endpoint

rise $w_H^\infty > w_H^{\text{pre-AI}}$, so humans invest more in s_H^H at the post-AI endpoint. Because \mathcal{C} -tasks have $\sigma^C < 1$, the same-agent cross-partial is

$$\frac{\partial^2 q_i}{\partial s_D^H \partial s_H^H} = \frac{1}{\sigma^C} \cdot \alpha_{iD} \alpha_{iH} (s_D^H s_H^H)^{-1/\sigma^C} q_i^{(2-\sigma^C)/\sigma^C} > 0, \quad (12)$$

strictly positive because $\sigma^C < 1$ and all factors are positive on \mathcal{C} -mixed tasks. Let $R_C(p_s)$ denote the measure-weighted return to s_D^H over \mathcal{C} -mixed tasks at the current equilibrium $s_H^H(p_s)$. Within-task complementarity gives $\dot{R}_C > 0$ as p_s falls: higher $w_H \Rightarrow$ higher $s_H^H \Rightarrow$ higher marginal return to same-agent s_D^H on \mathcal{C} -tasks via (12).

\mathcal{S} -task reassignment. On \mathcal{S} -mixed and pure-D tasks, AI becomes cost-efficient as p_s falls (Proposition 2(i)–(ii)). The measure of contestable tasks still served by humans, $\mu^{\text{rem}}(p_s) \equiv \mu(\mathcal{H}^{*c} \setminus \mathcal{A}(p_s))$, declines weakly and reaches zero as $p_s \rightarrow 0$. Let $R_S(p_s)$ denote the human D-return aggregated over this shrinking set. Reassignment gives $\dot{R}_S < 0$: the base of tasks over which s_D^H earns in the contestable set contracts.

Aggregate return and threshold. The total return to human s_D^H is

$$R(p_s) = \mu_C \cdot R_C(p_s) + \Pi_S(p_s), \quad \Pi_S(p_s) \equiv \mu^{\text{rem}}(p_s) \cdot R_S(p_s), \quad (13)$$

where Π_S is the aggregate (measure-weighted) \mathcal{S} -return and $\tau \equiv -\log p_s$ so “falling p_s ” corresponds to τ rising. The \mathcal{S} -reassignment force is conveniently summarised by the rate at which Π_S decays:

$$|\dot{R}_S|(p_s) \equiv -\frac{d\Pi_S}{d\tau} = -\dot{\mu}^{\text{rem}} \cdot R_S - \mu^{\text{rem}} \cdot \dot{R}_S \geq 0, \quad (14)$$

nonnegative asymptotically: $\Pi_S \rightarrow 0$ as AI absorbs \mathcal{H}^{*c} in measure (Proposition C.1(ii)). Pointwise monotonicity along the path is not invoked; Lemma C.1’s tail-bound construction uses only the asymptotic statement. The derivative of R with respect to τ is then

$$\dot{R} = \mu_C \dot{R}_C - |\dot{R}_S|, \quad (15)$$

with $\dot{R}_C > 0$ from within-task complementarity (when $\sigma^C < \eta$) and $|\dot{R}_S| \geq 0$ from \mathcal{S} -task reassignment contracting the contestable base.

Proposition 3 (Non-monotone D-investment). *Assume $\sigma^C < \eta$ (within-task complementarity operative; see Remark 2 otherwise). Let $\mu_C^*(p_s) \equiv |\dot{R}_S(p_s)|/\dot{R}_C(p_s)$, the instantaneous threshold defined via (14). As p_s falls:*

(i) *If $\mu_C^*(p_s) < \mu_C$ (within-task complementarity dominates \mathcal{S} -task reassignment), $\dot{w}_D > 0$ and $\dot{X}_D^{H*} > 0$.*

(ii) *If $\mu_C^*(p_s) > \mu_C$ (\mathcal{S} -task reassignment dominates), $\dot{w}_D < 0$ and $\dot{X}_D^{H*} < 0$.*

(iii) (Regime switch and U-shape.) *$\mu_C^*(p_s) \rightarrow 0$ as $p_s \rightarrow 0$, so if regime (ii) holds initially — $\mu_C^*(p_s^{pre-AI}) > \mu_C$ — then w_D and X_D^{H*} are non-monotone in $\tau \equiv -\log p_s$ with at least one regime switch from falling to rising. Under the additional condition that $\mu_C^*(p_s)$ is monotone in τ on $(0, p_s^{pre-AI}]$, the switch is unique and the path is strictly U-shaped.*

Proof in Appendix C; sufficient conditions for the single-crossing hypothesis in part (iii) are in Remark C.1 of the appendix, and the numerical equilibrium of Section 5.2 satisfies them.

Remark 2 (The $\sigma^C \geq \eta$ regime). When $\sigma^C \geq \eta$, within-task complementarity flips sign: a rise in s_H^H no longer raises the return to same-agent s_D^H on \mathcal{C} -tasks. \mathcal{S} -task reassignment is then unopposed, w_D is weakly decreasing, and the regime switch does not obtain.

Figure 3 traces the allocation of the Ben-Porath budget across the two bundles as p_s falls under the parameterized equilibrium. X_D^{H*} falls during the early \mathcal{S} -dominated phase and rises during the subsequent \mathcal{C} -dominated phase, tracing the strict U-shape of Proposition 3(iii); the single-crossing hypothesis on $\mu_C^*(\tau)$ holds in this parameterization. The budget constraint $X_D^{H*} + X_H^{H*} = \bar{X}$ forces X_H^{H*} into the mirror-image hump.

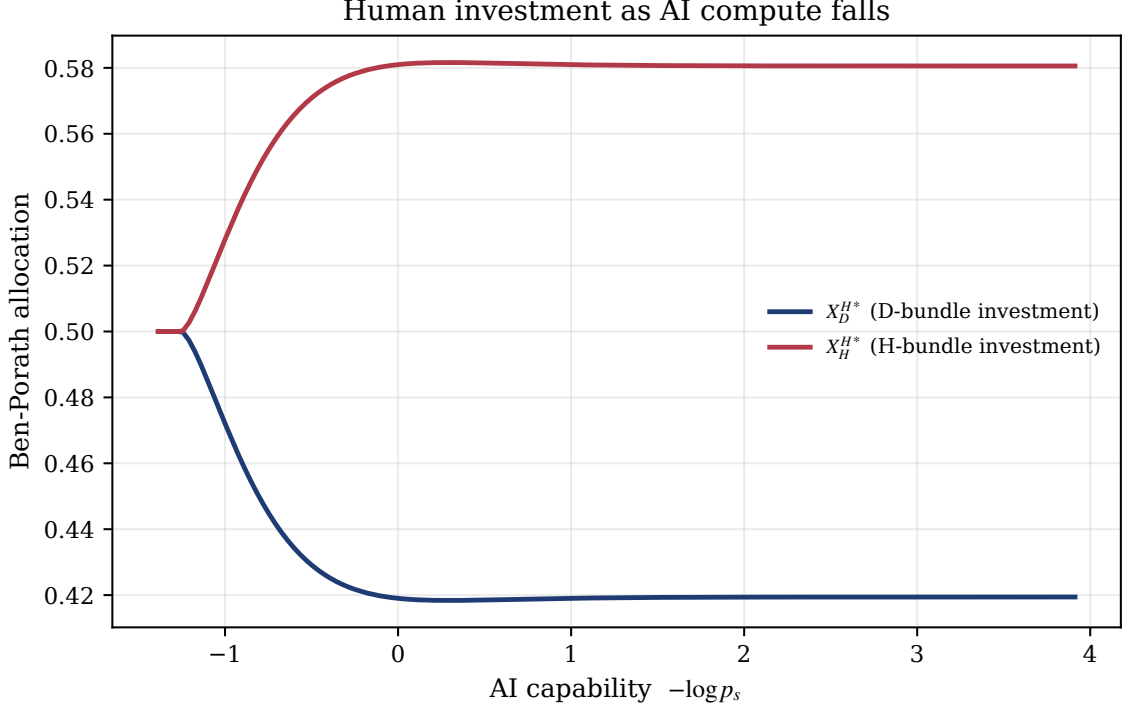


Figure 3: Optimal human investment allocation as p_s falls; $X_H^{H*} = \bar{X} - X_D^{H*}$ is the mirror-image hump. Horizontal axis is $\tau = -\log p_s$. Numerical GE solution; parameters in note.⁴

5.2. Wages and the aggregate Baumol channel

Wages inherit the forces of Proposition 3. The D -wage tracks its regime, and the relative price w_D/w_H is U-shaped — falling early as \mathcal{S} -task reassignment pulls w_D down while w_H is pulled up, recovering later as within-task complementarity reasserts itself on \mathcal{C} -tasks. Equivalently, the wage gap w_H/w_D is inverted-U over the transition (widens then narrows). The H -wage itself rises unambiguously between pre-AI and post-AI endpoints.

Corollary 2 (w_H unambiguous). *Under A4 ($\eta < 1$), as $p_s \rightarrow 0$:*

- (a) *The outer-CES expenditure share on \mathcal{H}^* converges to one while human H -skill supply is bounded, so w_H converges to a finite positive limit w_H^∞ .*
- (b) *Under a density-dominance condition on the joint task distribution plus a path-integrated bound on the cost-share swing (Appendix C), $w_H^\infty > w_H^{\text{pre-AI}}$ strictly. The benchmark parameterization (Section 6) satisfies these conditions analytically or by numerical verification*

Equilibrium skill prices as AI compute falls

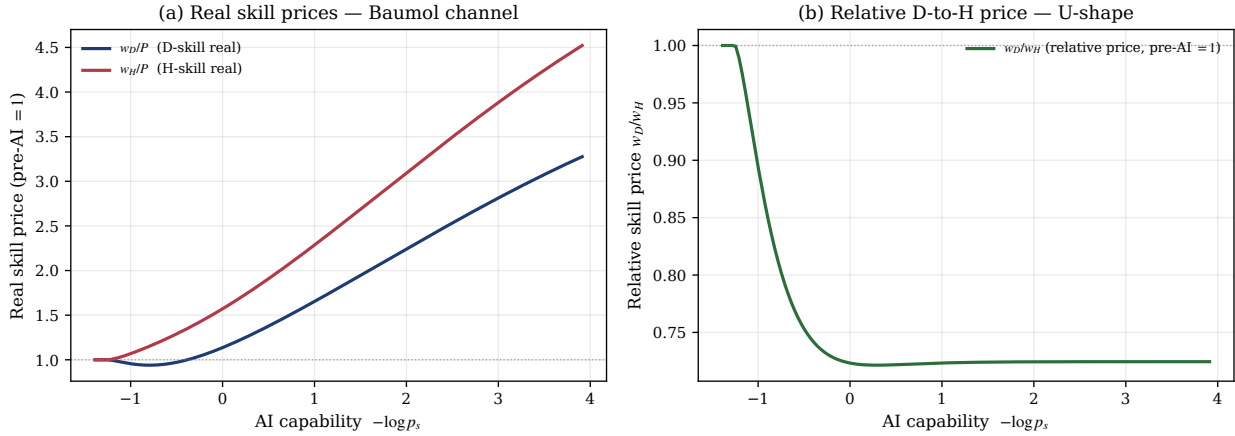


Figure 4: Equilibrium skill prices as p_s falls. Panel (a): $w_D/P, w_H/P$ normalised to pre-AI = 1. Panel (b): relative price w_D/w_H . Horizontal axis $\tau = -\log p_s$; parameters as in Figure 3.

(Table 2).

Within-human complementarity on C-tasks reinforces the channel in case (i) of Proposition 3 and does not overturn it in case (ii).

As $p_s \rightarrow 0$, s_D^{AI} saturates at its ceiling and the cost per unit of D -bundle output approaches zero, while the H -bundle is produced by humans subject to inelastic labor supply — one human attends to one interaction at a time. Under $\eta < 1$, demand for the more expensive composite is price-inelastic. The expenditure share on \mathcal{H}^* therefore rises toward one. A growing share of total expenditure chases a fixed supply of H -skill, bidding w_H up. This extends Baumol and Bowen (1966)’s cost disease from the sectoral to the factor level: the rising relative price of H -work reflects scarcity of an input with bounded supply.

Figure 4 traces the two skill prices. Panel (a): real skill prices $w_D/P, w_H/P$ both rise through the Baumol channel; w_H/P rises monotonically while w_D/P briefly dips before recovering above its pre-AI level. Panel (b): the U-shape in w_D/w_H described above.

The sectoral implication follows: relative prices of H -intensive components in healthcare, education, therapy, eldercare, and long-term care will rise faster than general wages as AI automates the documentation, diagnosis, content-delivery, and administrative layers of those sectors. The pattern is already in the data: real prices of personal services and education have grown roughly

0.6–1.3 percentage points per year faster than the price of consumer goods over the post-1950 U.S. record (Baumol, 2012, ch. 2, Tables 2.1–2.3), and an analogous wedge is documented for U.S. health care over 1960–2014 (Helland and Tabarrok, 2019, Fig. 1). The framework predicts this divergence will accelerate, not close, as AI diffuses.⁵

5.3. Short-run wage incidence at fixed skills

The long-run results of Subsections 5.1–5.2 all rely on one margin: workers reallocating their Ben-Porath budget toward s_H^H as w_H rises. Before that reallocation has occurred, the equilibrium response to falling p_s runs only through assignment, prices, and labour-market clearing at fixed skill stocks. We close the wage analysis by isolating this short-run partial equilibrium, both because it clarifies which long-run results require skill reinvestment to obtain and because it maps directly to the early empirical record.

Fix human skill stocks at their pre-AI equilibrium values $(s_D^{H,\text{pre-AI}}, s_H^{H,\text{pre-AI}})$, with $X_D^{H,\text{pre-AI}} + X_H^{H,\text{pre-AI}} = \bar{X}$ chosen optimally at $p_s^{\text{pre-AI}}$ via (11). The *short-run partial equilibrium at p_s* , denoted $(w_D^{SR}(p_s), w_H^{SR}(p_s))$, is the wage vector that clears labour markets at these fixed skills, with task assignment $\mathcal{A}(p_s)$ following cost-efficiency (17), task prices $\{p_i\}$ following the outer CES (16), and aggregate output Y adjusting at the resulting human skill supplies. The fixed-point argument of Theorem B.1 delivers existence: the worker-allocation step is replaced by an exogenous skill vector, and the remaining steps run unchanged.

Two forces act on the short-run wages. The \mathcal{S} -reassignment force $|\dot{\Pi}_S|$ in (14) operates identically to its long-run counterpart at the same skill point, since skills enter (17) on both sides of the contestability comparison. The \mathcal{C} -task force, by contrast, is reduced. The long-run object \dot{R}_C has two components: the Baumol price effect on p_i over \mathcal{C} -tasks, and the cross-partial amplification (12) through rising s_H^{H*} . The short-run version retains only the first.

⁵A demand-side force runs in the same direction: Comin, Lashkari, and Mestieri (2021) show that within-country sectoral reallocation is driven primarily by income effects, with spending shifting toward goods of income elasticity above one. If H -intensive goods — therapy, personalized education, eldercare, design consultation — have high income elasticity, AI-driven productivity growth raises demand for them independently of price effects.

Proposition 4 (Short-run wage incidence). *Under A1–A4 and the structural zero, with skills fixed at $(s_D^{H,pre-AI}, s_H^{H,pre-AI})$:*

- (i) (w_H rise is unconditional.) $w_H^{SR}(p_s)$ is weakly increasing in $\tau = -\log p_s$, with strict rise on any interval where $\mathcal{A}(p_s)$ is expanding. The endpoint comparison $w_H^{SR}(0^+) > w_H^{pre-AI}$ obtains under $\eta < 1$ alone, without the density-dominance condition required by Corollary 2.
- (ii) (No short-run U-shape in w_D .) The recovery of w_D in Proposition 3(iii) operates through the cross-partial (12), which is non-zero only when s_H^{H*} is rising; at fixed $s_H^{H,pre-AI}$ the recovery channel is silent. Whenever the long-run model is in regime (ii) of Proposition 3 at the pre-AI endpoint, $w_D^{SR}(p_s)$ is weakly decreasing on the early phase of the transition. The late-stage C-Baumol price effect can lift w_D^{SR} at the asymptote, but the U-shape of the long-run path collapses to a flatter trajectory amplified by no skill-reinvestment feedback.

Proof in Appendix C.

The Baumol channel of Corollary 2 actually strengthens at fixed $s_H^{H,pre-AI}$: with no offsetting supply expansion, the endpoint comparison holds without density-dominance. The U-shape of w_D , by contrast, requires the cross-partial through rising s_H^H — the channel switched off here.

The short-run prediction lines up with the early empirical record. Humlum and Vestergaard (2025) report essentially zero earnings effects for senior workers in AI-exposed Danish occupations through 2024 alongside contracting early-career hiring; Hampole et al. (2025) document within-firm labour-demand declines concentrated where workers’ portfolios skew toward AI-exposed task content, with sharper effects on the entry margin. Both patterns match the incidence of Proposition 4: workers whose pre-AI investments leaned toward D -bundle work on contestable tasks — disproportionately junior practitioners assigned to \mathcal{S} -mixed and pure- D tasks — bear the wage and employment incidence first, while incumbents already holding \mathcal{C} -task exposure ride the Baumol channel from the start. The recovery of w_D traced in the long-run analysis above requires a margin — skill reinvestment — that the cohort absorbing the brunt of the early transition cannot exercise instantly. The pipeline implication is taken up in Section 7.

5.4. A separate within- D specialization margin

A question distinct from the direction-of-investment results above: once a worker has decided how much D -skill to acquire, how should she compose it across D -sub-dimensions (front-end versus back-end, tax versus corporate, diagnostic versus procedural)? The relevant elasticity here is a within- D parameter ρ , which is logically distinct from σ_i . Whether AI pushes workers toward balance ($\rho < 1$) or specialization ($\rho > 1$) is a property of the within- D production technology, not pinned down by the model.⁶ The formal statement and proof are in Appendix D (Proposition D.1).

6. Task Composition and the Transition Path

Section 5 establishes the directional results; this section quantifies them. We trace the transition under a benchmark supported by existing evidence (Section 6.1), show that the directional predictions are robust to the parameters the theory does not pin down (Section 6.2), and characterise the conditions under which the long-run wage ordering can reverse (Section 6.3).

6.1. The benchmark

The benchmark makes two assumptions, each supported by existing evidence rather than by numerical convenience.

First, α_{iH} and σ_i are positively aligned: complement-type tasks draw α_{iH} from Beta(5, 2), concentrated at high α , and substitute-type tasks from Beta(2, 5), concentrated at low α . This alignment is the paper’s substantive claim about the task space, not a numerical convenience.

Sorting tasks along this alignment, the canonical substitutable work — routine information processing, basic coding, document drafting, summary writing, standardized forecasting — is both low- α and high- σ ; the canonical complement-type work — clinical presence, teaching, leadership,

⁶We treat ρ as a single parameter for tractability of the closed-form result; allowing ρ_i heterogeneity across tasks would preserve the qualitative statements of Proposition D.1 but lose the analytical comparison of corner versus interior portfolios.

therapy, long-term care — is both high- α and low- σ . The work AI has rapidly absorbed since the deployment of large language models is concentrated in the low- α , high- σ region (Dell’Acqua et al., 2023; Brynjolfsson et al., 2023; Noy and Zhang, 2023); we are not aware of comparable evidence on rapid absorption in the high- α , low- σ region.

Second, $\mu_C = 0.5$. We adopt this value because larger μ_C dampens the long-run H -premium, so a value at the conservative end of the plausible range serves the H-premium claim. O*NET relational-task shares and Dingel–Neiman irreducibility measures point toward μ_C in the 0.15–0.35 range at the task level; setting $\mu_C = 0.5$ understates the long-run H -premium relative to what those benchmarks would suggest.

Under these assumptions, the transition exhibits the qualitative shape derived in Section 5: a brief early dip in w_D/P as \mathcal{S} -task reassignment dominates, followed by a recovery as within-task complementarity on \mathcal{C} -tasks reasserts itself; a monotonic rise in w_H/P via the aggregate Baumol channel; and a long-run wage gap with w_H above w_D . Quantitatively, w_H/P rises to roughly $4.5\times$ its pre-AI level, w_D/P dips briefly to 0.94 of its pre-AI real level before recovering to roughly $3.3\times$, and the long-run wage ratio w_D/w_H settles at 0.72 (Figures 3 and 4). The optimal investment path X_D^{H*} traces the strict U-shape implied by Proposition 3(iii).

6.2. Robustness

Holding the benchmark’s positive α – σ alignment, every dial we examine moves magnitudes or timing but not the long-run sign of the predictions. We summarise the three main perturbations here; Appendix E reports the full sweeps and endpoint values.

The size of the permanently-human domain. Figure 5 varies $\mu_C \in \{0.25, 0.35, 0.50, 0.65\}$, holding shape and alignment at the benchmark. Endpoint w_H/P falls from roughly $15\times$ at $\mu_C = 0.25$ to roughly $2.6\times$ at $\mu_C = 0.65$ — a factor-of-six range, the dominant magnitude dial in the exercise. The $15\times$ endpoint at $\mu_C = 0.25$ is the upper extreme of a one-parameter sweep and should not be read as a precise prediction: pushing μ_C below the O*NET/Dingel–Neiman-supported range, holding all other dials fixed, mechanically inflates the endpoint via the inelastic-

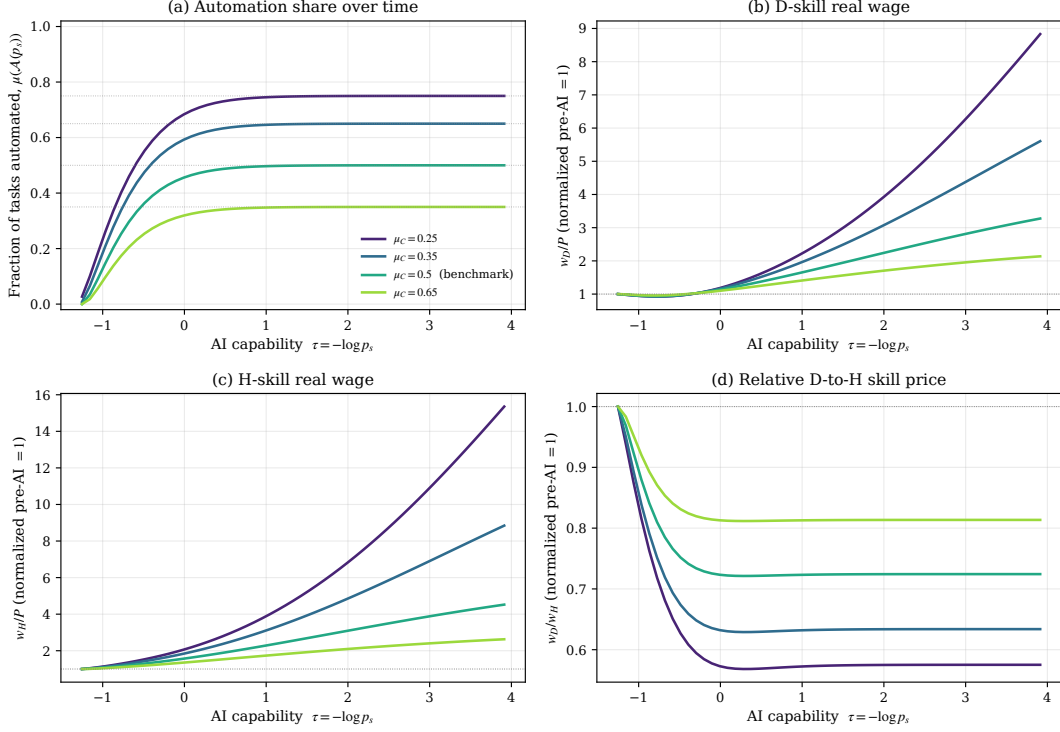


Figure 5: The size of the permanently-human domain. $\mu_C \in \{0.25, 0.35, 0.50, 0.65\}$, holding shape and alignment at the benchmark. Automation asymptote $1 - \mu_C$ (panel a); w_H/P endpoint scales with μ_C^{-1} (panel b); w_D/P remains non-monotone (panel c); wage ratio ordering preserved (panel d).

supply channel of Corollary 2. What is robust across the range is the qualitative pattern — a long-run H -premium and the early w_D/P dip; the level of the premium is the open empirical question the framework raises.

Concentration within regimes. Figure 6 varies the sharpness of f_C and f_S , holding μ_C and the alignment sign fixed. Sharper distributions push the automation boundary out faster at modest AI capability, because more of the \mathcal{S} -mass is concentrated at low α_{iH} where AI's comparative advantage is unambiguous; milder distributions stretch absorption over a longer range of p_s . The timing of the D -wage dip shifts correspondingly — shallower and earlier under sharper distributions, deeper and later under milder — but endpoint magnitudes differ only modestly and the ordering of w_D/w_H is preserved.

Non-compositional parameters. Five parameters are not features of the task composition: the outer CES elasticity η , the Ben-Porath curvature γ , the AI scale K^{AI} , the AI saturation \bar{s}_D^{AI} , and the inner within-task elasticities (σ^C, σ^S). Across all five, the directional predictions of Section 5

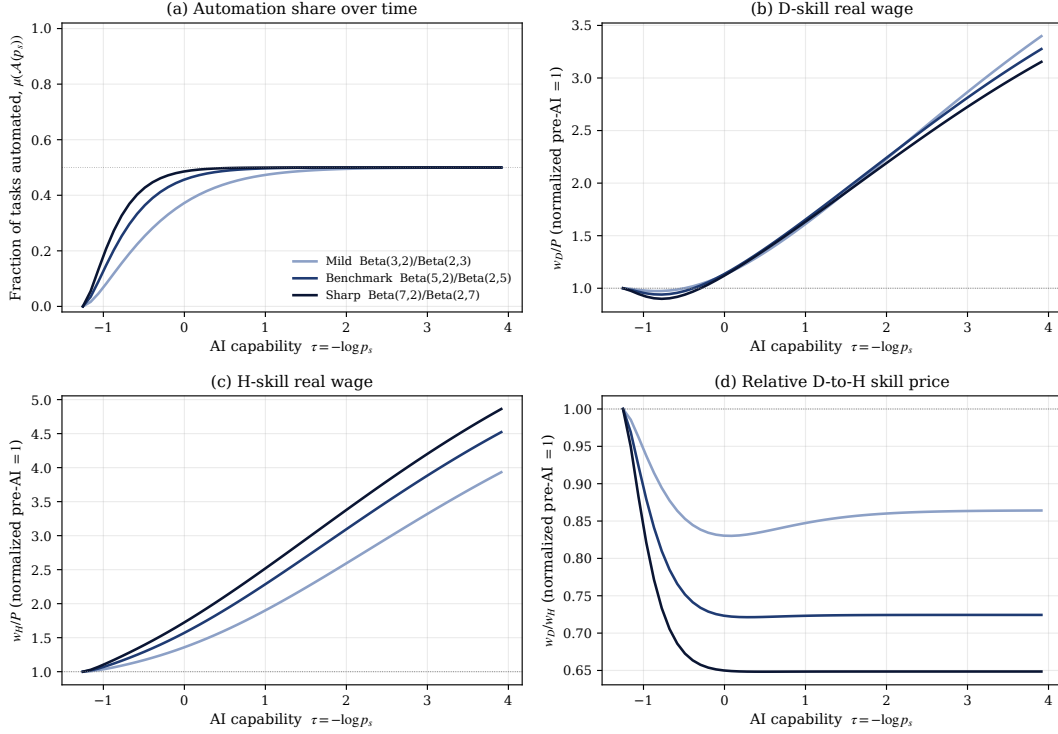


Figure 6: Concentration within regimes: sharpness of f_C and f_S (mild, benchmark, sharp) holding μ_C and the alignment sign fixed. Direction-preserving; sharper distributions give a faster early displacement wave and a shallower, earlier w_D/P dip.

retain their sign. Only η carries a non-trivial magnitude effect: endpoint w_H/P moves over roughly $[3.3, 6.8]$ across $\eta \in \{0.4, 0.6, 0.8\}$, with smaller η (stronger Baumol inelasticity) producing a larger long-run H -wage. The remaining four are second-order in the region where AI skill is close to saturation.

6.3. When the long-run wage ordering can reverse

The long-run sign of the wage gap is an empirical question, not a theoretical one: it depends on whether \mathcal{C} -tasks are first-personally intensive or distributionally intensive; both cases are consistent with the model's structure. Under positive α - σ alignment, the permanently-human domain is first-personally intensive and the aggregate Baumol channel bids up the price of first-personal skill. Under reversed alignment, the permanently-human domain is distributionally intensive and the Baumol channel bids up the price of D -skill instead.

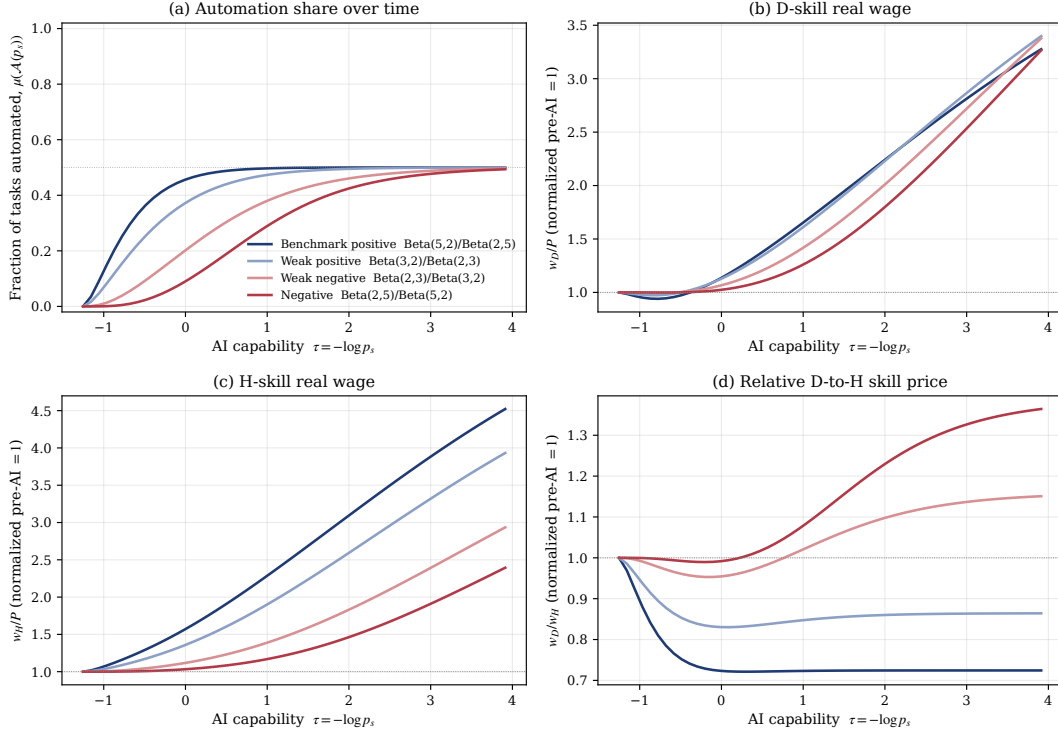


Figure 7: The alignment of first-personal weight and within-task complementarity. Four cases holding $\mu_C = 0.5$ fixed: the benchmark (H -premium, $w_D/w_H = 0.72$), weak positive (0.86), weak negative (1.15), and the full reversal (D -premium, $w_D/w_H = 1.36$).

Numerically, with $\mu_C = 0.5$ held fixed, the long-run w_D/w_H ranges from 0.72 under the benchmark Beta(5, 2)/Beta(2, 5) alignment to 1.36 under its full reversal; weaker versions (Beta(3, 2)/Beta(2, 3) and its reverse) bracket these endpoints at 0.86 and 1.15 (Figure 7). Under reversed alignment the D -wage rises essentially monotonically rather than dipping early. The structural floor itself is invariant: only μ_C governs its size.

We treat the positive-alignment case as the empirically relevant benchmark: AI has rapidly absorbed low- α , high- σ work while making slow progress on high- α , low- σ work. Under that alignment, w_D falls during the transition before recovering, with w_H above w_D in the long run. The reverse case is documented above for completeness.

7. Conclusion: Implications for Human Capital Formation

AI will not automate every task. The work that remains is the work whose generative inputs exist only first-hand — bodies, presences, and judgments that no external corpus contains. As AI cheapens, real wages rise; demand concentrates on what AI cannot do, while the supply of workers who can do it is fixed at one human per interaction. An aggregate Baumol channel raises the wage of first-personal skill, which persistently overtakes ordinary cognitive skill if human-exclusive tasks are sufficiently first-personal-skill intensive.

Two workers with identical resumes can land on opposite sides of the same shock. Where first-personal and cognitive inputs complement each other within the same human (a physician's medical knowledge paired with her capacity to read the patient), AI cannot perform the task at all; the Baumol channel raises its price as AI absorbs other work, and both skills rise in value. Where the first-personal layer is dispensable once the cognitive answer is good enough (a call-centre agent whose informational answer resolves the question), AI takes the whole task and the human is displaced. Standard AI exposure measures put the two kinds of work at the same point on the exposure axis: registered nurses, teachers, and managers next to office clerks, secretaries, and paralegals. Same nominal exposure, opposite incidence.

Implications for education

Modern education is organized around the efficient transmission of distributional knowledge: lectures convey codified content, examinations verify retention, credentials certify completion. This was the right organization for an economy in which distributional knowledge — codified facts, analytical techniques, established procedures — commanded substantial wage premiums. AI increasingly absorbs that layer of work. Experiential skill becomes the scarce factor — and current pedagogy is not built to develop it.

Distributional skill remains essential — as an input the permanent human domain itself requires — but its marginal value falls relative to first-personal skill, and educational investment

should follow. Pedagogy might shift from lecture-based transmission of codified content toward project-based learning, apprenticeships, and mentorship structures that develop judgment through repeated consequential action. Assessment might shift from standardized testing of distributional retention toward portfolio-based, performance-based, and long-form demonstrations of judgment in context. Credentialing might shift from certifying coursework completion toward recognizing experiential skill, so that practitioners whose value lies in accumulated first-personal engagement are credentialed for what they actually do.

The entry-level apprenticeship pipeline

The framework illuminates a current development whose long-run cost may be underappreciated: the contraction of the entry-level cognitive job, documented in U.S. data by Humlum and Vestergaard (2025). Junior consultants, paralegals, software developers, and analysts — the roles that have historically served as the apprenticeship rungs of professional careers — perform precisely the work AI is absorbing most rapidly. The standard reading is efficiency: the headcount that previously did this work was a labour cost, and AI substitution lowers it.

The framework here suggests a counter-implication. Entry-level cognitive work was where junior practitioners accumulated the experiential skill they would later deploy as seniors: the apprenticeship was a side-product of the work, not a detachable investment. If AI takes the entry rungs, the pipeline for senior experiential skill is choked off. The framework rationalises a worry that would otherwise look like nostalgia: structural interventions to preserve apprenticeship opportunities — protected on-ramps for junior workers, restructuring of senior roles around teaching and mentorship, floors on the human content of professional work — may be desirable not despite the productivity loss but because the long-run cost of a missing apprenticeship pipeline exceeds the short-run gain from skipping it.

References

- Daron Acemoglu. The simple macroeconomics of AI. *Economic Policy*, 40(121):13–58, 2025. doi: 10.1093/epolic/eiae042.
- 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.
- David H. Autor and Neil Thompson. Expertise. *Journal of the European Economic Association*, 23(4):1203–1271, 2025. doi: 10.1093/jeea/jvaf023.
- David H. Autor, Frank Levy, and Richard J. Murnane. The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, 118(4):1279–1333, 2003.
- William J. Baumol. *The Cost Disease: Why Computers Get Cheaper and Health Care Doesn't*. Yale University Press, 2012.
- William J. Baumol and William G. Bowen. *Performing Arts: The Economic Dilemma*. The Twentieth Century Fund, 1966.
- Yoram Ben-Porath. The production of human capital and the life cycle of earnings. *Journal of Political Economy*, 75(4):352–365, 1967.
- Erik Brynjolfsson, Danielle Li, and Lindsey R. Raymond. Generative AI at work. *NBER Working Paper No. 31161*, 2023.
- Diego Comin, Danial Lashkari, and Martí Mestieri. Structural change with long-run income and price effects. *Econometrica*, 89(1):311–374, 2021.

Fabrizio Dell’Acqua, Edward McFowland III, Ethan R. Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Kraymer, 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. *Harvard Business School Working Paper 24-013*, 2023.

David J. Deming. The growing importance of social skills in the labor market. *Quarterly Journal of Economics*, 132(4):1593–1640, 2017.

Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock. GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models. *arXiv preprint arXiv:2303.10130*, 2023.

Edward W. Felten, Manav Raj, and Robert Seamans. Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12):2195–2217, 2021.

Carl Benedikt Frey and Michael A. Osborne. The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114:254–280, 2017.

Luis Garicano. Hierarchies and the organization of knowledge in production. *Journal of Political Economy*, 108(5):874–904, 2000.

Luis Garicano and Luis Rayo. Training in the age of AI: A theory of apprenticeship viability. *Working Paper*, 2025.

Luis Garicano, Jin Li, and Yanhui Wu. Weak bundle, strong bundle: How AI redraws job boundaries. *CEPR Discussion Paper No. 21453*, 2026.

Menaka Hampole, Dimitris Papanikolaou, Lawrence D. W. Schmidt, and Bryan Seegmiller. Artificial intelligence and the labor market. *NBER Working Paper No. 33509*, 2025.

Eric Helland and Alex Tabarrok. Why are the prices so damn high? Health, education, and the baumol effect. *Mercatus Center Research Paper*, 2019.

Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*, 2022.

Anders Humlum and Emilie Vestergaard. Still waters, rapid currents: Early labor market transformation under generative AI. *NBER Working Paper No. 33777*, 2025.

Enrique Ide and Eduard Talamas. Artificial intelligence in the knowledge economy. *Journal of Political Economy*, 133(12), 2025. doi: 10.1086/737233.

Ingka Group. Annual Summary and Sustainability Report FY22. Ingka Group, Leiden, 2022. Reports remote interior design channel revenue of EUR 1.3 billion (3.3% of total revenue) in fiscal year 2022. Available at <https://www.ingka.com/reports/>.

Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.

Isabella Loaiza and Roberto Rigobon. The EPOCH of AI: Human-machine complementarities at work. *SSRN Working Paper No. 5028371*, 2024.

Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, 1995.

Shakked Noy and Whitney Zhang. Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654):187–192, 2023.

Helen Reid. IKEA bets on remote interior design as AI changes sales strategy. Reuters, June 13, 2023, 2023. Reports that Ingka Group’s AI chatbot Billie has handled approximately 47% of customer queries since 2021, and that 8,500 call centre workers globally were retrained as remote interior design advisers.

Catherine J. Weinberger. The increasing complementarity between cognitive and social skills.
Review of Economics and Statistics, 96(5):849–861, 2014.

A. Construction of the Empirical Indices

This appendix collects the construction details and validation diagnostics for the two O*NET indices used in Section 2.2.

Sample. O*NET database 30.1 (April 2025), Work Activities and Skills tables, importance ratings only. Merged at the SOC-6 occupation level with the Eloundou et al. (2023) occupation-level exposure file (human ratings) and the Felten et al. (2021) AIOE and LM AIOE files. Wages and employment from BLS OEWS national May 2021 (cross-occupation log-wage rankings are stable across years; this is a skill-level control rather than a level-of-wages object). Final intersection: 746 SOC-6 occupations against the Eloundou measure, 664 against the Felten measures (the latter exclude some narrow SOC categories). Replication code, raw data, and merged files are archived alongside the paper.

α_H **index.** Simple average of importance ratings (1–5 scale) on ten O*NET items chosen to capture interpersonal, care, coordination, and influence content of work:

- *Work Activities:* Assisting and Caring for Others; Establishing and Maintaining Interpersonal Relationships; Coaching and Developing Others; Resolving Conflicts and Negotiating with Others; Performing for or Working Directly with the Public.
- *Skills:* Social Perceptiveness; Service Orientation; Instructing; Persuasion; Negotiation.

PC1 over the standardised items explains 60.1% of variance with all loadings positive in [0.24, 0.37], led by Social Perceptiveness, Service Orientation, Persuasion, and Negotiation. Results in Section 2.2 are nearly identical with PC1 in place of the simple average.

Cognitive-complexity index. Simple average of importance ratings on fourteen O*NET items chosen to load on information processing, analysis, and abstract reasoning while not overlapping with the relational items:

- *Skills:* Critical Thinking; Complex Problem Solving; Mathematics; Reading Comprehension; Writing; Active Learning; Judgment and Decision Making; Systems Analysis; Systems Eval-

uation.

- *Work Activities*: Analyzing Data or Information; Processing Information; Making Decisions and Solving Problems; Updating and Using Relevant Knowledge; Thinking Creatively.

PC1 explains 70.2% of variance with all loadings positive in $[0.19, 0.30]$. The simple average correlates 0.78 with log mean wage and 0.65–0.74 with the three AI exposure measures, consistent with the interpretation in the main text that existing exposure measures load primarily on this dimension.

Eloundou exposure. $E_2 = \beta + 0.5\gamma$ from human raters in the Eloundou et al. (2023) replication file: β is the share of an occupation’s tasks where an LLM alone reduces time-to-task by at least 50%; γ is the additional share covered by LLM with tooling.

Felten AIOE and LM AIOE. Constructed by Felten et al. (2021) from O*NET ability requirements weighted by AI’s overlap with each ability. Used as published; the LM-specific variant restricts the application set to language-modelling capabilities.

B. Equilibrium Framework

This appendix collects the equilibrium framework on which the structural and transition results rest. The model is parameterized by the compute price $p_s > 0$, with wage and investment implications characterized as $p_s \rightarrow 0$. We first collect the technical assumptions invoked throughout the proofs, then state the equilibrium as a fixed-point problem in the wage vector, and establish existence at each p_s via Kakutani’s theorem. Uniqueness is imposed as a high-level condition (Assumption A5) and discussed below.

Technical assumptions

Throughout we maintain the following regularity, functional-form, and selection conditions. Each is hyperlinked at every appearance in the proofs below.

A1. Task distribution. The task weight distribution F over (α_{iH}, σ_i) pairs assigns positive measure to each of the three cells: \mathcal{C} -mixed ($\sigma_i = \sigma^C$, $\alpha_{iH} \in (0, 1)$), \mathcal{S} -mixed ($\sigma_i = \sigma^S$, $\alpha_{iH} \in (0, 1)$), and pure-D ($\alpha_{iH} = 0$). On \mathcal{C} - and \mathcal{S} -mixed tasks α_{iH} has continuous density on $[\underline{\alpha}, \bar{\alpha}] \subset (0, 1)$, bounded away from 0 and 1. For pure-D tasks the CES collapses to $q_i = s_D$ (with $\alpha_{iD} = 1$ by weight normalization) irrespective of σ_i ; hence σ_i on pure-D tasks is indeterminate and plays no role. The corner $\alpha_{iH} = 1$ (purely first-personal tasks) is treated as a limit of the \mathcal{C} -mixed case and is not a separate cell: analytically it delivers the same $q_i^{AI} = 0$ corner and adds no qualitative content. The partition of interest is the three-cell partition by $(\alpha_{iH}, \sigma_i|_{\text{mixed}})$.

A2. Skill functional forms. AI D-skill $s_D^{AI}(X_D^{AI}, C) = \bar{s}_D(1 - e^{-\gamma_D X_D^{AI} C})$; human D-skill $s_D^H(X_D^H) = \bar{s}_D^H(1 - e^{-\gamma_D^H X_D^H})$; human H-skill $s_H^H(X_H^H) = \bar{s}_H^H(1 - e^{-\gamma_H^H X_H^H})$; and $s_H^{AI} \equiv 0$ by the structural zero of Section 3.1.

A3. Market structure and pre-AI endpoint. Task production is one-agent-per-task (Section 3.2); labor markets are competitive; human labor supply is inelastic at \bar{L} ; the human investment budget satisfies $X_D^H + X_H^H \leq \bar{X}(T)$. The *pre-AI endpoint* is $p_s^{\text{pre-AI}} \equiv \sup\{p_s > 0 : \mathcal{A}(p_s) = \emptyset\}$; for $p_s \geq p_s^{\text{pre-AI}}$ humans produce every task, and under A1 and the structural zero this endpoint is finite. The *post-AI asymptotic endpoint* is $p_s \rightarrow 0$. Endpoint comparisons in the paper are between these two limits.

A4. Baumol condition ($\eta < 1$). The outer across-task elasticity η in (16) satisfies $\eta < 1$ strictly on the \mathcal{H}^* -intensive subindex; the knife-edge $\eta = 1$ is excluded, since at $\eta = 1$ expenditure shares would be constant and the Baumol channel would not operate.

A5. Equilibrium uniqueness. At each $p_s > 0$, the equilibrium wage vector $(w_D^*(p_s), w_H^*(p_s))$ is unique. Primitive sufficient conditions: (a) the Jacobian of excess demand at any candidate equilibrium has dominant diagonal; or (b) the aggregate best-response correspondence satisfies the gross-substitute / single-crossing property. A5 is invoked only to obtain

a single-valued continuous equilibrium selection $p_s \mapsto (w_D^*, w_H^*)(p_s)$ on $(0, p_s^{\text{pre-AI}}]$ via Berge's theorem. The main results (Corollary 2, Proposition C.1) use continuity together with the endpoint comparison $w_H^*(0^+) > w_H^*(p_s^{\text{pre-AI}})$; interior strict monotonicity of w_H is not claimed from A5 alone.

Note on elasticity ordering. The joint strict ordering $0 < \sigma^C < \eta < 1$ is not maintained as a standing assumption. It is a conditional regime introduced and characterized where it is needed, namely in the statement of Proposition 3, which describes the sign of within-task complementarity both when $\sigma^C < \eta$ (strictly positive; non-monotone D-wage result obtains) and when $\sigma^C \geq \eta$ (sign flips; the non-monotone result fails).

Equilibrium: definition, fixed-point statement, and existence

An equilibrium at a given p_s is a tuple $(w_D, w_H, Y, \{p_i\}_{i \in [0,1]}, \mathcal{A}(p_s))$ such that: (a) workers solve the static allocation problem (10)–(11) given (w_D, w_H) ; (b) tasks are assigned to the lower-cost producer via (9), with \mathcal{H}^* permanently human-assigned; (c) the task-level CES (6) and outer CES (16) imply task prices p_i and aggregate output Y at equilibrium factor allocations; (d) labor-market clearing determines (w_D, w_H) such that total demand for D-skill and H-skill equals supply. Existence of such an equilibrium is established below (Theorem B.1) via a Kakutani fixed-point argument on the wage vector (w_D, w_H) .

Aggregate production. Output is produced by combining a continuum of task outputs via a CES aggregator:

$$Y = \left(\int_0^1 q_i^{\frac{\eta-1}{\eta}} di \right)^{\frac{\eta}{\eta-1}} \quad (16)$$

where $\eta > 0$ is the elasticity of substitution *across tasks* in the aggregate production function — distinct from σ_i , the elasticity between the D-bundle and H-bundle *within* a task. The task price is $p_i = Y^{1/\eta} q_i^{-1/\eta}$, decreasing in q_i — tasks produced more cheaply command lower prices.

Task assignment and the structural floor. Assignment is one-agent-per-task: each task i is produced entirely by a single agent $a \in \{H, AI\}$ at per-unit costs $p_s c_i$ (AI) and $w(i) = w_D s_D^H +$

$w_H s_H^H$ (human). Let

$$i \in \mathcal{A}(p_s) \iff \frac{q_i^{AI}(p_s)}{p_s \cdot c_i} > \frac{q_i^H}{w(i)} \quad (17)$$

denote the contestable set AI captures. The structural floor \mathcal{H}^* defined in (8) has $q_i^{AI} = 0$ by Proposition 1(ii) and is never contestable, so $\mathcal{A}(p_s) \subseteq \mathcal{H}^{*c}$. The complement \mathcal{H}^{*c} comprises pure-D and \mathcal{S} -mixed tasks, both of which AI produces at bundle-compressed scale; as p_s falls, $\mathcal{A}(p_s)$ expands within \mathcal{H}^{*c} .

Existence at each p_s

Theorem B.1 (Existence of equilibrium). *Fix $p_s > 0$ and the structural zero $X_H^{AI} \equiv 0$ of Section 3.1 together with the technical assumptions A1–A4. There exists an equilibrium wage vector $(w_D^*, w_H^*) \in \mathbb{R}_{++}^2$ and associated allocation such that: (i) representative-agent investment (X_D^{H*}, X_H^{H*}) solves the worker’s allocation problem (10)–(11) given (w_D^*, w_H^*) ; (ii) the task-assignment set $\mathcal{A}(p_s)$ satisfies (9) pointwise; (iii) aggregate output Y and task prices $\{p_i\}$ solve the CES aggregators; (iv) labor markets for D-skill and H-skill clear.*

Proof. The argument is conducted at fixed $p_s > 0$ under the numeraire $PY = 1$. Both wage bounds below are constructed from primitives at the given p_s — the Walras factor-payment identity and the within-task CES cost-share floor on \mathcal{C} -mixed tasks — and no comparative-static result in p_s is invoked. We restrict the wage vector to a compact rectangle $K(p_s) = [\underline{w}_D(p_s), \bar{w}_D(p_s)] \times [\underline{w}_H(p_s), \bar{w}_H(p_s)] \subset \mathbb{R}_{++}^2$ constructed in two steps.

Cost-share floor on \mathcal{C} -mixed tasks. By A1, $\alpha_{iH} \in [\underline{\alpha}, \bar{\alpha}] \subset (0, 1)$ uniformly on \mathcal{H}^* ; combined with $\sigma^C < 1$, the equilibrium H-cost share

$$\theta_i^H = \frac{\alpha_{iH} (s_H^H)^{(\sigma^C - 1)/\sigma^C}}{\alpha_{iD} (s_D^H)^{(\sigma^C - 1)/\sigma^C} + \alpha_{iH} (s_H^H)^{(\sigma^C - 1)/\sigma^C}}$$

is bounded below by some $\underline{\theta}^H > 0$ on \mathcal{H}^* at any candidate equilibrium (with both skills positive by A2). Factor payments to H-skill on \mathcal{H}^* therefore satisfy $w_H \cdot L_H^H|_{\mathcal{H}^*}(p_s) \geq \underline{\theta}^H \cdot s_{H^*}(p_s) \cdot PY$,

where $s_{H^*}(p_s)$ is the expenditure share on \mathcal{H}^* at the candidate wage vector.

Interior-FOC restriction of $K(p_s)$. The cost-share floor delivers an *a priori* lower bound on w_H once we have a uniform lower bound $\underline{s}_{H^*} > 0$ on $s_{H^*}(p_s; w_D, w_H)$ over the candidate rectangle. Such a bound follows from primitives under a mild restriction of $K(p_s)$ to the region where the Ben-Porath optimum is interior. The Ben-Porath optimum is interior ($X_D^{H^*}, X_H^{H^*} > 0$) whenever the wage ratio satisfies $r_{\min} \leq w_H/w_D \leq r_{\max}$, where the bounds are determined by the functional forms of A2: at the corner $(X_D^H, X_H^H) = (\bar{X}, 0)$, the marginal condition $w_H(s_H^H)'(0) > w_D(s_D^H)'(\bar{X})$ excludes it as optimal, giving $r_{\min} \equiv \gamma_D^H \bar{s}_D^H e^{-\gamma_D^H \bar{X}} / (\gamma_H^H \bar{s}_H^H) > 0$; symmetrically $r_{\max} < \infty$ excludes the opposite corner. Restrict $K(p_s)$ to $\{(w_D, w_H) : r_{\min} \leq w_H/w_D \leq r_{\max}\}$, a compact convex polygon in \mathbb{R}_{++}^2 . On this restricted $K(p_s)$, the worker's objective $w_D s_D^H(X_D^H) + w_H s_H^H(X_H^H)$ is strictly concave in (X_D^H, X_H^H) by A2, so the argmax is single-valued and continuous in (w_D, w_H) by Berge's maximum theorem; the FOC pins each component strictly above zero. Compactness of $K(p_s)$ together with continuity of the unique argmax then delivers uniform interior bounds $X_D^{H^*} \geq \underline{X}_D > 0$ and $X_H^{H^*} \geq \underline{X}_H > 0$ for $(w_D, w_H) \in K(p_s)$. The restriction loses no equilibria: at a fixed point outside the wage-ratio range, the worker's corner allocation would force $s_{H^*} = 0$, contradicting the strict H-cost share on \mathcal{C} -mixed tasks under $\sigma^C < 1$ and positive α_{iH} .

Lower bounds on (w_D, w_H) . On the restricted $K(p_s)$, \mathcal{H}^* has positive measure $\mu_C > 0$ by A1, and on each $i \in \mathcal{H}^*$ the human output

$$q_i = [\alpha_{iD}(s_D^H)^{(\sigma^C-1)/\sigma^C} + \alpha_{iH}(s_H^H)^{(\sigma^C-1)/\sigma^C}]^{\sigma^C/(\sigma^C-1)}$$

is bounded below by a positive constant uniformly in (w_D, w_H) (skills bounded below; $\alpha_{iD}, \alpha_{iH} \geq \underline{\alpha} > 0$). Task prices p_i are bounded below by strictly positive competitive marginal cost; aggregate output is $PY = 1$. Hence $s_{H^*}(p_s; w_D, w_H) = \int_{\mathcal{H}^*} p_i q_i / PY \, di \geq \underline{s}_{H^*} > 0$ uniformly. With $L_H^H|_{\mathcal{H}^*}(p_s) \leq \bar{L}_H^H$ (A3, A2) and $PY = 1$, the cost-share floor delivers $\underline{w}_H(p_s) \equiv \frac{\theta^H}{\underline{s}_{H^*}} \cdot \underline{s}_{H^*} / (\bar{L}_H^H) > 0$. The lower bound on w_D follows symmetrically from the strictly positive D-cost share on \mathcal{C} -mixed tasks (since $\alpha_{iD} \in [1 - \bar{\alpha}, 1 - \underline{\alpha}]$ is also bounded away from zero).

Upper bounds on (w_D, w_H) . The upper bounds come from the aggregate resource constraint $PY = 1$: factor payments cannot exceed total revenue, so $\bar{w}_H(p_s) \leq 1/L_H^H|_{\min}$, where $L_H^H|_{\min}$ is a strictly positive lower bound on human H-skill supply to \mathcal{H}^* (positive by A1 and A3); $\bar{w}_D(p_s) \leq 1/\underline{L}_D$ analogously, where $\underline{L}_D > 0$ is the lower bound on D-labor supply to \mathcal{C} -mixed tasks (positive by A1). Both bounds are finite and strictly positive at any fixed $p_s > 0$.

Define the best-response correspondence $\Phi : K(p_s) \rightrightarrows K(p_s)$ as follows. Given $(w_D, w_H) \in K(p_s)$:

1. Solve the worker's allocation problem (10)–(11) to obtain $(X_D^{H^*}(w), X_H^{H^*}(w))$; these are single-valued and continuous in (w_D, w_H) by strict concavity of the skill functions and Berge's maximum theorem.
2. Compute equilibrium task assignments $\mathcal{A}(p_s; w)$ via (9). The assignment correspondence is upper hemicontinuous in w under the following regularity condition on F : the set of indifferent tasks $\{i : q_i^{AI}/(p_s c_i) = q_i^H/w(i)\}$ has Lebesgue measure zero at Lebesgue-a.e. w . Under the continuous-density assumption on α_{iH} in A1, the indifference condition is a single equation in continuous parameters and so cuts out a set of α_{iH} -density-zero at Lebesgue-a.e. w . Upper hemicontinuity of $\mathcal{A}(p_s; \cdot)$ as a parametric integral then follows from Berge's maximum theorem applied to the assignment program (the integrand is continuous in (i, w) off the measure-zero indifference set, and the task space is compact under A1).
3. Compute task prices $p_i(w)$ from the outer CES (16), which are continuous in w given continuous q_i .
4. Compute the implied market-clearing wage vector $\Phi(w) = (w'_D, w'_H)$ where w'_D equals integrated D-marginal-product over D-assigned tasks and w'_H equals integrated H-marginal-product over \mathcal{H}^* .

The correspondence Φ maps $K(p_s)$ into itself (by construction of the bounds via the resource constraint in step (b) above: any output of Φ respects the same total-revenue ceiling), is nonempty-valued, convex-valued (at any point), and upper hemicontinuous by (1)-(4). Kakutani's fixed-point

theorem applies: there exists $w^*(p_s) \in K(p_s)$ with $w^*(p_s) \in \Phi(w^*(p_s))$, which by construction is an equilibrium wage vector. Uniqueness is asserted separately in Assumption A5. Comparative statics in p_s — behavior of $w^*(p_s)$ as p_s varies, including the asymptotic comparison $w_H^*(p_s) \rightarrow w_H^\infty > w_H^{\text{pre-AI}}$ — are derived separately in Corollary 2 and Propositions C.1–C.4, using the fixed-point established here at each p_s as input. \square

Remark on the transition path. The proof establishes existence at each p_s separately. The comparative-statics results of Propositions C.1–C.4 and Proposition 3 describe movement along a family of equilibria indexed by p_s , not along an explicitly time-indexed transition path. A full dynamic equilibrium in which $p_s(t)$ falls exogenously and workers form rational expectations over future (w_D, w_H) paths would extend the static fixed-point of Appendix B to a time-varying-wage setting in which worker investment choices incorporate anticipated wage dynamics. The non-monotone D-wage prediction of Proposition 3 implies that expectations about when the economy will transition between C-dominated and S-dominated regimes matter for early-career investment choices; we leave characterization of the rational-expectations equilibrium to future work.

Uniqueness

Equilibrium uniqueness at each $p_s > 0$ is imposed as Assumption A5 above; the assumption is invoked whenever the comparative-statics results in Appendix C use continuity of the equilibrium selection $p_s \mapsto (w_D^*(p_s), w_H^*(p_s))$ obtained from Berge’s theorem. Primitive sufficient conditions are listed in the statement of A5; both a dominant-diagonal excess-demand condition and a gross-substitute property suffice. Under A5 the upper-hemicontinuous equilibrium correspondence Φ^{-1} is single-valued and therefore continuous on $(0, p_s^{\text{pre-AI}}]$.

C. Proofs

This appendix collects the proofs of all theorems, propositions, and corollaries stated in the main text, organized in three subsections mirroring the paper’s layers. Subsection C covers the structural results (Proposition 1, Proposition 2, Corollary 1, Remark 1), which do not invoke the fixed point and follow directly from the data-availability structure and the within-task CES. Subsection C covers the transition-path results (Proposition 3, Corollary 2), which use the equilibrium comparative statics in p_s . Subsection C states and proves the four comparative-statics propositions in p_s (Proposition C.1, Proposition C.2, Proposition C.3, Proposition C.4).

B.1 Structural results

Propositions 1–2, Corollary 1, and Remark 1 follow directly from the data-availability structure and the within-task CES with the one-agent-per-task assignment. We state the key steps for completeness.

Proof of Proposition 1(i) (component-level zero). $s_H^{AI}(X_H^{AI} = 0, C) = \bar{s}_H(1 - e^0) = 0$ for any C ; the α_{iH} -weighted term in (6) is identically zero for AI. \square

Proof of Proposition 1(ii) (task-level ceiling). Substituting $s_H^{AI} = 0$ into (6) and reading off each case of the partition (5) yields the three-case expression displayed in Proposition 1(ii); the \mathcal{C} -mixed case is the Leontief limit $\lim_{y \rightarrow 0^+} [\alpha_{iD}x^{(\sigma^C-1)/\sigma^C} + \alpha_{iH}y^{(\sigma^C-1)/\sigma^C}]^{\sigma^C/(\sigma^C-1)} = 0$ when $\alpha_{iH} > 0$ and $\sigma^C < 1$. Since $\partial s_H^{AI}/\partial C \equiv 0$, $\partial q_i^{AI}/\partial C = 0$ on \mathcal{C} -mixed tasks and strictly positive on \mathcal{S} -mixed and pure-D. \square

Proof of Proposition 2 (comparative advantage). Follows from Proposition 1 and the assignment condition (9). Part (i), pure-D tasks: $q_i^{AI} = s_D^{AI}$ rises without bound in $-\log p_s$, q_i^H is Ben-Porath-bounded, so $i \in \mathcal{A}(p_s)$ for small p_s . Part (ii), \mathcal{S} -mixed tasks: $q_i^{AI} = \alpha_{iD}^{\sigma^S/(\sigma^S-1)} s_D^{AI}$ rises without bound, so the same argument applies — bundle compression by $\alpha_{iD}^{\sigma^S/(\sigma^S-1)} < 1$ shifts the critical p_s but does not prevent eventual AI dominance. Part (iii), \mathcal{C} -mixed tasks: $q_i^{AI} = 0$ at every p_s by Proposition 1(ii), so the numerator of the AI cost ratio is zero and no human-offered

ratio is ever beaten. Corollary 1 is immediate: α_{iH} is orthogonal to any difficulty ordering on tasks by construction of the task space. \square

Remark 1 (common boundary). The claim rests on the same structural property that makes $X_H^{AI} \equiv 0$ — H -bundle capacities are constitutively first-personal and cannot be evaluated from outside the interaction any more than they can be trained from outside the organism — rather than on the mathematical structure of the model, which is why we state it as a remark rather than a proposition.

B.2 Transition-path results

The proof of Proposition 3 relies on a continuity-and-IVT argument that we package as a separate lemma to keep the regime-switch step short.

Lemma C.1 (Tail behavior of \dot{R} and existence of a regime-switch crossing). *Under A1–A4, A5, and the regime condition $\sigma^C < \eta$, define*

$$\dot{R}(\tau) \equiv \mu_C \dot{R}_C(\tau) - |\dot{R}_S(\tau)| \quad \text{on } (\tau_{\text{pre-AI}}, \infty),$$

where the constituent objects are as defined in the proof of Proposition 3. Then:

- (i) \dot{R} is continuous on $(\tau_{\text{pre-AI}}, \infty)$;
- (ii) there exist $\tau_1 < \infty$ and a constant $\underline{\dot{R}}_C > 0$, depending only on primitives, such that $\dot{R}(\tau) \geq \mu_C \cdot \underline{\dot{R}}_C > 0$ for all $\tau \geq \tau_1$;
- (iii) if additionally $\dot{R}(\tau_{\text{pre-AI}}^+) < 0$ (the S -dominated start), then there exists at least one $\bar{\tau} \in (\tau_{\text{pre-AI}}, \tau_1)$ with $\dot{R}(\bar{\tau}) = 0$.

Proof of Lemma C.1.

(i) *Continuity.* Under A5(a) the equilibrium selection $\tau \mapsto (w_D(\tau), w_H(\tau))$ is C^1 : the dominant-diagonal excess-demand Jacobian is invertible at every equilibrium, so the implicit function theorem applied to $E(p_s, w) = 0$ gives C^1 dependence of w on p_s , hence on $\tau = -\log p_s$. Under

the weaker A5(b), the equilibrium selection is Lipschitz on the bounded-wage rectangle of Theorem B.1, and Rademacher's theorem delivers differentiability on a full-measure subset on which the same continuity argument applies. \dot{R} is a difference of integrals whose integrands depend continuously on $(p_s, w_D, w_H, s_D^H, s_H^H, \dot{w}_D, \dot{w}_H)$ via the chain rule applied to $p_i \alpha_{iD} (s_D^H)^{-1/\sigma_i} q_i^{1/\sigma_i}$, so \dot{R} is continuous on $(\tau_{\text{pre-AI}}, \infty)$.

(ii) *Tail bound. Numerator.* By A1, $\alpha_{iH} \in [\underline{\alpha}, \bar{\alpha}] \subset (0, 1)$ uniformly on \mathcal{S} -mixed tasks, so the bundle-compression factor $g_i = \alpha_{iD}^{\sigma^S/(\sigma^S-1)}$ is bounded below by $\underline{g} \equiv (1 - \bar{\alpha})^{\sigma^S/(\sigma^S-1)} > 0$. The Proposition C.1(ii) uniform-boundedness argument (LHS of the assignment condition grows at rate $1/p_s$ times $g_i \geq \underline{g}$ while RHS $q_i^H/w(i)$ is bounded uniformly in i and p_s) delivers a uniform threshold $\underline{p}_s > 0$ with $\mathcal{A}(p_s) \supseteq \mathcal{H}^{*c}$ for all $p_s \leq \underline{p}_s$. Equivalently, $\mu^{\text{rem}}(p_s) \equiv 0$ on this tail, so $|\dot{R}_S(\tau)| = -d(\mu^{\text{rem}} R_S)/d\tau \equiv 0$ for all $\tau \geq \tau^* \equiv -\log \underline{p}_s$.

Denominator (Baumol price-effect channel). The lower bound on \dot{R}_C on the tail does not require interior strict monotonicity of w_H — which Step 5 of Corollary 2 explicitly disclaims — because the Baumol price effect on \mathcal{C} -mixed tasks delivers a uniform-in- τ contribution by itself. From Step 1 of Corollary 2, $P_{H^{*c}}(p_s) \rightarrow 0$ at rate p_s , while Step 2 establishes that $P_{H^*}(p_s)$ is bounded above and below; by the outer-CES dual (18) the expenditure share $s_{H^*}(p_s) \rightarrow 1$ as $p_s \rightarrow 0$ (Step 3 of Corollary 2), and since $P_{H^*} = s_{H^*}/Y_{H^*}$ with Y_{H^*} bounded, P_{H^*} is bounded below by a constant strictly above its pre-AI value on the tail. Each task-level price $p_i = Y^{1/\eta} q_i^{-1/\eta}$ on \mathcal{C} -mixed tasks is therefore bounded below by a primitive constant strictly above its pre-AI value, with $Y \geq \underline{Y} > 0$ uniformly (since $\mathcal{A}(p_s) \supseteq \mathcal{H}^{*c}$ delivers AI's bundle-compressed contribution on \mathcal{H}^{*c} on the tail). Differentiating along the equilibrium path, the Baumol contribution

$$\dot{R}_C^{(\text{Baumol})}(\tau) \equiv \int_{\mathcal{C}} \dot{p}_i \cdot \alpha_{iD} (s_D^H)^{-1/\sigma^C} q_i^{1/\sigma^C} di$$

admits a uniform lower bound $\underline{\dot{R}_C^{(B)}} > 0$ on the tail, by A1, A2, and the uniform wage-ratio bounds in Theorem B.1 (which keep s_D^H, s_H^H, q_i in compact intervals on \mathcal{C} -tasks). The relevant rate is explicit: from Step 1, $1 - s_{H^*}(p_s) = O(p_s^{1-\eta})$ on the tail, so $\dot{s}_{H^*}/(1 - s_{H^*}) \rightarrow (1 - \eta) > 0$ under

A4, and \dot{p}_i/p_i inherits a positive lower bound from this rate via the chain $p_i = (s_{H^*}/Y_{H^*})^{1/\eta} \cdot Y^{1/\eta} q_i^{-1/\eta} \cdot \mu(\mathcal{H}^*)^{-1/\eta}$ on \mathcal{C} -tasks. The within-human cross-partial channel under $\sigma^C < \eta$ (with A4, $\eta < 1$) and the integrand-level elasticity $1/\sigma^C - 1/\eta > 0$ reinforces \dot{R}_C wherever $\dot{s}_H^H > 0$ but is not required for the lower bound. Hence $\dot{R}_C(\tau) \geq \underline{\dot{R}}_C \equiv \underline{\dot{R}}_C^{(B)} > 0$ uniformly on the tail. Combining numerator and denominator, $\dot{R}(\tau) \geq \mu_C \cdot \underline{\dot{R}}_C > 0$ for $\tau \geq \tau_1$.

(iii) *IVT*. Continuity of \dot{R} on $(\tau_{\text{pre-AI}}, \tau_1]$ together with the hypothesis $\dot{R}(\tau_{\text{pre-AI}}^+) < 0$ and the tail-bound conclusion $\dot{R}(\tau_1) > 0$ gives, by the intermediate value theorem, at least one $\bar{\tau} \in (\tau_{\text{pre-AI}}, \tau_1)$ with $\dot{R}(\bar{\tau}) = 0$. \square

Remark C.1 (Sufficient conditions for single-crossing in Proposition 3(iii)). A sufficient condition for the monotonicity of $\mu_C^*(\tau)$ on $(0, p_s^{\text{pre-AI}}]$ is that both (a) $|\dot{R}_S(\tau)|$ is nonincreasing in τ and (b) $\dot{R}_C(\tau)$ is nondecreasing in τ . Condition (a) follows when the AI-absorption rate decelerates as $\alpha_S^*(p_s) \rightarrow 1$ (AI struggles with the marginal \mathcal{S} -tasks that are highest in α_{iH} ; the frontier slows, $\dot{\mu}^{\text{rem}} \rightarrow 0$, and Π_S approaches zero convexly). Condition (b) follows when $s_H^H(\tau)$ rises roughly concavely in τ , consistent with the Ben-Porath FOC under monotonically rising w_H . Neither condition is automatic from the regularity assumptions alone; both hold in the numerical equilibrium of Section 5.2, which exhibits a single-minimum U-shape in $X_D^{H^*}(\tau)$. Without single-crossing, the regime switch in Proposition 3(iii) still occurs at least once, but multiple switches are in principle possible.

Proof of Proposition 3 (Non-monotone D-investment)

Let $s_D^{AI}(p_s) \equiv \bar{s}_D^{AI}(1 - e^{-\gamma_D^{AI} S/p_s})$ denote AI's D-bundle skill on tasks it serves, monotonically increasing in $-\log p_s$ and bounded above by \bar{s}_D^{AI} . Let $\tau \equiv -\log p_s$ so that $\dot{x} \equiv dx/d\tau$ and “ p_s falling” corresponds to τ rising.

Marginal product decomposition. Under one-agent-per-task (Section 3.2), human D-skill earns on the two subsets of tasks humans serve: (a) the structural floor \mathcal{H}^* in (8), the set of \mathcal{C} -mixed tasks (where $\alpha_{iD} \in (0, 1)$); (b) the still-human portion of the contestable set \mathcal{H}^{*c} , which is $\mathcal{H}^{*c} \setminus \mathcal{A}(p_s)$ of measure $\mu^{\text{rem}}(p_s)$. The revenue-weighted marginal product of human D-skill, holding H-skill at

its optimized level $s_H^{H^*}(p_s)$, is

$$R(p_s) = \mu_C \cdot R_C(p_s) + \mu^{\text{rem}}(p_s) \cdot R_S(p_s),$$

where μ_C denotes the measure of \mathcal{C} -mixed tasks, and

$$R_C(p_s) \equiv \mathbb{E}_{i \in \mathcal{C}, \alpha_{iD} > 0} \left[p_i \alpha_{iD} (s_D^H)^{-1/\sigma^C} q_i^{1/\sigma^C} \right], \quad R_S(p_s) \equiv \mathbb{E}_{i \in \mathcal{H}^{*c} \setminus \mathcal{A}(p_s)} \left[p_i \alpha_{iD} (s_D^H)^{-1/\sigma_i} q_i^{1/\sigma_i} \right].$$

Sign of R_C at the endpoint (within-task complementarity). By Corollary 2 (proved below), $w_H^\infty > w_H^{\text{pre-AI}}$ strictly; by the Ben-Porath FOC (11), the optimal $X_H^{H^*}$ at the post-AI endpoint strictly exceeds its pre-AI value, and hence s_H^H is strictly larger at the endpoint. For interior p_s , the same sign holds whenever $w_H(p_s) > w_H^{\text{pre-AI}}$ (i.e. on any subinterval where the continuous path of w_H has already risen above its pre-AI level); we denote this local sign-information by $\dot{w}_H > 0$, $\dot{X}_H^{H^*} > 0$, $\dot{s}_H^H > 0$ without claiming strict interior monotonicity at every p_s . On \mathcal{C} -mixed tasks (permanently human), $q_i = [\alpha_{iD} (s_D^H)^{(\sigma^C - 1)/\sigma^C} + \alpha_{iH} (s_H^H)^{(\sigma^C - 1)/\sigma^C}]^{\sigma^C / (\sigma^C - 1)}$; since $\partial q_i / \partial s_H^H > 0$ on \mathcal{C} -mixed tasks, $\dot{q}_i > 0$. Substituting the equilibrium price $p_i = Y^{1/\eta} q_i^{-1/\eta}$ into the integrand $p_i \alpha_{iD} (s_D^H)^{-1/\sigma^C} q_i^{1/\sigma^C}$ gives a q_i -elasticity of $1/\sigma^C - 1/\eta$. For this to be strictly positive we need $\eta > \sigma^C$; this is exactly the conditional regime $\sigma^C < \eta$ introduced in the statement of Proposition 3, under which $1/\sigma^C - 1/\eta > 0$. Under A4 ($\eta < 1$ strict) combined with this regime condition, the integrand rises strictly in q_i , and within-task complementarity yields $\dot{R}_C > 0$: rising $s_H^H \Rightarrow$ rising q_i on \mathcal{C} -tasks \Rightarrow rising marginal product of same-agent s_D^H via the within-human cross-partial (12). (In the complementary regime $\sigma^C \geq \eta$, the elasticity is non-positive and $\dot{R}_C \leq 0$, as characterized in the statement of Proposition 3.) An additional direct contribution to \dot{R}_C comes from Y rising as $\mathcal{A}(p_s)$ expands, which raises the Baumol price index on \mathcal{H}^* ; this reinforces \dot{R}_C .

Sign of \dot{R}_S (\mathcal{S} -task reassignment). R_S is an integral over the set $\mathcal{H}^{*c} \setminus \mathcal{A}(p_s)$, whose measure $\mu^{\text{rem}}(p_s)$ weakly shrinks as p_s falls (Proposition C.1). Tasks exit this set permanently as AI takes over. For tasks remaining in this set, s_D^H earns with the same integrand form; the log-elasticity $1/\sigma_i - 1/\eta$ on remaining tasks takes mixed signs (negative on \mathcal{S} -mixed tasks where $\sigma^S > \eta$

is plausible), but the dominant effect as τ rises is the measure contraction: $\mu^{\text{rem}} \cdot R_S$ decreases because the set shrinks. Concretely, $d(\mu^{\text{rem}} R_S)/d\tau = \dot{\mu}^{\text{rem}} \cdot R_S + \mu^{\text{rem}} \cdot \dot{R}_S$ where $\dot{\mu}^{\text{rem}} \leq 0$ strictly when the boundary of $\mathcal{A}(p_s)$ sweeps across a positive-density set. Absorbing both terms, define $|\dot{R}_S| \equiv -d(\mu^{\text{rem}} R_S)/d\tau \geq 0$ so that the second term in $R(p_s) = \mu_C R_C + \mu^{\text{rem}} R_S$ has derivative $-|\dot{R}_S|$, with strict inequality whenever $\dot{\mu}^{\text{rem}} < 0$ and $R_S > 0$.

Threshold and cases (i)–(ii). Writing

$$\dot{R} = \mu_C \dot{R}_C - |\dot{R}_S|,$$

the sign of \dot{R} equals the sign of $\mu_C - \mu_C^*(p_s)$, where $\mu_C^*(p_s) \equiv |\dot{R}_S(p_s)|/\dot{R}_C(p_s)$. Since the labor-market-clearing wage w_D is a monotone transformation of R (holding s_H^{H*} at its optimized level; labor supply inelastic by A3), \dot{w}_D shares the sign of \dot{R} . The interior Ben-Porath FOC (11) pins X_D^{H*} as an increasing function of w_D , so \dot{X}_D^{H*} shares the sign of \dot{w}_D . Cases (i) and (ii) follow.

Case (iii): regime switch. Work in the regime $\sigma^C < \eta$; the alternative regime $\sigma^C \geq \eta$ is characterized in the statement of the proposition and does not produce regime switching. Suppose the economy begins S-dominated at pre-AI: $\mu_C^*(\tau_{\text{pre-AI}}^+) > \mu_C$, equivalently $\dot{R}(\tau_{\text{pre-AI}}^+) < 0$. By Lemma C.1(iii), there exists $\bar{\tau} \in (\tau_{\text{pre-AI}}, \tau_1)$ with $\dot{R}(\bar{\tau}) = 0$. Cases (i)–(ii) give $\text{sign}(\dot{w}_D) = \text{sign}(\dot{R})$, so $\dot{w}_D < 0$ before $\bar{\tau}$ and $\dot{w}_D > 0$ on the tail $[\tau_1, \infty)$; \dot{X}_D^{H*} inherits the sign change via the Ben-Porath FOC. If instead the economy begins C-dominated at pre-AI ($\mu_C^* < \mu_C$), no crossing is forced and the path is rising throughout. \square

Case (iv): strict U-shape under single-crossing. Impose additionally the single-crossing hypothesis: $\mu_C^*(\tau)$ is weakly monotonically nonincreasing in τ on $(\tau_{\text{pre-AI}}, \infty)$ (sufficient conditions are discussed in Remark C.1; both (SC1) $|\dot{R}_S|(\tau)$ nonincreasing and (SC2) $\dot{R}_C(\tau)$ nondecreasing make the ratio nonincreasing; the numerical equilibrium in Section 5.2 satisfies both). Combined with the lemma's IVT and the strict endpoint values $\mu_C^*(\tau_{\text{pre-AI}}) > \mu_C$ and $\mu_C^*(\tau) = 0 < \mu_C$ for $\tau \geq \tau_1$, monotonicity forces the crossing $\bar{\tau}$ to be unique. On $(\tau_{\text{pre-AI}}, \bar{\tau})$, $\mu_C^*(\tau) > \mu_C$ (nonincreasing plus unique crossing plus positive initial gap) and therefore $\dot{w}_D(\tau) < 0$ strictly; on $(\bar{\tau}, \infty)$,

$\mu_C^*(\tau) < \mu_C$ and $\dot{w}_D(\tau) > 0$ strictly. Hence $w_D(\tau)$ is strictly decreasing on $(\tau_{\text{pre-AI}}, \bar{\tau})$ and strictly increasing on $(\bar{\tau}, \infty)$, i.e. strictly U-shaped with a unique turning point at $\bar{\tau}$. By the Ben-Porath FOC, $X_D^{H^*}(\tau)$ inherits the strict U-shape. \square

Sufficient condition for $\dot{R}_C > 0$ throughout. Under A1 the subset of \mathcal{C} -mixed tasks with $\alpha_{iD}\alpha_{iH} > 0$ has positive measure. Under the regime condition $\sigma^C < \eta$ (with A4 giving $\eta < 1$), the integrand-level elasticity $1/\sigma^C - 1/\eta$ is strictly positive, and Corollary 2 guarantees $\dot{s}_H^H > 0$ at the endpoint (and locally wherever $w_H(p_s) > w_H^{\text{pre-AI}}$). Hence $\dot{R}_C > 0$ strictly on such subintervals. \square

Proof of Corollary 2 (w_H unambiguous)

Proof strategy. The proof uses an *aggregate* Baumol argument based on expenditure shares under bounded factor supply, rather than an integrand-by-integrand approach on $R^H = \int_{\mathcal{H}^*} p_i \cdot (\partial q_i / \partial s_H^H) di$, which fails because in case (ii) of Proposition 3 $\partial q_i / \partial s_H^H$ can fall on \mathcal{C} -mixed tasks as s_D^H falls and s_H^H rises.

Setup: aggregate CES. Partition task output into two composites via the outer CES (16): Y_{H^*} (aggregated over $i \in \mathcal{H}^*$) and $Y_{H^{*c}}$ (aggregated over $i \in \mathcal{H}^{*c}$), with associated price indices P_{H^*} and $P_{H^{*c}}$. The outer-CES dual gives

$$s_{H^*}(p_s) \equiv \frac{P_{H^*} Y_{H^*}}{P Y} = \omega_{H^*} \left(\frac{P_{H^*}}{P} \right)^{1-\eta}, \quad P^{1-\eta} = \omega_{H^*} P_{H^*}^{1-\eta} + \omega_{H^{*c}} P_{H^{*c}}^{1-\eta}, \quad (18)$$

where $\omega_{H^*}, \omega_{H^{*c}}$ are the CES weights. Under A4 ($\eta < 1$ strict), $1 - \eta > 0$, so s_{H^*} strictly increases in P_{H^*}/P .

Step 1: $P_{H^{*c}}(p_s) \rightarrow 0$ as $p_s \rightarrow 0$. On each $i \in \mathcal{A}(p_s) \subset \mathcal{H}^{*c}$, AI's per-task input cost is $p_s c_i$ and AI output is q_i^{AI} , so the equilibrium task price under competitive entry equals AI's per-output marginal cost $p_i = p_s c_i / q_i^{AI}$. Since $q_i^{AI} \rightarrow \bar{s}_D^{AI} \cdot g_i$ with $g_i \in [\underline{g}, 1]$ bounded and positive (Step 2 below relies on the same lower bound on g_i on \mathcal{S} -mixed tasks), $p_i = O(p_s)$ uniformly in i on the tail. Aggregating via the inner CES on \mathcal{H}^{*c} yields $P_{H^{*c}}(p_s) = O(p_s)$, hence $P_{H^{*c}}(p_s) \rightarrow 0$ at rate

p_s .

Step 2: $P_{H^*}(p_s)$ is bounded and positive. By A3, Y_{H^*} is bounded above by a constant \bar{Y}_{H^*} (skill saturation); by A1, $\mathcal{H}^* = \{i : \sigma_i = \sigma^C, \alpha_{iH} \in (0, 1)\}$ has positive measure $\mu_C > 0$ and α_{iD}, α_{iH} are bounded away from 0 and 1, so Y_{H^*} is bounded below by a positive constant $\underline{Y}_{H^*} > 0$ whenever $X_D^{H^*}, X_H^{H^*} > 0$ (both bundles enter the C-mixed CES). Under Walras-consistent normalization $PY = 1$, the outer-CES optimization gives $P_{H^*}Y_{H^*} = s_{H^*}$, hence $P_{H^*} = s_{H^*}/Y_{H^*}$. Since $s_{H^*} \in (0, 1)$ and Y_{H^*} is bounded above and below, P_{H^*} is bounded above and below as well.

Step 3: $s_{H^*}(p_s) \rightarrow 1$ as $p_s \rightarrow 0$. From (18), $s_{H^{*c}} = \omega_{H^{*c}}(P_{H^{*c}}/P)^{1-\eta}$. Since $P_{H^{*c}} \rightarrow 0$ (Step 1) while P stays bounded below (dominated by the $P_{H^*}^{1-\eta}$ term which is bounded below positive by Step 2), $(P_{H^{*c}}/P)^{1-\eta} \rightarrow 0$ at rate $p_s^{1-\eta}$. Hence $s_{H^{*c}}(p_s) \rightarrow 0$ and $s_{H^*}(p_s) = 1 - s_{H^{*c}}(p_s) \rightarrow 1$.

Step 4: lower bound on w_H rises strictly; endpoint comparison $w_H^\infty > w_H^{\text{pre-AI}}$. Factor payments on \mathcal{H}^* equal revenue on \mathcal{H}^* (under one-agent-per-task, humans receive all revenue on \mathcal{H}^* since $q_i^{AI} = 0$ there by Proposition 1):

$$w_D \cdot L_D^H|_{\mathcal{H}^*}(p_s) + w_H \cdot L_H^H|_{\mathcal{H}^*}(p_s) = s_{H^*}(p_s).$$

On \mathcal{C} -mixed tasks, $\sigma^C < 1$ and $\alpha_{iH} \in (0, 1)$ imply that competitive pricing delivers a strictly positive H-skill revenue share. Specifically, writing the H-skill share on task i as

$$\theta_i^H \equiv \frac{\alpha_{iH}(s_H^H)^{(\sigma^C-1)/\sigma^C}}{\alpha_{iD}(s_D^H)^{(\sigma^C-1)/\sigma^C} + \alpha_{iH}(s_H^H)^{(\sigma^C-1)/\sigma^C}} \in (0, 1),$$

which is the standard CES cost-share expression, we have $w_H s_H^H \cdot (\partial \ln q_i / \partial \ln s_H^H) = \theta_i^H \cdot p_i q_i$ by Euler's theorem. Under A1 (the support of α_{iH} on \mathcal{C} -mixed tasks is bounded away from 0 and 1) combined with A2 and A3 (bounded skills, inelastic supply), there exists $\underline{\theta}^H > 0$ such that $\theta_i^H \geq \underline{\theta}^H$ uniformly in p_s on \mathcal{C} -mixed tasks. Therefore

$$w_H \cdot L_H^H|_{\mathcal{H}^*}(p_s) = \int_{\mathcal{H}^*} \theta_i^H \cdot p_i q_i di \geq \underline{\theta}^H \cdot s_{H^*}(p_s).$$

With $L_H^H|_{\mathcal{H}^*}(p_s) \leq \bar{L}\bar{s}_H^H$ (A3, A2), define the lower bound

$$g(p_s) \equiv \frac{\theta^H \cdot s_{H^*}(p_s)}{\bar{L}\bar{s}_H^H} \leq w_H(p_s).$$

Endpoint properties: (a) As $p_s \rightarrow 0$, $s_{H^*}(p_s) \rightarrow 1$ (Step 3), so $g(p_s) \rightarrow g^\infty \equiv \underline{\theta}^H/(\bar{L}\bar{s}_H^H) > 0$; (b) at pre-AI ($\mathcal{A} = \emptyset$), humans also produce \mathcal{H}^{*c} , so $s_{H^*}(p_s^{\text{pre-AI}}) < 1$ (since \mathcal{H}^{*c} has positive measure by A1).

Explicit pre-AI computation and strict endpoint comparison. We work with the H-skill expenditure share rather than the lower bound g . Let $s_{\text{earn}}^H(p_s) \equiv w_H(p_s) \cdot L_H^H/PY = w_H(p_s) \cdot L_H^H$ under $PY = 1$, where $L_H^H \equiv \bar{L}\bar{s}_H^H$ is the inelastic total H-skill supply (A3, A2). At pre-AI the factor-payment identity decomposes as

$$s_{\text{earn}}^{H,\text{pre}} = s_{H^*}^{\text{pre}} \cdot \bar{\theta}_{\mathcal{H}^*,\text{pre}}^H + s_S^{\text{pre}} \cdot \bar{\theta}_{S,\text{pre}}^H + 0,$$

where $\bar{\theta}_{R,\cdot}^H$ is the revenue-weighted within-task H-cost share on region R and the pure-D contribution vanishes because $\alpha_{iH} = 0$ there. At the post-AI limit, $s_{H^*}^\infty = 1$ (Step 3) and AI pays no H-rent on \mathcal{H}^{*c} (since $s_H^{AI} \equiv 0$), so

$$s_{\text{earn}}^{H,\infty} = \bar{\theta}_{\mathcal{H}^*,\infty}^H.$$

Substituting the adding-up constraint $s_{H^*}^{\text{pre}} = 1 - s_S^{\text{pre}} - s_{\text{pureD}}^{\text{pre}}$ and collecting terms,

$$s_{\text{earn}}^{H,\infty} - s_{\text{earn}}^{H,\text{pre}} = \underbrace{(\bar{\theta}_{\mathcal{H}^*,\infty}^H - \bar{\theta}_{\mathcal{H}^*,\text{pre}}^H)}_{\text{(i): sign ambiguous}} + \underbrace{s_S^{\text{pre}} (\bar{\theta}_{\mathcal{H}^*,\text{pre}}^H - \bar{\theta}_{S,\text{pre}}^H)}_{\text{(ii)} \geq 0} + \underbrace{s_{\text{pureD}}^{\text{pre}} \cdot \bar{\theta}_{\mathcal{H}^*,\text{pre}}^H}_{\text{(iii)} > 0 \text{ strictly}}.$$

The three terms are respectively:

(i) *Composition effect, sign ambiguous in general, magnitude bounded by a path-integral.* The CES cost-share formula $\theta_i^H = \alpha_{iH}/[\alpha_{iD}(s_D^H/s_H^H)^{(\sigma^C-1)/\sigma^C} + \alpha_{iH}]$ has negative exponent $(\sigma^C - 1)/\sigma^C < 0$ under $\sigma^C < 1$, so θ_i^H is *strictly decreasing* in the ratio s_H^H/s_D^H . Ben-Porath reallocation between pre-AI and the asymptotic endpoint can move this ratio in either direction depending on

the relative change in (w_D, w_H) . To bound the magnitude of term (i), differentiate θ_i^H along the equilibrium path: $d\theta_i^H/d\log(s_H^H/s_D^H) = -\theta_i^H(1 - \theta_i^H) \cdot |1 - 1/\sigma^C|$. Integrating along the path, the change in the revenue-weighted average $\bar{\theta}_{\mathcal{H}^*}^H$ between pre-AI and the asymptotic endpoint is bounded:

$$|\bar{\theta}_{\mathcal{H}^*,\infty}^H - \bar{\theta}_{\mathcal{H}^*,\text{pre}}^H| \leq \frac{1}{4} \cdot |1 - 1/\sigma^C| \cdot |\Delta \log(s_H^H/s_D^H)| \equiv B^*, \quad (19)$$

where $\Delta \log(s_H^H/s_D^H)$ is bounded uniformly because (w_D, w_H) lies in the compact wage-ratio rectangle $[r_{\min}, r_{\max}]$ of Theorem B.1 and (s_D^H, s_H^H) are the resulting Ben-Porath optima. The path-integrated bound B^* is generally tighter than the loose envelope $B \equiv \bar{\theta}^H - \underline{\theta}^H$ used by setting one endpoint to the supremum and the other to the infimum; we use the tighter bound below.

(ii) *Weakly positive under density dominance.* Under the equilibrium pre-AI factor allocation (Theorem B.1, which pins (s_D^H, s_H^H) in a compact range), the CES cost-share formula delivers a *pointwise in α_{iH} comparison*: with $e^C \equiv (\sigma^C - 1)/\sigma^C < 0$ and $e^S \equiv (\sigma^S - 1)/\sigma^S > 0$ (under $\sigma^S > 1$), $\theta_C^H(\alpha_{iH}; s_D^H, s_H^H) \geq \alpha_{iH} \geq \theta_S^H(\alpha_{iH}; s_D^H, s_H^H)$ holds whenever $s_D^H \geq s_H^H$ pre-AI; *the latter inequality is an equilibrium property in the regime $\mu(\text{pure-D}) > 0$, where pure-D tasks pay only D-rent and tilt the Ben-Porath FOC toward $X_D^H > X_H^H$.*⁷ Integrating the revenue-weighted averages:

$$\bar{\theta}_{\mathcal{H}^*,\text{pre}}^H \geq \bar{\alpha}_{\mathcal{H}^*,\text{pre}}^{\text{rev}}, \quad \bar{\theta}_{\mathcal{S},\text{pre}}^H \leq \bar{\alpha}_{\mathcal{S},\text{pre}}^{\text{rev}},$$

where $\bar{\alpha}_{R,\text{pre}}^{\text{rev}} \equiv \int_R \alpha_{iH} d\mu^{\text{rev}}(i)$ is the pre-AI revenue-weighted mean of α_{iH} on region R . Hence $\bar{\theta}_{\mathcal{H}^*,\text{pre}}^H \geq \bar{\theta}_{\mathcal{S},\text{pre}}^H$ whenever the *density dominance condition* $\bar{\alpha}_{\mathcal{H}^*,\text{pre}}^{\text{rev}} \geq \bar{\alpha}_{\mathcal{S},\text{pre}}^{\text{rev}}$ holds — the substantive content of the positive α - σ alignment documented empirically in Section 2.2 and adopted as the benchmark in Section 6 (\mathcal{C} : Beta(5, 2), mean ≈ 0.71 ; \mathcal{S} : Beta(2, 5), mean ≈ 0.29). When density dominance fails, term (ii) can be negative but is bounded by $|(ii)| \leq \bar{\alpha} - \underline{\alpha}$. The combined comparison below treats Case A (density dominance) as the substantive case and weakens to a sign-robust statement otherwise.

⁷If the task distribution is sufficiently H-intensive that $w_H > w_D$ pre-AI, the pointwise comparison reverses for $s_H^H > s_D^H$. The endpoint comparison below then runs through the same algebra with the inequality direction flipped and remains weakly positive under the analogous density-dominance condition.

(iii) *Strictly positive.* $s_{\text{pureD}}^{\text{pre}} > 0$ (by A1, pure-D has positive measure) and $\bar{\theta}_{\mathcal{H}^*, \text{pre}}^H \geq \underline{\theta}^H > 0$, so (iii) $\geq s_{\text{pureD}}^{\text{pre}} \cdot \underline{\theta}^H > 0$.

Combined strict comparison. The substantive claim is delivered under density dominance (the benchmark case), with the path-integrated bound B^* from (19) replacing the loose envelope B in the sufficient condition.

Case A (substantive: density dominance). Term (ii) ≥ 0 , so (i) + (ii) + (iii) $\geq -B^* + s_{\text{pureD}}^{\text{pre}} \cdot \underline{\theta}^H$. The strict comparison holds under

$$s_{\text{pureD}}^{\text{pre}} \cdot \underline{\theta}^H > B^*. \quad (20)$$

Equivalently, the pure-D pre-AI revenue share must exceed the path-integrated cost-share swing $B^* = \frac{1}{4} \cdot |1 - 1/\sigma^C| \cdot |\Delta \log(s_H^H/s_D^H)|$ (using $\theta^H(1 - \theta^H) \leq 1/4$). This condition is empirically reasonable when pre-AI cognitive professional work (lawyers, analysts, software engineers) accounts for substantial revenue and the wage-ratio movement across the transition is bounded.

Case B (sign-robust, without density dominance). If density dominance fails, term (ii) contributes at most $-(\bar{\alpha} - \underline{\alpha})$, so the sign-robust condition strengthens to $s_{\text{pureD}}^{\text{pre}} \cdot \underline{\theta}^H > B^* + (\bar{\alpha} - \underline{\alpha})$. Under the wide cost-share-swing regime $\sigma^C \rightarrow 0$, $|1 - 1/\sigma^C| \rightarrow \infty$ and B^* can grow large; in that regime the analytic sufficient condition above may fail.

The condition (20) is not vacuous in plausible parameter ranges but is restrictive when σ^C is small. The benchmark sits in the latter regime ($\sigma^C = 0.30$); the strict comparison there rests on numerical verification rather than on (20) (Table 2: w_H/P rises from 1 to 4.52, with the same sign throughout the parameter sweeps of Section 6.2).

Dividing through by the constant L_H^H yields $w_H^\infty > w_H^{\text{pre-AI}}$ strictly under (20) plus density dominance, and via numerical verification under the benchmark parameterization. The economic content: pure-D revenue, which pays zero H-rent pre-AI, is redirected by the Baumol channel toward \mathcal{H}^* post-AI where it pays strictly positive H-rent. Term (ii) provides a reinforcing contribution under density dominance; term (i) captures a composition effect from Ben-Porath reallocation whose sign depends on the case of Prop 3 but whose magnitude is bounded by the path-integral B^* . Interior strict monotonicity of $w_H(p_s)$ in $-\log p_s$ is not claimed from this argument alone; see

Step 5.

Step 5: continuity of equilibrium selection in p_s . The equilibrium correspondence at each $p_s > 0$ (Theorem B.1) has upper hemicontinuous best-response map. Under Assumption A5 (*equilibrium uniqueness at each $p_s > 0$*), the upper hemicontinuous correspondence Φ^{-1} (mapping p_s to the set of equilibrium wage vectors) becomes single-valued; by the standard topological fact that an upper hemicontinuous compact-valued correspondence that is single-valued is continuous (a direct corollary of Berge’s maximum theorem applied to the parametric fixed-point problem), the equilibrium selection $p_s \mapsto (w_D^*(p_s), w_H^*(p_s))$ is continuous on $(0, p_s^{\text{pre-AI}}]$. Steps 1–4 then imply: (i) $w_H^*(p_s)$ is bounded below by a strictly positive Baumol floor uniformly in p_s (from A4 applied to \mathcal{C} -mixed tasks, whose $\alpha_{iH} \in (0, 1)$ delivers a positive H-skill cost share); (ii) $w_H^*(p_s) \rightarrow w_H^\infty \leq 1/(\bar{L}\bar{s}_H^H)$ as $p_s \rightarrow 0$; and (iii) $w_H^*(p_s^{\text{pre-AI}}) < w_H^\infty$ strictly, so w_H^* rises from its pre-AI value to its asymptotic limit along *any* continuous path. Strict interior monotonicity is not claimed; local reversals are logically possible under only continuity. The downstream results invoke only the endpoint comparison $w_H^*(0^+) > w_H^*(p_s^{\text{pre-AI}})$ and continuity.

Remark on within-human channel. In case (i) of Proposition 3 (X_D^{H*} rises), the within-agent cross-partial (12) reinforces the aggregate Baumol channel on \mathcal{C} -mixed tasks; in case (ii), it can partially offset. The aggregate Baumol argument (Steps 1–4) does not depend on within-task dynamics on \mathcal{C} -tasks — only on the aggregate expenditure-share shift driven by $P_{H^*c} \rightarrow 0$ and the bounded factor supply. Hence $w_H \uparrow$ strictly in both cases. \square

Proof of Proposition 4 (Short-run wage incidence)

Hold $(s_D^H, s_H^H) = (s_D^{H,\text{pre-AI}}, s_H^{H,\text{pre-AI}})$ throughout; otherwise the equilibrium structure is that of Theorem B.1, with the worker-allocation step replaced by an exogenous skill vector.

Part (i). As p_s falls, $\mathcal{A}(p_s)$ expands monotonically over \mathcal{H}^{*c} by Proposition C.1. AI’s bundle-compressed marginal cost on absorbed tasks falls toward zero, so the relative share of expenditure on \mathcal{H}^{*c} falls and the share on \mathcal{H}^* rises. Under A4 ($\eta < 1$), the \mathcal{H}^* -expenditure share converges to one as $p_s \rightarrow 0$. With $PY = 1$ as numeraire and $L_H^H|_{\mathcal{H}^*} = \bar{L} \cdot s_H^{H,\text{pre-AI}}$ pinned by the fixed

skill vector, $w_H^{SR}(p_s) \cdot \bar{L} \cdot s_H^{H, \text{pre-AI}}$ converges to a finite positive limit strictly above its pre-AI value, since the pre-AI \mathcal{H}^* -expenditure share is bounded strictly below one. Monotonicity in τ on any interval of expanding $\mathcal{A}(p_s)$ follows from the same argument applied locally. The density-dominance condition of Corollary 2 is not invoked because no offsetting supply expansion enters when s_H^H is held fixed.

Part (ii). The cross-partial (12) enters \dot{R}_C along the long-run path multiplied by $\dot{s}_H^{H^*}$; at fixed $s_H^{H, \text{pre-AI}}$ this multiplier is zero, so $\dot{R}_C^{SR}(p_s)$ retains only the Baumol price-effect component. On the early phase of the transition where $\mathcal{A}(p_s)$ is absorbing \mathcal{S} -mixed and pure- D tasks, the \mathcal{S} -reassignment force $|\dot{\Pi}_S|$ operates as in Proposition 3; the \mathcal{C} -task force is unambiguously smaller because the cross-partial channel is silent. Pointwise, the short-run threshold $\mu_C^{*,SR}(p_s) \equiv |\dot{\Pi}_S(p_s)|/\dot{R}_C^{SR}(p_s)$ satisfies $\mu_C^{*,SR} \geq \mu_C^*$, so whenever $\mu_C^*(p_s^{\text{pre-AI}}) > \mu_C$ also $\mu_C^{*,SR}(p_s^{\text{pre-AI}}) > \mu_C$, placing the short-run path in the falling regime at the pre-AI endpoint. As $p_s \rightarrow 0$, $|\dot{\Pi}_S| \rightarrow 0$ (the contestable set shrinks to a null set under Proposition C.1) while $\dot{R}_C^{SR} > 0$ retains the Baumol price-effect component; the late-stage lift in w_D^{SR} is therefore possible, but the strict-U conclusion of Proposition 3(iii) does not extend without an additional magnitude condition on the Baumol price effect on \mathcal{C} -tasks. \square

B.3 Comparative statics in p_s

Proposition C.1 (Automation threshold). *As $p_s \rightarrow 0$:*

- (i) *The set of AI-assigned tasks $\mathcal{A}(p_s)$ expands monotonically when the equilibrium wage vector $(w_D(p_s), w_H(p_s))$ is held fixed; under full equilibrium wage feedback, monotonicity along the path is not asserted.*
- (ii) *$\mathcal{A}(p_s)$ converges to $[0, 1] \setminus \mathcal{H}^*$: AI eventually dominates all tasks outside the structural human set.*
- (iii) *The measure of human-assigned tasks converges to $\mu(\mathcal{H}^*) > 0$, which is strictly positive whenever the task weight distribution assigns positive measure to \mathcal{C} -mixed tasks ($\sigma_i = \sigma^C$,*

$$\alpha_{iH} \in (0, 1).$$

Asymptotic convergence (parts (ii) and (iii)) is what the paper’s main results require; pointwise monotonicity of $\mathcal{A}(p_s)$ along the equilibrium path is not invoked anywhere downstream.

This is the formal statement of the paper’s central claim: falling compute prices asymptotically expand AI’s domain up to the structural floor \mathcal{H}^* determined by the H -bundle structure. The floor is not a quantitative threshold that more compute can overcome; it is determined by a property of the task weight distribution and the architecture of human minds.

Proposition C.2 (Wage effects). *As p_s falls:*

- (i) **D-wage.** *The equilibrium D-wage w_D is non-monotone in p_s per Proposition 3: in the S-dominated regime it falls, in the C-dominated regime it rises, and along the transition it can be U-shaped as $\mu^{rem}(p_s) \rightarrow 0$ exhausts S-task reassignment.*
- (ii) **H-wage.** *The equilibrium H-wage w_H rises from pre-AI value w_H^{pre-AI} to a strictly larger finite limit w_H^∞ as $p_s \rightarrow 0$ per Corollary 2. Under Assumption A4 (outer across-task elasticity $\eta < 1$ strictly for the H-intensive subindex), the Baumol channel delivers the strict endpoint comparison $w_H^\infty > w_H^{pre-AI}$; the within-human complementarity channel on C-mixed tasks reinforces this when $X_D^{H^*}$ rises. Interior strict monotonicity is not claimed without further regularity (A5(b)).*
- (iii) **No across-agent complementarity.** *There is no “collaboration premium” within a task under one-agent-per-task assignment. The apparent complementarity economists observe in mixed occupations (physicians, engineers, researchers) operates across tasks: AI absorbs S-mixed and pure-D tasks formerly served by the worker, freeing the worker to specialize in C-mixed tasks where the within-human cross-partial (12) amplifies earnings. Workers whose occupation bundle is disproportionately C-mixed thus experience rising total earnings via both the Baumol channel (raising w_H) and the within-human channel (raising the return to s_D^H on C-tasks).*

Proposition C.3 (Investment reallocation). *In the interior of the worker's allocation problem, optimal human investment responds to wages via the FOC (11):*

- (i) **H-bundle investment rises at the endpoint.** *Since $w_H^\infty > w_H^{\text{pre-AI}}$ (Corollary 2), optimal $X_H^{H^*}(p_s)$ in the Ben-Porath FOC rises from its pre-AI level to a strictly higher asymptotic level; interior $X_H^{H^*}(p_s)$ follows the same (possibly non-monotone interior) path as $w_H(p_s)$.*
- (ii) **D-bundle investment is non-monotone.** *Since w_D can be non-monotone (Proposition 3), optimal $X_D^{H^*}$ inherits the regime structure: falls in the S-dominated phase, rises in the C-dominated phase. Workers entering the labor market during an S-dominated transition who optimally invest in D-skill can find themselves over-invested if the economy later switches to C-dominated (part (iii) of Proposition 3 renders this strand endogenously U-shaped).*
- (iii) **Horizon-bounded capacity.** *Affective and co-constructive capacity responds to deliberate investment; mortality-grounded authority accrues constitutively through living and cannot be accelerated. Workers targeting high-mortality-weighted occupations face transition paths bounded by calendar time, not budget.*

Proposition C.4 (Polarization). *As p_s falls, the wage distribution polarizes along the (α_{iH}, σ_i) -dimensions rather than along conventional skill-level dimensions:*

- (i) *Workers whose task bundle lies predominantly in the structural floor \mathcal{H}^* (i.e., C-mixed tasks) experience strictly higher asymptotic earnings than under pre-AI wages: $w_H^\infty > w_H^{\text{pre-AI}}$ (Corollary 2 endpoint comparison), and on C-mixed tasks their s_D^H earns the within-human complementarity premium captured by (12).*
- (ii) *Workers whose task bundle lies predominantly in the contestable set \mathcal{H}^{*c} (S-mixed and pure-D) face shrinking task measure as AI absorbs $\mathcal{A}(p_s) \rightarrow \mathcal{H}^{*c}$; their earnings trajectory depends on whether and how quickly they reallocate to \mathcal{H}^* -intensive occupations.*
- (iii) *With $\eta < 1$, this generates a bimodal distribution in occupation-level earnings with a growing gap between \mathcal{H}^* -intensive and \mathcal{H}^{*c} -intensive occupations. The gap widens with the rate of*

compute-price decline.

(iv) *This polarization pattern differs from the task-content polarization of Autor, Levy, and Murnane (2003) in a testable way: the Autor, Levy, and Murnane mechanism predicts hollowing of middle-skill routine work with growth at both skill extremes; the present framework predicts hollowing of \mathcal{H}^{*c} -intensive work at all skill levels with growth concentrated in \mathcal{H}^* -intensive work, independent of the worker’s measured skill level.*

The polarization result has a direct implication for the standard “skill-biased technological change” narrative. The relevant dimension is not skill level but the (α_{iH}, σ_i) architecture of the worker’s task bundle. A low-wage care worker whose bundle is heavily \mathcal{C} -mixed (therapy, end-of-life support, teaching under relational closeness) is less substitutable by AI than a high-wage analyst whose bundle is heavily pure-D or \mathcal{S} -mixed (standardized forecasting, bench-marked drafting). The AI economy does not simply favor the educated; it favors task bundles with high \mathcal{H}^* -weight and within-task complementarity, and these two groups do not coincide with the educated–less-educated partition.

Proof of Proposition C.1 (Automation threshold)

Part (i) (weak monotonicity of the AI-served set, generic). We prove the weaker claim that generically $\mathcal{A}(p_s)$ expands as p_s falls; this claim holds when wages are taken as fixed, but can fail under full equilibrium wage feedback. Fix for this argument the equilibrium wage vector $(w_D(p_s), w_H(p_s))$ as a function of p_s . For $i \in \mathcal{H}^{*c}$, by Proposition 1(ii), $q_i^{AI}(p_s)$ depends on p_s only through $s_D^{AI}(p_s) = \bar{s}_D(1 - e^{-\gamma_D X_D^{AI} S/p_s})$, which is strictly increasing in $-\log p_s$; hence $q_i^{AI}/(p_s c_i)$ is strictly increasing in $-\log p_s$. On the right-hand side, q_i^H is bounded by skill saturation (A2); $w(i) = w_D s_D^H + w_H s_H^H$ evolves with equilibrium (w_D, w_H) . On pure-D tasks ($\alpha_{iH} = 0$), $w(i) = w_D \cdot s_D^H$, and w_D is non-monotone (Proposition 3) but bounded, so the RHS is bounded; the LHS eventually dominates, giving strict expansion for small p_s . On \mathcal{S} -mixed tasks, $w(i) = w_D s_D^H + w_H s_H^H$ is uniformly bounded above across all $p_s \in (0, p_s^{\text{pre-AI}}]$ because: (a) $w_H(p_s) \leq w_H^\infty \leq 1/(\bar{L} \bar{s}_H^H)$ uniformly by

Corollary 2; (b) $w_D(p_s)$ is uniformly bounded by the factor-payment identity $w_D L_D + w_H L_H \leq 1$ combined with the permanent human employment floor on \mathcal{C} -mixed tasks (see Part (ii) below); (c) s_D^H, s_H^H are bounded by skill saturation (A2). Hence $w(i) \leq \bar{w} < \infty$ uniformly, so the RHS of the assignment condition is uniformly bounded while the LHS grows as $1/p_s$; LHS dominates for p_s small enough. *Strict monotonicity* of $\mathcal{A}(p_s)$ in p_s for every i is not proven; non-monotone wage feedback could in principle cause tasks to exit and re-enter \mathcal{A} over a subinterval. What is proven is asymptotic convergence (Part (ii)). \square

Part (ii) ($\mathcal{A}(p_s) \rightarrow \mathcal{H}^{*c}$ as $p_s \rightarrow 0$). For $i \in \mathcal{H}^{*c}$, $q_i^{AI}(p_s) \rightarrow s_D^{AI}(p_s) \cdot g_i$ where $g_i = 1$ on pure-D tasks and $g_i = \alpha_{iD}^{\sigma^S/(\sigma^S-1)} > 0$ on \mathcal{S} -mixed tasks. As $p_s \rightarrow 0$, $s_D^{AI} \rightarrow \bar{s}_D < \infty$, so $q_i^{AI} \rightarrow \bar{s}_D g_i < \infty$, and the assignment-ratio LHS $q_i^{AI}/(p_s c_i) \rightarrow \infty$ at rate $1/p_s$. On the RHS, $q_i^H \leq q_i^H(\bar{X}(T))$ is bounded by A3. For $w(i) = w_D s_D^H + w_H s_H^H$, Corollary 2 establishes that $w_H(p_s)$ converges to a finite limit w_H^∞ as $p_s \rightarrow 0$ (under numeraire $PY = 1$); for $w_D(p_s)$, the factor-payment identity $w_D L_D + w_H L_H \leq PY = 1$ combined with the permanent human employment floor on \mathcal{C} -mixed tasks (which have positive measure by A1 and use human D-skill with $\alpha_{iD} > 0$) gives a strictly positive lower bound $\underline{L}_D > 0$ on total human D-labor demand at every p_s ; hence $w_D(p_s) \leq 1/\underline{L}_D < \infty$ uniformly. Correspondingly, $w(i)$ is bounded above by $\bar{w} = (1/\underline{L}_D)\bar{s}_D^H + w_H^\infty \bar{s}_H^H < \infty$, independent of p_s . Further, $w(i)$ is bounded below by $\underline{w} \equiv \underline{w}_H s_H^H > 0$ where \underline{w}_H is the positive lower bound on w_H from A1 and A4 (\mathcal{C} -mixed tasks have $\alpha_{iH} \in (0, 1)$ bounded away from 0, guaranteeing a strictly positive H-skill cost share and a Baumol floor on w_H at every $p_s > 0$). Hence $q_i^H/w(i) \leq q_i^H(\bar{X}(T))/\underline{w} < \infty$ uniformly. The LHS grows unboundedly while the RHS stays bounded, so for any $i \in \mathcal{H}^{*c}$ there exists $p_s(i) > 0$ such that the assignment condition holds for all $p_s < p_s(i)$. Hence $\mathcal{A}(p_s) \rightarrow \mathcal{H}^{*c}$ in measure as $p_s \rightarrow 0$. \square

Part (iii) ($\mu(\mathcal{H}^*)$ invariant and positive). \mathcal{H}^* is defined in (8) purely in terms of (α_{iH}, σ_i) ; its measure is a property of F alone, independent of p_s . Positivity follows from A1: F assigns positive measure to \mathcal{C} -mixed tasks, which constitute \mathcal{H}^* . For $i \in \mathcal{H}^*$, $q_i^{AI} = 0$ by Proposition 1(ii), so $i \notin \mathcal{A}(p_s)$ at every $p_s > 0$. \square

Proof of Proposition C.2 (Wage effects)

Part (i) (w_D non-monotone). This is exactly Proposition 3; its proof is given in Appendix C. The non-monotonicity follows from (15): \dot{R} is the difference $\mu_C \dot{R}_C - |\dot{R}_S|$, and the threshold $\mu_C^*(p_s) \rightarrow 0$ as reassignment completes forces the eventual sign of \dot{R} to be positive. \square

Part (ii) (w_H monotone). This is exactly Corollary 2; its proof is given in Appendix C. Under Assumption A4 with $\eta < 1$ strict, the Baumol channel is strict; the within-human complementarity channel on \mathcal{C} -tasks reinforces it (Case (i) of Proposition 3) or does not overturn it (Case (ii)), as shown in the proof of Corollary 2. \square

Part (iii) (no across-agent complementarity premium). Under one-agent-per-task, a “collaboration premium” would require a single task to receive inputs from both agents, which is ruled out by the indivisibility primitive in Section 3.2. The empirical phenomenon previously called a complementarity premium operates across tasks: AI absorption of \mathcal{S} -mixed and pure-D tasks leaves workers specializing in \mathcal{C} -mixed tasks, where both their s_D^H and s_H^H earn rising returns via the within-human cross-partial (12) and the Baumol channel on w_H . \square

Proof of Proposition C.3 (Investment reallocation)

By the Ben-Porath FOC (11), optimal investment equalizes marginal value across bundles: $w_D \cdot (s_D^H)'(X_D^{H*}) = w_H \cdot (s_H^H)'(X_H^{H*})$ at the interior optimum.

Part (i) (X_H^{H} rising at the endpoint).* $w_H^\infty > w_H^{\text{pre-AI}}$ (Corollary 2). Holding $(s_H^H)'$ strictly decreasing in X_H^H (concavity, A2), the FOC $w_H \cdot (s_H^H)'(X_H^{H*}) = \mu$ yields X_H^{H*} strictly larger at the post-AI endpoint than at the pre-AI endpoint; the interior path of $X_H^{H*}(p_s)$ inherits the continuity and endpoint-strict structure of $w_H(p_s)$. \square

Part (ii) (X_D^{H} non-monotone).* w_D can be non-monotone in p_s (Proposition 3). The FOC $w_D \cdot (s_D^H)'(X_D^{H*}) = \mu$ implies that X_D^{H*} inherits the sign pattern of w_D : in the S-dominated phase X_D^{H*} falls; in the C-dominated phase it rises. When the transition crosses the threshold \bar{p}_s of Proposition 3(iii), X_D^{H*} reverses direction; a worker who entered early in the transition optimally

accumulates less than one who enters later, inheriting the U-shape. \square

Part (iii) (horizon-bounded capacity). Affective and co-constructive sub-components of the H-bundle respond to deliberate investment (positive argument in the skill function). Mortality-grounded authority, in contrast, accrues constitutively through lived time and has no investment argument; its accumulation is bounded by calendar time rather than by the Ben-Porath budget. Workers targeting occupations whose H-bundle is mortality-weighted therefore face transition paths that deliberate investment cannot accelerate. \square

Proof of Proposition C.4 (Polarization)

Parts (i) and (ii). Let occupation o have a task bundle summarized by the occupation-level share $\beta_o^* \equiv \int \mathbf{1}\{i \in \mathcal{H}^*\} dF_o(i)$ of \mathcal{H}^* -intensive tasks. Under one-agent-per-task, the worker in occupation o earns on $\mathcal{H}^* \cap F_o$ throughout and on $\mathcal{H}^{*c} \cap F_o$ only while the corresponding tasks remain human-served. Total occupation earnings decompose as

$$E_o(p_s) = \beta_o^* \cdot [w_D \cdot \bar{s}_D^H + w_H \cdot \bar{s}_H^H]_{\mathcal{H}^*} + (1 - \beta_o^*) \cdot \mathbb{E}_{i \in \mathcal{H}^{*c}} [\mathbf{1}\{i \notin \mathcal{A}(p_s)\} \cdot \text{human earnings on } i].$$

The first term is strictly increasing in $-\log p_s$ via both channels (Corollary 2: w_H rises; within-task complementarity: return to s_D^H on \mathcal{C} -tasks rises). The second term falls strictly as $\mathcal{A}(p_s)$ expands. For high β_o^* the first term dominates; for low β_o^* the second. \square

Part (iii). Let $T_1(p_s)$ denote the *region-average* of human earnings on \mathcal{H}^* (revenue per unit \mathcal{H}^* -measure, integrated over the α_{iH} density and skill profile under A1, A2), and $T_2(p_s)$ the corresponding region-average of human earnings on \mathcal{H}^{*c} (with the indicator $\mathbf{1}\{i \notin \mathcal{A}(p_s)\}$ folded into the integration). Both are well-defined region-averages, not pointwise quantities:

$$T_1(p_s) \equiv \frac{1}{\mu(\mathcal{H}^*)} \int_{\mathcal{H}^*} (\text{human earnings on } i) di,$$

$$T_2(p_s) \equiv \frac{1}{\mu(\mathcal{H}^{*c})} \int_{\mathcal{H}^{*c}} \mathbf{1}\{i \notin \mathcal{A}(p_s)\} \cdot (\text{human earnings on } i) di.$$

Occupation earnings then decompose as $E_o = \beta_o^* T_1 + (1 - \beta_o^*) T_2$ where β_o^* is the occupation's share of \mathcal{H}^* -intensive tasks, and the cross-sectional gap between occupations o^h, o^l with $\beta^h > \beta^l$ is

$$\text{Gap}(p_s) \equiv E_{o^h} - E_{o^l} = (\beta^h - \beta^l) \cdot [T_1(p_s) - T_2(p_s)].$$

Under one-agent-per-task and the Walras normalization $PY = 1$, the factor-payment identity at pre-AI ($\mathcal{A} = \emptyset$) gives $T_1^{\text{pre}} \mu(\mathcal{H}^*) + T_2^{\text{pre}} \mu(\mathcal{H}^{*c}) = 1$, while at the asymptotic limit humans receive all revenue on \mathcal{H}^* (where $q_i^{\text{AI}} = 0$) and zero on \mathcal{H}^{*c} (where $\mathcal{A}(p_s) \rightarrow \mathcal{H}^{*c}$), so $T_1^\infty \mu(\mathcal{H}^*) = s_{H^*}^\infty = 1$ (Corollary 2 Step 3) and $T_2^\infty = 0$. Combining,

$$T_1^\infty - T_1^{\text{pre}} = \frac{T_2^{\text{pre}} \mu(\mathcal{H}^{*c})}{\mu(\mathcal{H}^*)} > 0$$

strictly (since $T_2^{\text{pre}} > 0$ by pre-AI positive human earnings on \mathcal{H}^{*c} , and $\mu(\mathcal{H}^{*c}) > 0$ by A1). The *change* in the cross-sectional earnings gap is therefore

$$\Delta \text{Gap} \equiv \text{Gap}^\infty - \text{Gap}^{\text{pre}} = (\beta^h - \beta^l) \cdot [T_1^\infty - T_1^{\text{pre}} + T_2^{\text{pre}}] = \frac{(\beta^h - \beta^l) \cdot T_2^{\text{pre}}}{\mu(\mathcal{H}^*)} > 0,$$

strictly positive whenever $\beta^h > \beta^l$, with magnitude equal to $T_2^{\text{pre}}/\mu(\mathcal{H}^*)$ per unit of β^* -differential — a quantity determined entirely by pre-AI human earnings on the contestable set and the \mathcal{H}^* -measure, neither of which depends on the D-wage path (thereby sidestepping the case-dependent sign of $(w_D^\infty - w_D^{\text{pre-AI}})$ in Proposition 3). The *level* of the post-AI gap equals $(\beta^h - \beta^l) \cdot T_1^\infty = (\beta^h - \beta^l)/\mu(\mathcal{H}^*)$, strictly positive whenever the β^* -difference is positive. The resulting occupation-level distribution of $E_o(p_s)$ bifurcates along β_o^* . \square

Part (iv). The Autor, Levy, and Murnane (2003) task-content mechanism predicts polarization on the skill-complexity dimension (U-shape across the skill distribution). The present framework predicts polarization on the (α_{iH}, σ_i) dimension, specifically on β_o^* , which is orthogonal to measured skill level. Low-wage \mathcal{C} -mixed occupations (residential care, teaching, clinical social work) gain; high-wage \mathcal{S} -mixed or pure-D occupations (routine drafting, bench-marked forecast-

ing, professional writing to template) lose. These are distinct empirical signatures testable against occupational panel data. \square

D. Portfolio Concentration under Inner CES

The following local-curvature result applies to any bundle with an inner CES nest over two sub-dimensions A and B . In the D -bundle case used in Section 5.4, the result governs whether a worker optimally balances or concentrates her D -bundle investment across sub-dimensions such as front-end and back-end programming, tax and corporate law, or diagnostic and procedural medicine.

Fix the total D -bundle investment $\ell \equiv X_D^H = d_A + d_B$, with the worker's D -investment split between sub-dimensions A and B in shares $\phi \in [0, 1]$ and $1 - \phi$. The worker takes the task wage schedule as given and chooses ϕ to maximize income.

Let $r \equiv (\rho - 1)/\rho$ and $f(\phi, \ell) \equiv \lambda_A s_A (\phi \ell)^r + \lambda_B s_B ((1 - \phi)\ell)^r$ so that the composite D -bundle skill is $s_D^H(\phi, \ell) = f^{1/r}$.

Local curvature at symmetry. Differentiating once,

$$\frac{\partial s_D^H}{\partial \phi} = \frac{1}{r} f^{1/r-1} \cdot f'(\phi, \ell), \quad f'(\phi, \ell) = r\ell \left[\lambda_A s_A (\phi \ell)^{r-1} s'_A(\phi \ell) - \lambda_B s_B ((1 - \phi)\ell)^{r-1} s'_B((1 - \phi)\ell) \right].$$

Under symmetry ($\lambda_A = \lambda_B \equiv \lambda$, $s_A = s_B \equiv s$, $s'_A = s'_B \equiv s'$), $f'|_{\phi=1/2} = 0$, confirming $\phi = 1/2$ as an interior critical point. Differentiating again and evaluating at $\phi = 1/2$ (so that $f' = 0$),

$$\left. \frac{\partial^2 s_D^H}{\partial \phi^2} \right|_{\phi=1/2} = \frac{1}{r} f^{1/r-1} \cdot f''(\phi, \ell) \Big|_{\phi=1/2}.$$

A direct computation gives

$$f''(\phi = 1/2, \ell) = 2\lambda r \ell^2 s^{r-1} \left[s'' + (r - 1) \frac{(s')^2}{s} \right].$$

Substituting $r - 1 = -1/\rho$ and collecting,

$$\left. \frac{\partial^2 s_D^H}{\partial \phi^2} \right|_{\phi=1/2} = 2\lambda \ell^2 f^{1/r-1} s^{r-1} \left[s'' - \frac{1}{\rho} \frac{(s')^2}{s} \right].$$

The r -factor from f'' cancels the $1/r$ prefactor from the chain rule. The bracket $[s'' - \rho^{-1}(s')^2/s]$ is strictly negative for all $\rho > 0$, since $s'' < 0$ by concavity of the skill function and $(s')^2/s > 0$. Therefore $\phi = 1/2$ is always a local maximum of s_D^H , for all $\rho > 0$.

Numerical verification. For $s(x) = s_0(1 - e^{-\gamma x})$, $s_0 = \gamma = \ell = 1$: at $\rho = 1/2$, Taylor expansion around $\phi = 1/2$ gives $s_D^H(1/2 + \varepsilon) \approx s_0 - s_0\varepsilon^2 + O(\varepsilon^4)$; at $\rho = 2$, $s_D^H(1/2 + \varepsilon) \approx s_0 - (s_0/4)\varepsilon^2 + O(\varepsilon^4)$. Both confirm $\phi = 1/2$ as a local max.

Proposition D.1 (Portfolio concentration). *Consider a worker with inner-CES D -bundle $s_D^H(\phi, \ell) = f(\phi, \ell)^{1/r}$ as defined above and total D -investment $\ell = X_D^H$ fixed, choosing portfolio share $\phi \in [0, 1]$ to maximize income at equilibrium task prices $\{p_i(p_s)\}$. Assume **(HI)**: the task weight distribution F assigns positive measure to tasks i with $|\lambda_A(i) - \lambda_B(i)| > \delta$ for some $\delta > 0$.*

- (i) ($\rho < 1$.) *The interior portfolio $\phi = 1/2$ is a global maximum of the worker's wage function on \mathcal{C} -mixed tasks, and the margin of preference for interior over corner portfolios strengthens as p_s falls.*
- (ii) ($\rho > 1$.) *Under HI, there exists $p_s^* > 0$ such that for all $p_s < p_s^*$, a corner portfolio ($\phi \in \{0, 1\}$) strictly dominates the interior $\phi = 1/2$ in equilibrium wage. Without HI — if F places all mass on tasks with $\lambda_A = \lambda_B$ — corner and interior portfolios deliver equal value in the limit and ϕ^* is indeterminate.*
- (iii) ($\rho = 1$.) *The optimal portfolio is pinned down by the sub-dimension weights (λ_A, λ_B) and is invariant to p_s .*

Proof. The proof has three parts, one per regime. The local-curvature computation above establishes $\phi = 1/2$ as a local maximum of s_D^H for all $\rho > 0$; what varies across regimes is whether the interior is also a *global* maximum of the wage function.

Part (i) [$\rho < 1$: interior optimum]. The local curvature result establishes $\phi = 1/2$ as a local maximum of s_D^H . We show this is also a *global* maximum of worker wages for $\rho < 1$, and that the pull toward $1/2$ strengthens as p_s falls.

When $\rho < 1$, the inner CES nest exhibits supermodularity in (s_A, s_B) : $\partial^2 s_D^H / \partial s_A \partial s_B > 0$ for $\rho \in (0, 1)$ (see, e.g., Mas-Colell, Whinston, and Green, 1995). This is the Leontief-leaning regime in which the two D -sub-dimensions are productive complements within the task.

Consider a corner $\phi = 1$: the worker has $s_A(\ell)$ in sub-dimension A and 0 in sub-dimension B . On \mathcal{C} -mixed tasks (the operative locus for the portfolio-composition question in Section 5.4), the human's D -skill is complementary within the task with her H -skill; the aggregate D -skill entering the task CES is $s_D^H(\phi, \ell)$, computed via the inner nest. As p_s falls, AI saturates the D -bundle components it supplies on neighboring contestable tasks, and the marginal value of the worker's residual D -role — integrating both sub-dimensions within a \mathcal{C} -mixed task — rises. At $\phi = 1/2$ the worker supplies coverage on both sub-dimensions; at $\phi = 1$ the worker's sub-dimension B contribution is zero, which under $\rho < 1$ collapses the inner CES term on B and delivers a strictly lower s_D^H than interior allocation.

The wage function $w(\phi, p_s) = \int p_i(p_s) \cdot (\partial q_i / \partial s_D^H) \cdot (\partial s_D^H / \partial d_k) di$ therefore strictly prefers the interior $\phi = 1/2$ to any corner under $\rho < 1$, and the strength of this preference grows as p_s falls because AI's absorption of \mathcal{S} -mixed and pure- D tasks concentrates the worker's remaining D -role on \mathcal{C} -mixed tasks where the within-task complementarity binds. \square

Part (ii) [$\rho > 1$: specialization]. The local curvature result establishes $\phi = 1/2$ as a local maximum of s_D^H also for $\rho > 1$. The proposition's claim — that under H1, for sufficiently small p_s , corner wages exceed the interior wage — is therefore *not* a local second-order statement but a *global* comparison of the wage function $w^*(\phi, p_s) \equiv \int p_i(p_s) \cdot (\partial q_i / \partial s_D^H) \cdot (\partial s_D^H / \partial d_k(\phi)) di$ across portfolios.

Sketch via an explicit two-task example. Consider a binary support: half of \mathcal{C} -mixed tasks have $(\lambda_A, \lambda_B) = (1, 0)$ and the other half $(\lambda_A, \lambda_B) = (0, 1)$, with all other primitives symmetric. Under $\rho > 1$, the inner-CES collapses to $s_D^H(\phi, \ell)$ that on A -only tasks equals $s_A(\phi \ell)$ and on B -only tasks

equals $s_B((1 - \phi)\ell)$ (the contribution of the unused sub-dimension is zero on a task that weights it not at all). At $\phi = 1$ the worker's D -skill is $s_A(\ell)$ on A -tasks and 0 on B -tasks; symmetrically at $\phi = 0$. At $\phi = 1/2$ the worker has $s_A(\ell/2)$ on A -tasks and $s_B(\ell/2)$ on B -tasks. Concavity of s gives $s(\ell) > 2s(\ell/2)$ for the exponential form, so $s_A(\ell) + 0$ on the half of tasks weighted toward A exceeds the symmetric portfolio's contribution there; humans matter only on C -mixed tasks where AI's contribution is zero; and on the other half of tasks the corner worker earns zero while the symmetric worker earns $s_B(\ell/2)$. The ratio depends on how rapidly AI absorbs the S -mixed and pure- D regions; for p_s small enough that $\mathcal{A}(p_s) \approx \mathcal{H}^{*c}$ and the worker's earnings are concentrated on C -mixed tasks, the corner allocation dominates because $s_A(\ell) > s_A(\ell/2) + s_B(\ell/2)$ when one mixes at the cost of severely concave underlying skill functions. The general claim of Part (ii) extends this two-mass example to any distribution satisfying H1 — positive mass at $|\lambda_A - \lambda_B| > \delta$.

The argument is a sketch, not a closed-form proof; we leave the general comparison to future work. \square

Part (iii) [$\rho = 1$: Cobb-Douglas]. At $\rho = 1$: $s_D^H(\phi, \ell) = s_A(\phi\ell)^{\lambda_A} s_B((1 - \phi)\ell)^{\lambda_B}$. The interior FOC gives $\lambda_A s'_A(\phi\ell)/s_A(\phi\ell) = \lambda_B s'_B((1 - \phi)\ell)/s_B((1 - \phi)\ell)$, or equivalently $\lambda_A \cdot e_A(\phi\ell) = \lambda_B \cdot e_B((1 - \phi)\ell)$ where $e_k(x) \equiv x s'_k(x)/s_k(x)$ is the elasticity of skill with respect to investment. Since e_k depends only on d_k through the human skill function (3) and is independent of p_s , the solution ϕ^* is determined solely by $(\lambda_A, \lambda_B, \ell)$, invariant to p_s . \square

E. Robustness Figures

This appendix collects the five non-compositional sensitivities summarised in Section 6.2. Each figure perturbs a single parameter around the benchmark parameterization of Section 6.1, holding all other features fixed. In every panel the horizontal axis is $\tau = -\log p_s$, so AI capability grows left to right.

The three compositional axes (μ_C , the α - σ alignment, and within-regime shape concentration) appear inline in Section 6 as Figures 7, 5, and 6. The five non-compositional axes are reproduced

Table 2: Benchmark endpoint values. Pre-AI: $p_s = 4.0$. Post-AI endpoint: $p_s = 0.02$. Real wages w/P normalised to pre-AI = 1.

Quantity	Pre-AI	Post-AI endpoint
$\mu(\mathcal{A}(p_s))$	0.00	0.50
w_H/P	1.00	4.52
w_D/P	1.00	3.28
w_D/w_H	1.00	0.72
X_H^{H*}	0.50	0.58
P	2.00	0.44

below. Across all five, the directional predictions of Section 5 retain their sign. The endpoint magnitudes shift non-trivially with the outer CES elasticity η , modestly with the inner elasticities (σ^C, σ^S) , and are second-order in the remaining three axes.

Table 2 reports the benchmark endpoint values that anchor the perturbations.

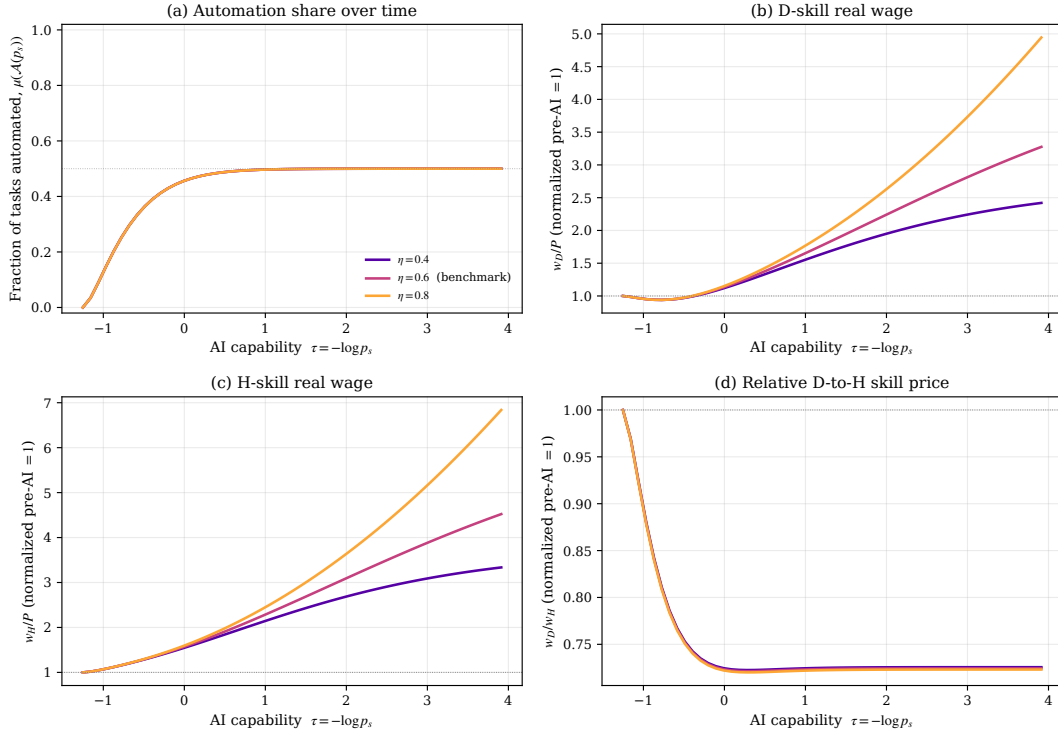


Figure 8: Sensitivity to the outer CES elasticity $\eta \in \{0.4, 0.6, 0.8\}$. Direction-preserving; endpoint w_H/P moves across roughly $[3.3, 6.8]$. This is the second-largest magnitude dial after μ_C ; smaller η (stronger Baumol) produces a larger endpoint H -wage.

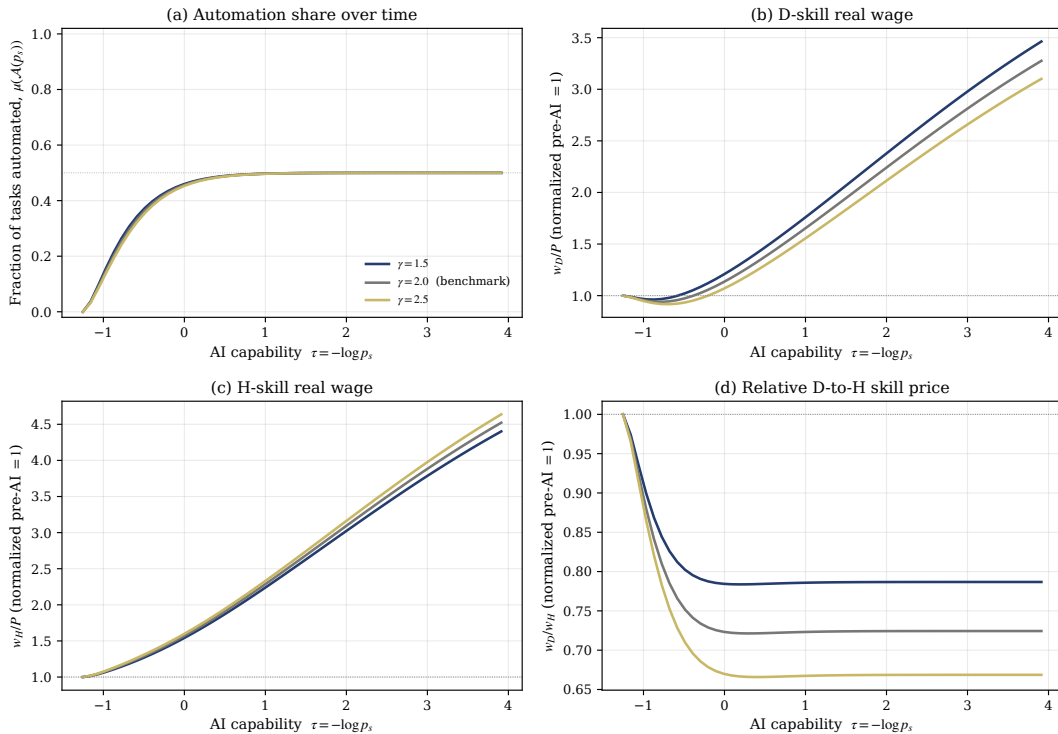


Figure 9: Sensitivity to the Ben-Porath curvature $\gamma \in \{1.5, 2.0, 2.5\}$. Direction-preserving; small magnitude sensitivity.

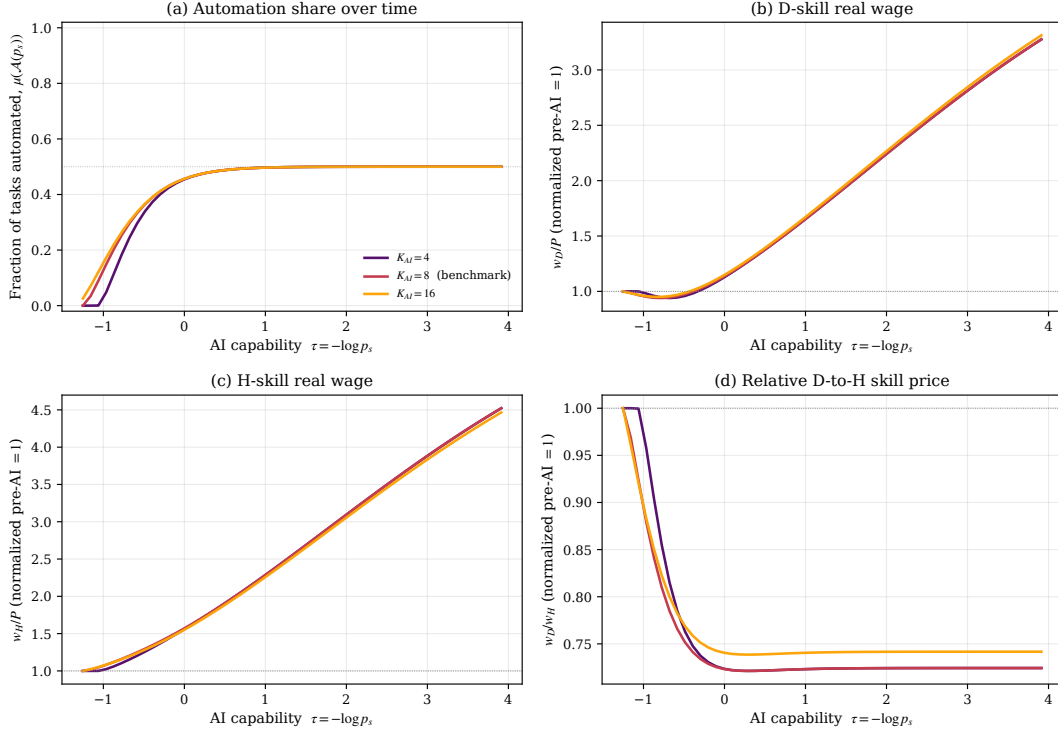


Figure 10: Sensitivity to the AI scale $K^{AI} \in \{4, 8, 16\}$. Direction-preserving; magnitude sensitivity is negligible at the endpoint because s_D^{AI} saturates at its ceiling \bar{s}_D^{AI} for all three values.

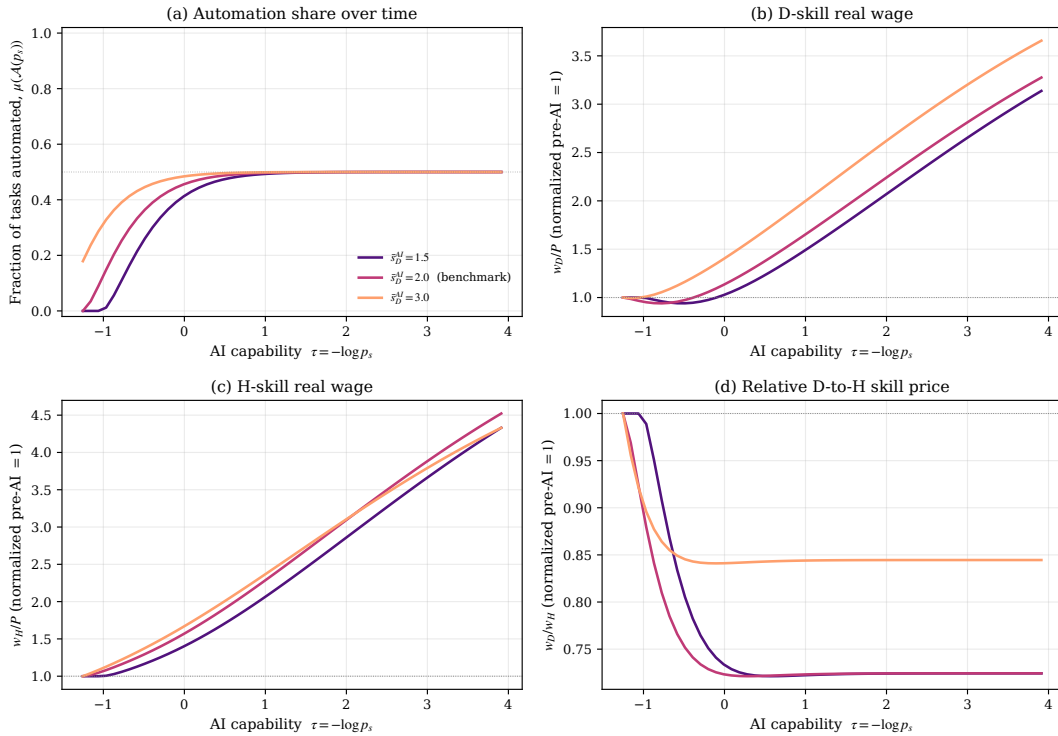


Figure 11: Sensitivity to the AI saturation level $\bar{s}_D^{AI} \in \{1.5, 2.0, 3.0\}$. Direction-preserving; small magnitude sensitivity.

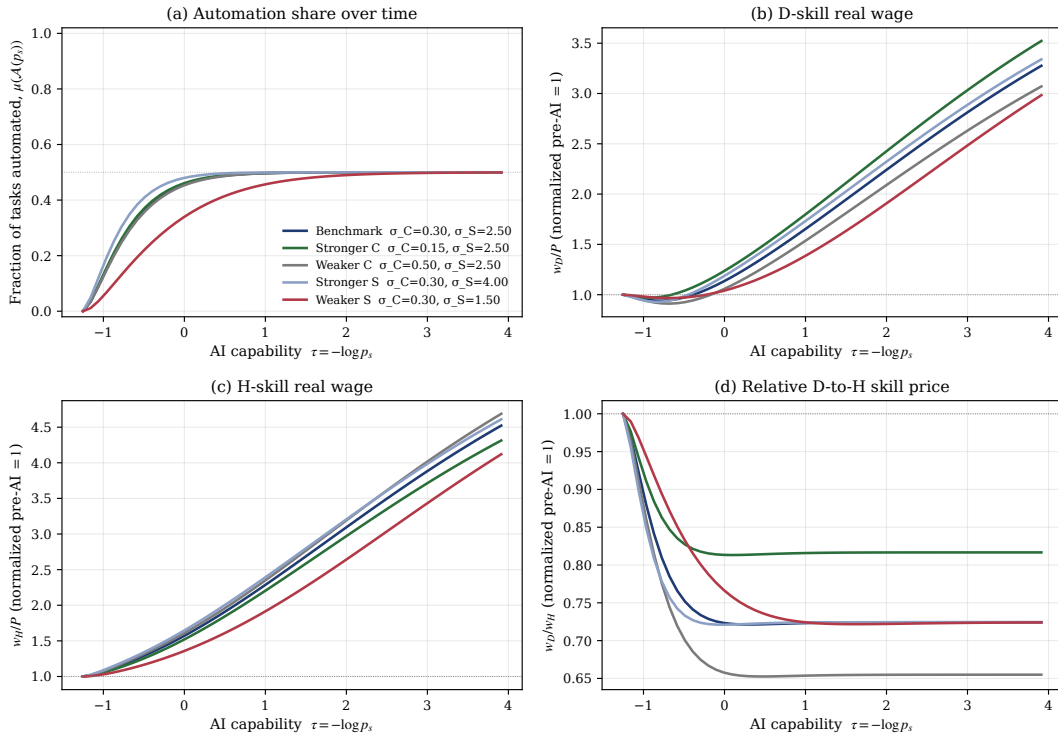


Figure 12: Sensitivity to the inner within-task elasticities σ^C and σ^S under ± 0.1 perturbations around the benchmark $(\sigma^C, \sigma^S) = (0.30, 2.50)$. Direction-preserving; small-to-moderate sensitivity on the wage ratio, small on the levels.