FEDERAL RESERVE BANK *of* ATLANTA **WORKING PAPER SERIES**

Automation, Learning, and Career Dynamics

Hassan Afrouzi, Andrés Blanco, Andres Drenik, and Erik Hurst

Working Paper 2026-6
May 2026

Abstract: We study how an automating technology affects career dynamics, human capital, and welfare in an economy where workers acquire skill through the tasks they perform. In a continuous-time general equilibrium model, learning-by-doing is determined jointly with the share of tasks automated, the frontier of tasks managers maintain, and the worker-to-manager career transition. Economies with high learning capacity admit pairs of stationary equilibria strictly ranked by the aggregate learning rate. Cheaper technology has opposite effects across the two: in the high-learning equilibrium, it raises welfare *through* the learning channel itself; in the low-learning equilibrium, it tips the economy into a human-capital trap. The planner's first-best combines a tax on automation profits with a subsidy on frontier maintenance expenditures at a common rate.

JEL classification: E23, E24, J24

Key words: human capital, learning-by-doing, automation, AI

https://doi.org/10.29338/wp2026-06

Hassan Afrouzi is with Columbia University. Andrés Blanco is with the Federal Reserve Bank of Atlanta. Andres Drenik is with the University of Texas. Erik Hurst is with the University of Chicago Booth School of Business. The authors thank seminar participants at the University of Pennsylvania and the 2026 Pacific Macro Conference for useful comments.

"Learning is the product of experience. Learning can only take place through the attempt to solve a problem and therefore only takes place during activity."... *Kenneth Arrow (1962)*

1 Introduction

A defining regularity across a broad set of white-collar occupations is that senior workers began as junior workers who accumulated relevant expertise by executing domain specific entry level tasks. Senior lawyers advise clients only after years of reviewing contracts and researching precedents; data scientists develop the implicit knowledge about data quality and model reliability only after working extensively with such data and models; management consultants build the judgment to advise senior leadership in the companies who hire them only by doing the analytical legwork through which they learn both institutional details and the tacit craft of advising. The tasks that fill entry-level positions are not merely low-value work — they are the curriculum through which workers accumulate the human capital that makes them productive later in their careers. In the tradition of Arrow (1962a) and Stokey (1988a), learning is the product of experience, and in these occupations, it operates through the tasks workers actively perform rather than through instruction.[1]

Some technologies, like generative artificial intelligence (AI), can affect on-the-job human capital accumulation through two opposing effects. First, a reduction in the price of this technology may incentivize firms to automate a higher share of existing entry-level tasks narrowing the pipeline through which workers acquire skill. Second, the technology may also complement the human capital of existing managers. In this case, a fall in the price of the technology may help managers expand the task frontier increasing opportunities for entry-level workers to learn. A decline in the price of such technology drives both effects within a *given* worker's career: it substitutes for workers on entry-level tasks *and* complements senior workers in creating the new tasks the occupation performs. Systematic evidence has already begun to line up with both sides of this duality: entry-level hiring has declined sharply in AI-exposed white-collar occupations since late 2022, while employment for older workers in the same occupations has remained

[1]Models with learning-by-doing are common in the economic growth literature; see, for example, Lucas (1988b), Chari and Hopenhayn (1991), Young (1991a), Young (1993a), and Jovanovic and Nyarko (1996). The view that learning comes from experience is also pervasive in education and cognitive psychology literature; see Dewey (1938), Anderson (1982), and Ericsson, Krampe, and Tesch-Römer (1993) for canonical references.

stable or grown (Brynjolfsson, Chandar, and Chen, 2025a); and firms that invest heavily in AI grow primarily through product innovation rather than labor replacement (Babina, Fedyk, He, and Hodson, 2024).[2]

Yet both the measure of tasks left to workers to learn from and the frontier of new tasks the economy creates are objects whose effect on human capital accumulation are not internalized by atomistic workers or firms. If the human capital is general rather than firm-specific, a given atomistic firm will not internalize how their automation and task creation decisions will affect worker learning. Likewise, atomistic workers cannot credibly offer a wage concession large enough to change any firm's automation decision; the two channels therefore generate a pecuniary externality that neither individual firms or individual workers has the leverage to close.[3] The central question of this paper is how the introduction of a technology with these features — one that simultaneously substitutes for entry-level workers and complements senior workers in creating new tasks, both through channels that neither firms or workers internalize — affects career dynamics, human capital accumulation, and welfare in the economy. The answer, we argue, depends on which of two stationary equilibria the economy coordinates on: cheaper technology can be welfare-improving or welfare-destroying on the same primitives.

We study this question in a continuous-time general equilibrium framework where the extent of *learning-by-doing* is determined in interaction with three key ingredients: endogenous automation, an endogenous frontier of tasks, and endogenous career choice with human capital accumulation. Each of these ingredients captures a distinct channel through which the technology interacts with learning. Endogenous automation mechanically takes tasks away from

[2] Industry leaders have described this duality in stark terms. Anthropic's CEO Dario Amodei has claimed that AI could wipe out half of all entry-level white-collar jobs over the next few years (Cooper, 2025); OpenAI's CEO Sam Altman, in turn, has argued that continued AI advances will make "figuring out what questions to ask even more important than figuring out the answer" (Grant, 2025). See Roose (2025) and Telford (2026) for broader discussion in the *New York Times* and *Washington Post*, respectively. Fossen, McLemore, and Sorgner (2024) surveys the growing literature on AI and entrepreneurial activity, and Brynjolfsson, Chandar, and Chen (2025b) presents preliminary evidence that AI is already reducing employment in entry-level white-collar jobs.

[3] This externality is distinct from the bilateral-contracting friction in the Becker worker training problem and its AI-era variants (Garicano and Rayo, 2017a, 2025a, Ide, 2025a). In those frameworks, an individual worker can pay for training by accepting a lower wage, so the relevant friction is the extent to wages can adjust. In our setting, the relevant friction is the coordination required among many workers and firms to internalize the learning rate, not the bilateral wage-setting problem. To that end, Shen and Tamkin (2026) provides direct experimental evidence that AI usage among participants reduced their accumulation of skills needed for career progression. By randomly assigning workers to use AI, they offer preliminary evidence for the learning-by-doing mechanism at the heart of our analysis.

workers — the tasks through which they would have learned — and the automation share is determined in equilibrium by the worker wage and the technology price. Similar to Acemoglu and Restrepo (2022a), workers and the technology are perfect substitutes at the task level, generating a cutoff rule: tasks below some threshold are performed by the technology, while the rest are performed by workers.[4]

The endogenous frontier of tasks adjusts to the technology and gives the economy the opportunity to create new tasks for workers to learn from.[5] Specifically, managers combine their accumulated human capital with the technology to maintain the frontier, and a cheaper technology lowers the unit cost of the frontier and, for any given supply of managerial skill, sustains a wider range of tasks the occupation performs.[6] Endogenous career choice and human capital accumulation govern the labor supply going to each of these roles: households start their careers as workers, accumulate human capital at a rate proportional to the measure of worker-performed tasks, and transition permanently to management at an human-capital threshold of their choice that is shaped by the learning rate, the worker wage, and the manager wage premium. Together, the three ingredients are linked through a single technology price and a single economy-wide learning rate, and it is this link that is at the heart of the paper's results.

Our first result concerns the structure of stationary equilibria. We show that in economies where the learning channel is not strong — i.e., the learning capacity of the economy, governed by the intensity of on-the-job learning, the potential scale of the task frontier, and the returns to scale in production of each task, is low — the stationary equilibrium is unique and follows the familiar task-based pattern: no automation at high technology prices, full automation at low prices, and interior automation in between. This case recovers the standard task-based

[4]A large literature studies how such technological shifts reallocate workers *across* occupations; see, for example, Autor (2015), Mokyr, Vickers, and Ziebarth (2015), and Acemoglu and Johnson (2024) on the power loom; Acemoglu and Restrepo (2020a) on industrial robots; and Autor, Levy, and Murnane (2003a) and Goos, Manning, and Salomons (2014) on the polarization generated by computerization. Our paper differs from these in focusing on how a technology can affect career dynamics *within* a given occupation, where the endogeneity of human capital accumulation is central.

[5]See Acemoglu, Autor, and Johnson (2026) for a discussion of why it is important to incorporate both how AI may automate existing tasks and how AI may facilitate the creation of new tasks when examining the effects of AI on labor market outcomes.

[6]The hierarchical structure of workers and managers within a firm has been studied in a large literature including Garicano (2000), Garicano and Rossi-Hansberg (2006), and Caicedo, Lucas, and Rossi-Hansberg (2019). Our framework shares their distinction between line production and managerial judgment but features endogenous human capital accumulation through learning by doing, managers who maintain the task frontier, and a technology that both complements managerial human capital and substitutes for worker human capital.

benchmark in papers such as Acemoglu and Restrepo (2022a) and does not deliver the paper's novel results. However, when the learning capacity is high enough, stationary equilibria — when they exist — come in pairs.[7] Each pair consists of a high-learning equilibrium with limited automation and depressed worker wages, and either a low-learning interior equilibrium with extensive automation and elevated worker wages, or a full-automation equilibrium with zero on-the-job learning. Coexisting equilibria are strictly ranked by the aggregate learning rate: both aggregate real value added and the expected lifetime welfare of an entering household are strictly higher in the high-learning equilibrium. Firm profits, by contrast, are approximately invariant across coexisting equilibria — the additional automation surplus in the low-learning equilibrium approximately offsets the smaller task frontier — so the welfare cost of coordinating on low learning falls almost entirely on workers through the human-capital channel.

Our second result concerns how the coexisting equilibria respond to a decline in the technology price. These two stationary equilibria respond with opposite signs. In the high-learning equilibrium, cheaper technology raises the aggregate learning rate through a virtuous cycle: lower unit cost of the frontier enlarges the range of entry-level tasks, accelerates learning, and produces more skilled managers who sustain a still wider frontier. In the low-learning equilibrium, the direct automation effect dominates the wage channel, and cheaper technology instead pushes the learning rate toward zero. At sufficiently low technology prices, the full-automation zero-learning equilibrium is always admissible. However, in high-capacity economies it coexists with a high-learning interior equilibrium.

Because automation equilibria exist in pairs, the economy cannot smoothly transition between them. If coordination falls on the zero-learning branch, entry-level work is fully automated, the learning pipeline shuts down, and cheap technology does not deliver broad productivity gains; we refer to this scenario as a human-capital trap. However, a more optimistic view emerges from the duality of stationary equilibria in this region: whenever the economy is coordinated on the high-learning equilibrium, cheaper technology raises both aggregate value added and household welfare *through* the learning channel itself, not despite it. Numerical examples confirm that multiplicity, the dual comparative statics, and the trap all survive in our full model which

[7] In economies with sufficiently large learning capacity, stationary equilibria may fail to exist altogether, because workers' option value to continue learning exceeds their effective discount rate.

includes both AI's substitutability with entry level tasks and a positive complementarity between managerial human capital and the technology in the creation of new tasks.

Our third result characterizes the planner's optimal policy response. The competitive economy is distorted in two directions simultaneously — firms automate too many existing tasks and invest too little in maintaining the frontier — and both distortions trace to the same pecuniary externality. This result cautions against the pessimistic view of the effects of technology that would prescribe taxing it uniformly: a blanket tax on the technology price would deter automation, but because the same technology also complements managerial skill in maintaining the task frontier, such a tax would simultaneously deter the technology's positive contribution and could backfire, leaving the under-investment distortion unaddressed. The first-best allocation is implemented instead by two instruments targeting the two distinct *uses* of the technology: a proportional tax on profits from the automation use, and a proportional subsidy on expenditures dedicated to maintaining the frontier. Both instruments take the same rate, pinned down by a single shadow wedge that measures the gap between the social and private marginal value of human capital accumulation, and the rate is largest when the economy is closest to the human-capital trap.

Related Literature. Our paper contributes to three strands of the literature. It is most directly connected to the literature on learning-by-doing and human-capital externalities, dating to Arrow (1962b) and developed in growth models by Romer (1986), Stokey (1988b), Lucas (1988a), Young (1991b), and Young (1993b). The key force in these classical models is a technological spillover in which the aggregate stock of knowledge directly enters the economy's production function; the source of multiple balanced growth paths in Lucas (1988a), for instance, is precisely the external effect of average human capital in goods production. We depart from this tradition by generating multiple stationary equilibria with *no* knowledge spillover in aggregate production. Our externality is pecuniary: the measure of worker-performed tasks determines the economy-wide learning rate, which shapes expected future managerial wages, which in turn pins down the worker's optimal stopping problem for promotion, which feeds back into the measure of tasks firms find profitable to leave to workers. This places us in the family of Acemoglu (1996), Redding (1996), Ciccone and Matsuyama (1996), Galor and Tsiddon (1997), and Hassler and Rodríguez Mora (2000), who all show that pecuniary externalities alone can generate multiple steady states in atomistic economies without appealing to aggregate increasing returns. The novel element in our setting is that the threshold governing equilibrium selection is the career transition from worker to manager, and the feedback loop runs through the endogenous on-the-job learning rate. In addition, we argue that the form of the corrective instrument matters: targeted taxation on the two uses of the technology, rather than a blanket tax on the technology itself, is the natural counterpart to the pecuniary externality we identify.

Our framework also contributes to the task-based theory of automation initiated by Zeira (1998), Autor, Levy, and Murnane (2003b), and Acemoglu and Autor (2011), and developed as a theory of endogenous automation and new-task creation by Acemoglu and Restrepo (2018, 2019, 2020b, 2022b). In each of these frameworks, the upper task frontier is expanded by a separate innovation sector whose inputs respond to factor prices. Our contribution is to endogenize the frontier through a third, distinct channel: managerial human capital. Managers combine their human capital with the same automating technology in a CES aggregator to maintain the new tasks that lie beyond the automation frontier. Because managerial human capital is itself the output of the learning pipeline, automation shocks propagate to the frontier through the career ladder rather than through profit-directed innovation.

Our paper is most directly comparable to three contemporaneous papers that theoretically examine how AI can potentially affect human-capital development over a worker's career: Garicano and Rayo (2025a), Ide (2025a), and Friebel, Huang, Li, Shukla, and Zhang (2026). All three share our core intuition that AI substitutes for entry-level tasks and that becoming a senior requires first performing those tasks earlier in one's career, so that AI adoption threatens to unravel the human-capital pipeline. The first two frame this through the lens of the Becker training problem. Garicano and Rayo (2025a) build on the dynamic contracting model of Garicano and Rayo (2017b) to derive conditions under which the bilateral apprenticeship relationship between a "master" and a "novice" survives the introduction of AI: AI simultaneously raises the novice's outside option (an entry-task floor) and the productivity of the expert at the top. Critically, their model contains no externality: training viability is fully determined by the bilateral surplus relative to onboarding costs. Ide (2025a) embeds a similar training mechanism in an overlapping-generations framework and shows that a liquidity constraint together with a floor on novice compensation creates a market incompleteness: even when the training surplus is positive, novices cannot compensate experts sufficiently to sustain efficient training when the wage floor binds, generating socially excessive automation.[8] Friebel, Huang, Li, Shukla, and Zhang (2026), in turn, study the transitional dynamics of the junior–senior task assignment under AI productivity and learning shocks, generating oscillations between "pyramid" and "diamond" organizational shapes in a model with a fixed task pool. Our framework differs from these papers on three dimensions that together carve a unique coordinate: first, our task frontier is endogenous and jointly produced by managerial human capital and the technology; second, we endogenize the worker–manager transition so that today's automation decisions affect tomorrow's managerial supply through the learning rate; third, our externality is pecuniary and operates in general equilibrium rather than through a bilateral contracting friction.

A further complementary and contemporaneous paper is Acemoglu, Kong, and Ozdaglar (2026), who study how agentic AI affects collective human knowledge in a dynamic cognition-based model. They share with us the structural feature that AI can generate multiple steady states, including one in which a human capital channel shuts down. Our framework differs from theirs in three respects: our model is task-based rather than cognition-based; our externality runs through equilibrium wages and endogenous career transitions rather than through the public good of general knowledge; and our first-best policy is targeted taxation on the two distinct uses of the technology rather than information-design on AI's accuracy.

2 Model

Time is continuous and indexed by $t \geq 0$. The economy is populated by a unit measure of households who each supply one unit of labor inelastically and accumulate human capital, denoted x, through a process described below. Individuals transition from being workers early in their career to managers later in their career. Workers execute "tasks" while managers develop new "tasks". Tasks can be thought of as activities done by the firm and can represent new ideas, new product lines, or new varieties of an existing product. All agents discount the future at rate $\rho > 0$ and face a constant death rate $\chi > 0$; upon death, each household is immediately replaced by a new entrant who starts with the minimum level of human capital. A representative

[8] Ide's learning mechanism, like the apprenticeship model of Garicano and Rayo (2017a), follows a bilateral knowledge diffusion tradition in the spirit of Lucas and Moll (2014) in which experts pass on their knowledge to novices by interacting with them. When experts interact less with novices, less knowledge is passed along.

firm produces a final good from a continuum of differentiated products and transfers the profits to a measure zero set of households.

2.1. Firm's Problem

Atomistic firms decide how many workers to hire, how many managers to hire and their use of technology in production. There is one common technology with price q but the firm can use the technology in two ways. First, the technology can help managers create new tasks. In this way, the technology is a complement with managerial ability. Second, the technology can be used to execute existing tasks. Here, the technology is a substitute with worker human capital. As we discuss in greater detail below, our combined assumptions that (i) firms are atomistic within an industry, (ii) human capital accrued through on-the-job learning is industry specific as opposed to firm specific, and (iii) firms hire many workers simultaneously generates the key externality in our model. Specifically, the atomistic firms do not internalize how their decisions to use technology—either their decision to create new tasks or their decision to automate existing tasks—affect individual human capital accumulation.

Product frontier. The firm produces a set of differentiated products indexed by $\tau \in [0, f(t)]$, where the *product frontier* $f(t)$ is determined jointly by managerial input and technology:

$$f(t) \equiv \bar{\tau} \left[\psi^{\omega^{-1}} a_m(t)^{1-\omega^{-1}} + (1-\psi)^{\omega^{-1}} s(t)^{1-\omega^{-1}} \right]^{\frac{1}{1-\omega^{-1}}} \tag{1}$$

where $\bar{\tau}$ is a constant governing the managerial productivity of the economy, $a_m(t)$ is the firm's demand for managerial technology, $\psi \in [0,1]$ governs the relative weight of technology versus managerial human capital in frontier production, and $\omega > 0$ is the elasticity of substitution between the two inputs. The managerial input $s(t) \equiv \int h(x) dN_m^d(x,t)$ aggregates the human capital of managers, with $N_m^d(x,t)$ denoting the CDF of human capital demanded for management at time t. We assume $h(x) = x^{\xi}$ with $\xi > 1$, so that a manager's productivity in maintaining the frontier is convex in their human capital x.

Three things are worth noting about our specification of the product frontier. First, $a_m(t)$ is how much of the common technology that a firm uses to complement managerial talent in expanding the task frontier. The price per unit of technology is q. Second, the CES specification nests several important special cases. The limiting case $\psi = 0$ yields $f(t) = \bar{\tau}\, s(t)$, in which the frontier depends only on managerial human capital with no role for technology.

Production technology. Total output of the firm aggregates across all products within the frontier:

$$y(t) = \int_0^{f(t)} y(\tau, t)^{1-\eta^{-1}} d\tau, \quad \eta > 1 \tag{2}$$

Similar to Acemoglu and Restrepo (2022a), we assume each product τ can be produced either by workers or by technology:

$$y(\tau, t) = \max\left\{\int x dN_e(x, t; \tau), z(\tau)a(\tau, t)\right\} \tag{3}$$

Here, $N_e(x, t; \tau)$ is the cumulative measure function of the firm's demand for human capital at time t for production of product τ, $a(\tau, t)$ is its demand for technology for the same product at price q, and $z(\tau)$ is the productivity of technology in producing product τ. We order products so that $z(\tau)$ is decreasing—technology has a comparative advantage in lower-indexed products—and assume the functional form $z(\tau) = e^{-\varepsilon\tau}$. In other words, tasks are ordered so that lower-index tasks are more amenable to automation.

Two things are worth noting about our specification of the production technology. First, $a(\tau, t)$ is distinct from $a_m(t)$ described above; $a(\tau, t)$ is the technology used to produce existing tasks (as opposed to generating new tasks) and is a substitute for worker human capital. Second, ε is another key parameter in our analysis. It governs how many existing tasks will be automated as q falls; a lower value of ε implies that task automation is more elastic when the price of technology falls.

To summarize, a technology in our framework can be described by a combination of ω, ψ, and ε where, in response to a decline in the technology's price, ω describes the substitutability between technology and managerial talent in generating new ideas (tasks), ψ describes how important the technology is in maintaining the task frontier, and ε describes how that same technology will automate existing tasks.

Labor markets and wages. Labor markets are competitive and indexed by position (worker or manager). We let $w_e(t)$ denote the equilibrium wage per unit of human capital for a worker and $w_m(t)$ the wage per unit of $h(x) = x^{\xi}$ for a manager.

Firm's optimization. Therefore, the firm chooses an allocation:

$$\mathscr{A}_f := \{N_e(x, t; \tau), a(\tau, t), N_m(x, t), s(t), a_m(t)\}_{\tau\in[0,f(t)], x\in\mathbb{R}_+, t\geq 0}$$

and its problem is

$$\Pi(t) \equiv \max_{\mathscr{A}_f} \Big\{ \int_0^{f(t)} \max\Big\{ x(\tau,t)^{1-\eta^{-1}} - w_e(t)x(\tau,t), (z(\tau)a(\tau,t))^{1-\eta^{-1}} - qa(\tau,t) \Big\} d\tau \tag{4}$$

$$- w_m(t) \int h(x) dN_m(x,t) - qa_m(t) \Big\}$$

$$s.t. \quad x(\tau,t) \equiv \int x dN_e(x,t;\tau)$$

$$s(t) = \int h(x) dN_m(x,t), \quad f(t) = \bar{\tau} \left[\psi^{\omega^{-1}} a_m(t)^{1-\omega^{-1}} + (1-\psi)^{\omega^{-1}} s(t)^{1-\omega^{-1}} \right]^{\frac{1}{1-\omega^{-1}}}$$

where the firm takes wages $w_e(t)$ and $w_m(t)$ along with the price of technology, q, as given.

A natural implication of the production technology is that each product is produced entirely by workers or entirely by technology. We let $\mathscr{T}(t) \in \mathscr{B}\left([0, f(t)]\right)$ denote the Borel set of products at time t that are produced by workers and write $|\mathscr{T}(t)|$ for its measure. Formally, given the firm's demand $N_e(x,t;\tau)$ for workers with human capital of x or less to produce task τ, we have $\mathscr{T}(t) = \text{supp}_\tau \left(N_e(\infty, t; \cdot)\right)$ in equilibrium.

2.2. Households

Households are born with one unit of human capital ($x = 1$) and have linear utility. At each instant, a household with human capital x at time t chooses whether to supply its labor for production at wage $w_e(t)x$ or manage at wage $w_m(t)h(x)$. We denote this decision as $l(x,t) : [1,\infty) \to \{e, m\}$.

The key feature of the model is that human capital accumulates through learning by doing, where a worker's human capital grows at rate $\varphi(t)$, proportional to the measure of products produced by workers:

$$\dot{x}(t) = \varphi(t)x(t), \quad \varphi(t) \equiv \delta|\mathscr{T}(t)| \tag{5}$$

This learning-by-doing function is the heart of our model. A few comments are worth noting. First, we are assuming that learning is proportional to the range of tasks performed by workers in the firm. As workers, collectively, perform more tasks they learn faster. The rate of learning is governed by δ; δ will be another crucial parameter in our analysis because it will govern the strength of learning in the model. Second, learning is proportional to a worker's current level of human capital such that a more skilled worker learns faster for a given level of tasks. Third, we assume the worker human capital accrued through on-the-job learning is industry specific (as opposed to firm-specific). As a result, atomistic firms within the industry will not internalize how their actions affect either φ or equilibrium prices in their optimization process. This is the key externality in our model. Finally, we also assume atomistic workers. When a firm hires many workers, no individual worker is able to incentivize the firm to provide more learning-by-doing opportunities by taking a lower wage; there needs to be coordination among the many workers to alter firm decisions. The needed coordination among the atomistic workers results in a free rider problem that sustains the pecuniary externality.[9] As a result, our assumptions that (i) firms are atomistic within an industry, (ii) human capital accrued through on-the-job learning is industry specific, and (iii) firms hire many workers simultaneously generates the key externality in our model.

A manager's human capital, by contrast, remains constant:[10]

$$\dot{x}(t) = 0$$

The value of a household with human capital x at time t, who is not a shareholder of firms, satisfies

$$(\rho + \chi)V(x,t) = \max\Big\{ \underbrace{w_e(t)x + \varphi(t)x\partial_x V(x,t)}_{l(x,t)=e}, \underbrace{w_m(t)h(x)}_{l(x,t)=m} \Big\} + \partial_t V(x,t) \tag{6}$$

This value function shows that working pays the flow wage $w_e(t)x$ plus a capital gain $\varphi(t)x\partial_x V(x,t)$ from rising human capital; management pays $w_m(t)h(x)$ with no further accumulation. We assume households are atomostic relative to firms and thus cannot affect the automation decisions of their employers. Thus, neither firms nor households internalize how their decisions affect the learning rate $\varphi(t)$. This is the key externality of our framework.

Finally, the measure zero set of shareholders also receive the profits of the firm as a lump-sum transfer on the right hand side of (6). The inclusion of this term does not affect these household's decision on whether to work or manage as it only affects the level of the value function. Thus, the value of a the typical household in the model is defined in terms of their future labor income rather than firm profits.

[9] Our assumption of atomistic workers is distinct from the bilateral contracting problem of worker-firm trainings studied in papers like Garicano and Rayo (2017a), Garicano and Rayo (2025b), and Ide (2025b). In these other papers, individual workers can incentivize firms to provide them training opportunities by taking a lower wage.

[10] This is a simplifying assumption. If managers continued to accumulate human capital, the distribution at $x*$ (formalized below) would not be a mass point, but the threshold policy and the key externality would remain.

2.3. Distribution of Human Capital

Let $N(x,t)$ denote the CDF of human capital in the economy at time t. Given the occupational choice of households, let

$$N_e(x,t) = \int_1^x \mathbf{1}_{\{l(x',t)=e\}} N(dx',t) \tag{7}$$

denote the CDF of those choosing production. The distribution of human capital, $N(x,t)$, thus evolves according to

$$\partial_t N(x,t) = \chi(1 - N(x,t)) - \varphi(t) x \partial_x N_e(x,t) \tag{8}$$

with no mass below $x = 1$.

2.4. Equilibrium

A *competitive equilibrium* is a household allocation $\mathscr{A}_h \equiv \{V(x,t), l(x,t) : [1,\infty) \to \{e,m\}\}_{t\geq 0}$; a firm allocation $\mathscr{A}_f \equiv \{N_e(x,t;\tau), a(\tau,t), N_m(x,t), s(t), a_m(t)\}_{t\geq 0, \tau\in[0,f(t)], x\in\mathbb{R}_+}$; and prices with a rate of human capital growth $\mathscr{P} \equiv \{w_e(t), w_m(t), \varphi(t)\}$ such that: *(i) Optimality:* given $\mathscr{P}$, $\mathscr{A}_h$ and $\mathscr{A}_f$ solve the household and firm problems. *(ii) Market clearing:* labor markets for both positions clear and the rate of human capital growth is consistent with equilibrium automation:

$$\int_0^{f(t)} N_e(x,t;\tau) d\tau = N_e(x,t), \quad \forall x \geq 1, \forall t \geq 0$$

$$N_m(x,t) = N(x,t) - N_e(x,t) \quad \forall x \geq 1, \forall t \geq 0$$

$$\varphi(t) = \delta |\mathscr{T}(t)|, \quad \mathscr{T}(t) = \mathscr{T}(t) = \text{supp}_\tau (N_e(\infty, t; \cdot))$$

where the distribution of human capital evolves according to the equilibrium decisions of households in Equation (8).

Stationary equilibrium. An equilibrium is *stationary* if all allocations, prices, and distributions are time-invariant.

3 Characterizing Stationary Equilibria

We now characterize stationary equilibria, dropping time subscripts throughout. We first state the optimality conditions for firms and households, then collect the equilibrium conditions, and finally establish our main theoretical results under a set of simplifying assumptions.

3.1. Optimality Conditions

Automation threshold. Consider a product $\tau \in [0, f]$. If the firm produces it with technology, the implied inverse demand function is

$$(1-\eta^{-1})\,y(\tau)^{-\eta^{-1}} = \frac{q}{z(\tau)} = qe^{\varepsilon\tau} \tag{9}$$

whereas the implied inverse demand function, if the firm produces it with workers, is

$$(1-\eta^{-1})\,y(\tau)^{-\eta^{-1}} = w_e \tag{10}$$

Comparing (9) and (10), the firm automates product τ whenever $qe^{\varepsilon\tau} < w_e$. Three cases arise:

(i) If $w_e > qe^{\varepsilon f}$, the firm automates all products within the frontier.

(ii) If $q < w_e < qe^{\varepsilon f}$, there exists a threshold

$$\tau^* \equiv \frac{1}{\varepsilon}\ln\frac{w_e}{q} \tag{11}$$

below which all products are automated, with workers producing products in $(\tau^*, f]$.

(iii) If $w_e \leq qe^{\varepsilon 0} = q$, no products are automated.

Output and profits. Using the inverse demand functions above, output per product is

$$y(\tau) = \left(\frac{\min\{qe^{\varepsilon\tau},\, w_e\}}{1-\eta^{-1}}\right)^{-\eta} \tag{12}$$

and profits per product are

$$\pi(\tau) = \eta^{-1}\left(\frac{\min\{qe^{\varepsilon\tau},\, w_e\}}{1-\eta^{-1}}\right)^{1-\eta} \tag{13}$$

All worker-produced products share the same output level since they face a common marginal cost w_e, and output is continuous at τ^*.

Product frontier. Under the CES frontier technology (1), the firm minimizes the cost of sustaining the frontier by choosing a_m and s. Standard CES cost minimization yields the unit cost of the frontier p_f:

$$p_f = \frac{1}{\bar{\tau}}\left[\psi\, q^{1-\omega} + (1-\psi)\, w_m^{1-\omega}\right]^{\frac{1}{1-\omega}} \tag{14}$$

The optimal frontier equates the marginal value of an additional product to this unit cost:

$$\underbrace{\eta^{-1}\left(\frac{\min\{qe^{\varepsilon f}, w_e\}}{1-\eta^{-1}}\right)^{1-\eta}}_{\text{marginal value}} = \underbrace{\frac{1}{\bar{\tau}}\left[\psi\, q^{1-\omega} + (1-\psi)\, w_m^{1-\omega}\right]^{\frac{1}{1-\omega}}}_{\text{marginal cost}} \tag{15}$$

A lower technology price q affects the frontier through two channels. First, holding wages fixed, it reduces the unit cost of expanding the frontier through the CES cost function. Second, it raises the automation threshold τ^* in (11), which feeds back into the marginal value of frontier products through the worker wage in general equilibrium.

Household's problem. The household's problem (6) requires each household to choose whether to work or manage. Since human capital is frozen during management, the optimal policy is a stopping time: once a household begins managing, it never returns to working.

Lemma 1 ***(Threshold Policy).*** *Consider a stationary equilibrium and suppose $\xi\varphi < \rho + \chi$. Then there exists a threshold $x^* \in [1, \infty)$ such that $l(x) = e$ for $x \in [1, x^*)$ and $l(x) = m$ for $x \geq x^*$, with*

$$\frac{w_e x^*}{\rho + \chi - \xi\varphi} = \frac{w_m x^{*\xi}}{\rho + \chi} \quad \textit{or} \quad x^* = \left(\frac{w_e}{w_m}\frac{\rho + \chi}{\rho + \chi - \xi\varphi}\right)^{\frac{1}{\xi - 1}} \tag{16}$$

Proof. See Appendix A. ■

To understand the worker's optimal policy, consider first the case $\varphi = 0$. In this case, the human capital threshold to become a manager equates the present discounted value of working and managing: $w_e x^*/(\rho + \chi) = w_m x^{*\xi}/(\rho + \chi)$. When $\varphi > 0$, the value of remaining a worker is higher because of the expected capital gain from learning, so the threshold x^* exceeds its no-learning counterpart. This effect is amplified when ξ is larger, since a higher ξ raises the convexity of the managerial payoff, making the option to continue accumulating human capital more valuable.

When $\xi\varphi \geq \rho + \chi$, the option value of remaining a worker always dominates the flow payoff from management; workers never transition, and the supply of managers is zero at any finite wage. Since a positive supply of managers is necessary for sustaining the task frontier, no stationary equilibrium can exist in this case. The condition $\xi\varphi < \rho + \chi$ is therefore necessary for the worker's value function to be finite.

The stationary distribution of human capital, implied by the Kolmogorov forward equation (8) with $\partial_t N = 0$, is

$$N(x) = \begin{cases} 1 - x^{-\chi/\varphi} & x < x^* \\ x^{*-\chi/\varphi} & x = x^* \end{cases} \tag{17}$$

Factor supplies and market clearing. From the stationary distribution (17), the aggregate supply of managerial human capital is $s = x^{*\xi-\chi/\varphi}$ and the CES demand for frontier technology is $a_m = \frac{\psi}{1-\psi} s(q/w_m)^{-\omega}$ (derived in Appendix B). The product frontier is therefore

$$\varphi = \delta(f - \tau^*), \quad f = \bar{\tau} \left[\psi^{\omega^{-1}} a_m^{1-\omega^{-1}} + (1-\psi)^{\omega^{-1}} s^{1-\omega^{-1}} \right]^{\frac{1}{1-\omega^{-1}}} \tag{18}$$

Each worker-produced product demands the same human capital input (from (12)), so market clearing for worker human capital requires

$$\delta^{-1} \varphi \left(\frac{w_e}{1-\eta^{-1}} \right)^{-\eta} = \int_1^{x^*} x \, dN_e(x) = \frac{\chi}{\varphi - \chi} \left(x^{*1-\chi/\varphi} - 1 \right) \tag{19}$$

Therefore, a stationary equilibrium is characterized by the automation threshold (11), the frontier FOC (15), worker human capital market clearing (19), the human capital growth rate (18), and the stopping condition (16). The full derivation of factor supplies is in Appendix B.

3.2. Theoretical Results

We now simplify the framework to establish our main theoretical results. We set $\psi = 0$, so that the frontier depends only on managerial human capital: $f = \bar{\tau} s$. We also take the limits $\rho \downarrow 0$ and $\xi \downarrow 1$. Under these simplifications, any stationary equilibrium with both occupations active must have $\varphi < \chi$; otherwise, workers would never transition to management, causing the frontier to collapse. Moreover, the stopping condition simplifies to $w_m = w_e(1 - \varphi/\chi)^{-1}$, the aggregate supply of managerial human capital is $s = x^{*1-\chi/\varphi}$, and the product frontier is $f = \bar{\tau} x^{*1-\chi/\varphi}$. Throughout this subsection, we assume strictly positive prices ($w_e > 0$, $w_m > 0$, $q > 0$).

It is useful to define the *normalized technology price*

$$\bar{q} \equiv \left[\frac{q}{\left(\bar{\tau}\eta^{-1}\right)^{\eta^{-1}} \left(1-\eta^{-1}\right)^{1-\eta^{-1}}} \right]^{1/(\varepsilon\bar{\tau})} \tag{20}$$

We refer to the class of interior or full automation equilibria as "automation equilibria."

No automation. The necessary and sufficient condition for no automation is $w_e \leq q$ (equivalently, $\tau^* = 0$). In this case, the frontier FOC and market clearing conditions can be solved in closed form, yielding an implied equilibrium growth rate of:

$$\varphi = \bar{\varphi} \equiv \frac{\delta \bar{\tau}}{\eta} \tag{21}$$

With no automation, the growth rate of human capital is independent of the technology price q and increasing in the learning parameter δ and the ratio $\bar{\tau}/\eta$ of managerial productivity to demand curvature. Intuitively, a larger η lowers the value of an additional product, reducing the growth rate of human capital.

However, recall from (16) that a stationary equilibrium requires the growth rate of human capital to be smaller than the discounting adjusted by convexity of the managerial productivity, which in this case requires $\varphi < \chi$. Thus, $\varphi = \bar{\varphi}$ constitutes a stationary equilibrium as long as $\bar{\varphi} < \chi$ and $w_e \leq q$. Conditional on the former holding, the latter holds if and only if (see Appendix C),

$$\ln \bar{q} \geq \bar{\varepsilon}^{-1} \ln\left(1 - \frac{\bar{\varphi}}{\chi}\right), \qquad \bar{\varepsilon} \equiv \eta \bar{\tau} \varepsilon \tag{22}$$

Here the new parameter $\bar{\varepsilon}$ is a convenient rescaling of the elasticity parameter ε that simplifies our derivations below.

Full automation. For full automation to be a stationary equilibrium, all tasks must be automated ($\tau^* \geq f = \bar{\tau}$), so $\varphi = 0$, and all households choose management at birth. Since every agent is a manager with $h(1) = 1$, the frontier is $f = \bar{\tau}$. The frontier FOC becomes $w_m = \bar{\tau}\eta^{-1}(qe^{\varepsilon\bar{\tau}}/(1 - \eta^{-1}))^{1-\eta}$. For the firm to prefer automation of the marginal task and for workers to prefer management, we need $qe^{\varepsilon\bar{\tau}} \leq w_e \leq w_m$. Such w_e exist if and only if

$$\ln \bar{q} \leq -1 \tag{23}$$

Comparing with (22), we note that when $\bar{\varphi} < \chi$—the only regime in which a no-automation equilibrium is admissible—if $\bar{\varepsilon}^{-1} \ln\left(1 - \frac{\bar{\varphi}}{\chi}\right) \leq -1$ then no-automation and full automation are simultaneously admissible. Below, in explaining the taxonomy of equilibria, we will refer back to this condition and exclude this region of the parameter space for now. The formal assumption we will use is as follows.

Assumption 1. *In the case of $\bar{\varphi} < \chi$, assume $\bar{\varepsilon}^{-1} \ln\left(1 - \frac{\bar{\varphi}}{\chi}\right) > -1$, so that full-automation and no-automation equilibria are not simultaneously admissible.*

When $\bar{\varphi} \geq \chi$, no-automation is not admissible and the assumption is vacuous.

Interior automation. For a stationary equilibrium with $0 < \tau^* < f$ and $\varphi > 0$, combining the growth rate equation (18) with worker human capital market clearing (19) yields one expression for τ^*:

$$\frac{\tau^*}{\bar{\tau}} = 1 - \frac{\varphi}{\bar{\varphi}} \tag{24}$$

Since the learning rate is proportional to the range of worker-produced tasks ($\varphi = \delta(f - \tau^*)$), automating a larger share of the frontier leaves fewer tasks for workers and reduces human capital accumulation. Market clearing pins down the frontier, yielding a linear trade-off between the automation share and the learning rate. Furthermore, since both $\varphi > 0$ and $\tau^* > 0$ in any interior equilibrium, (24) requires

$$0 < \varphi < \bar{\varphi} \tag{25}$$

A second relationship between the automation threshold and the learning rate comes from substituting the automation threshold (11) and the stopping condition into the frontier FOC (15):

$$\frac{\tau^*}{\bar{\tau}} = \bar{\varepsilon}^{-1} \ln\left(1 - \frac{\varphi}{\chi}\right) - \ln \bar{q} \tag{26}$$

Intuitively, a higher learning rate requires a higher manager wage premium for workers to give up the human capital gains of working through the stopping condition. Since the FOC for the frontier requires manager wage to be equalized to the variety premium produced by workers on the supply side, the firm responds by reducing automation threshold, which in turn reduces worker wages (through $w_e = qe^{\varepsilon\tau^*}$) and raises the new variety premium paid to managers. We note that this argument implicitly assumes the existence of an interior τ^* so such an adjustment can be made, which requires $\ln(\bar{q})$ to be small enough. This be seen formally in (26), where if $\ln(\bar{q})$ is large enough the right-hand side is negative and a positive automation threshold does not exist. This already signals that interior equilibria are not guaranteed to exist.

Setting them equal, any interior equilibrium corresponds to a root of the function

$$g(\varphi) \equiv (1 + \ln \bar{q}) - \frac{\varphi}{\bar{\varphi}} - \bar{\varepsilon}^{-1} \ln\left(1 - \frac{\varphi}{\chi}\right) \tag{27}$$

Combining (25) with the requirement that $\varphi < \chi$ for both occupations to be active, an interior automation equilibrium requires

$$0 < \varphi < \min\{\chi, \bar{\varphi}\} \tag{28}$$

The function g is globally convex with $g''(\varphi) = \bar{\varepsilon}^{-1}(\chi - \varphi)^{-2} > 0$ and is minimized at $\varphi^* \equiv \chi - \bar{\varphi}/\bar{\varepsilon}$. Setting $g(\varphi^*) = 0$ defines a critical technology price

$$\Lambda \equiv \frac{\chi}{\bar{\varphi}} - 1 - \bar{\varepsilon}^{-1}\left[1 + \ln\frac{\chi\bar{\varepsilon}}{\bar{\varphi}}\right] \tag{29}$$

at which two interior roots of g collide. To summarize, an interior automation equilibrium exists if and only if g has a root in the interval (28). It is useful to study the cases $\bar{\varphi} > \chi$ and $\bar{\varphi} < \chi$ separately, as the binding constraint on φ and the nature of equilibrium multiplicity differ across the two regimes.

3.2.1. High human capital growth capacity ($\bar{\varphi} > \chi$). Recall that $\bar{\varphi} = \delta\bar{\tau}/\eta$, so the condition $\bar{\varphi} > \chi$ holds when the economy has a large capacity for learning by doing: a high learning intensity δ, a wide product frontier $\bar{\tau}$, or a low substitution elasticity η that increases the marginal value of additional varieties for a given set of prices.

Under $\bar{\varphi} > \chi$, the no-automation growth rate exceeds χ, so a no-automation equilibrium cannot be sustained: the implied learning rate would violate the requirement that workers eventually transition to management. Any stationary equilibrium must therefore feature strictly positive automation. The binding constraint on the learning rate is $\varphi < \chi$, and interior automation equilibria correspond to roots of g in $(0, \chi)$.

Proposition 1. *Suppose a stationary equilibrium with interior automation exists. Then there is also another stationary equilibrium, either with full automation and zero human capital growth, or with a strictly positive learning rate. In the latter case, the two equilibria have different growth rates except at a knife-edge value of q where they collapse to a single equilibrium.*

Proof. See Appendix D. ■

The source of multiplicity is a complementarity between the labor market and the automation decision that operates through the stopping condition. When workers learn quickly, they command a high option value from continued on-the-job training, which raises the manager wage premium needed to induce the transition to management. This premium increases the cost of managerial input relative to the value of worker-produced tasks, depressing the worker wage and reducing automation. With less automation, a wider range of tasks remains for workers, and it is precisely this wider range that sustains the high learning rate through the resource constraint (24). The reverse configuration is equally self-consistent: when workers learn slowly, the manager premium is small, worker wages are high, automation is extensive, and the narrow range of remaining worker tasks confirms the low learning rate. Both configurations satisfy the firm's pricing condition (26) and the resource constraint simultaneously, but at different learning rates and wage structures. If neither interior equilibrium features positive learning, the stagnant outcome takes the form of full automation with zero human capital growth.

Proposition 2. *Suppose* $\ln \bar{q} < -1$*, so that a stationary equilibrium with full automation and zero learning exists strictly. Then there is exactly one other stationary equilibrium featuring both strictly positive learning and interior automation.*

Proof. See Appendix D. ∎

The economic content is that an economy trapped in full automation is not trapped by technological limitations—the technology is cheap enough to be transformative—but by the absence of learning opportunities. When all tasks are automated, workers have nothing to learn from, and the economy settles into a stationary state with zero human capital growth. Yet an alternative is always available: if workers were to begin accumulating human capital, the rising option value of on-the-job training would push up the manager wage premium, depressing worker wages and making automation less attractive at the margin. The resulting reduction in automation would open up worker-produced tasks, sustaining exactly the learning rate that initiated the process. This alternative equilibrium is unique because there is exactly one learning rate at which the wage structure implied by the stopping condition is consistent with the resource constraint on worker-produced tasks.

Corollary 1. *Automation equilibria always exist in pairs: either both with strictly positive learning, or one with strictly positive learning and one with zero learning.*

The pairing result follows directly from Propositions 1 and 2. Economically, whenever automation plays a role, the complementarity between learning and the wage structure always generates two possible outcomes: one in which the high option value of on-the-job training sustains a large manager premium, low worker wages, and limited automation, and another in which low learning keeps the manager premium small, worker wages high, and automation extensive. No parameter configuration supports a single automation equilibrium in isolation—the economy always faces a coordination problem between a high-learning and a low-learning configuration.

Proposition 3 ***(Comparative Statics).*** *Automation equilibria have opposing comparative statics with respect to q: the growth rate in the equilibrium with high learning rises as q falls, and the growth rate in the equilibrium with low or zero learning weakly decreases as q falls.*

Proof. See Appendix D. ■

A decline in the price of technology has opposite effects on the two equilibria. Cheaper technology directly raises the automation threshold, making it profitable to automate tasks that were previously worker-produced. In the high-learning equilibrium, this reduction in worker-produced tasks is offset by the wage channel: less demand for workers depresses w_e, which raises the manager premium through the stopping condition, making workers willing to learn longer before transitioning to management. The higher learning rate restores the range of worker-produced tasks needed for the resource constraint (24) to hold, and the net effect is a rise in φ. In the low-learning equilibrium, the wage channel is too weak to offset the direct automation effect: the manager premium is small, wages barely adjust, and the additional automation crowds out worker-produced tasks. Learning falls, confirming the low-φ configuration. The same technological change thus strengthens whichever wage and learning configuration the economy is already in.

A further implication of Corollary 1 and Proposition 3 is the possibility of a *human capital trap*. Because automation equilibria always exist in pairs, the economy cannot smoothly transition from one equilibrium to the other. For sufficiently low technology prices, the full-automation zero-learning equilibrium is always admissible and coexists with a high-learning interior equilibrium. In such a trap, the lack of on-the-job learning severely limits the productivity gains associated with cheap technology, even though the technology is potentially transformative.

3.2.2. Low human capital growth capacity ($\bar{\varphi} < \chi$). Recall that $\bar{\varphi} = \delta\bar{\tau}/\eta$, so the condition $\bar{\varphi} < \chi$ holds when the economy has limited capacity for learning by doing: a low learning intensity δ, a narrow product frontier $\bar{\tau}$, or a high demand elasticity η that reduces the marginal value of additional tasks fixing prices. Under $\bar{\varphi} < \chi$, the no-automation growth rate is compatible with stationarity, so a no-automation equilibrium can be sustained. Under Assumption 1, the full automation cutoff (23) and the no-automation cutoff (22) satisfy $-1 < \bar{\varepsilon}^{-1}\ln(1-\bar{\varphi}/\chi)$, so as $\ln\bar{q}$ rises the economy transitions from full automation through interior automation to no automation. Interior automation equilibria correspond to roots of g in $(0,\bar{\varphi})$. Whether such roots are unique depends on the shape of g, which is governed by $\bar{\varepsilon}$ relative to $\bar{\varphi}/(\chi-\bar{\varphi})$.

Proposition 4 ***(Equilibria when*** $\bar{\varphi} < \chi$***).*** *Suppose* $\bar{\varphi} < \chi$ *and Assumption 1 holds.*

(i) If $\bar{\varepsilon} \geq \bar{\varphi}/(\chi-\bar{\varphi})$*, the stationary equilibrium is unique for each* $\bar{q}$*: full automation when* $\ln\bar{q} \leq -1$*, a unique interior automation equilibrium when* $\ln\bar{q} \in (-1, \bar{\varepsilon}^{-1}\ln(1-\bar{\varphi}/\chi))$*, and no automation when* $\ln\bar{q} \geq \bar{\varepsilon}^{-1}\ln(1-\bar{\varphi}/\chi)$*.*

(ii) If $\bar{\varepsilon} < \bar{\varphi}/(\chi-\bar{\varphi})$*, the cutoff* Λ *in (29) satisfies* $\Lambda > \bar{\varepsilon}^{-1}\ln(1-\bar{\varphi}/\chi)$*. For* $\ln\bar{q} \in (\bar{\varepsilon}^{-1}\ln(1-\bar{\varphi}/\chi), \Lambda)$*, the no-automation equilibrium coexist with two interior automation equilibria that collpase to a single equilibrium in the knife-edge case of* $\ln\bar{q} = \Lambda$*. For* $\ln\bar{q} < \bar{\varepsilon}^{-1}\ln(1-\bar{\varphi}/\chi)$ *and for* $\ln\bar{q} > \Lambda$*, the stationary equilibrium is unique, following the same pattern as case (i).*

Proof. See Appendix D. ■

The distinction between the two cases turns on whether the wage channel operating through the stopping condition is strong enough to sustain multiple equilibria. In case (i), $\bar{\varepsilon}$ is large relative to $\bar{\varphi}/(\chi-\bar{\varphi})$: automation costs rise steeply with the threshold, so the firm's automation decision is insensitive to the wage adjustments induced by changes in the learning rate. The manager premium moves too little to generate a second self-consistent configuration. As q falls, the economy transitions smoothly from no automation through a unique interior equilibrium to full automation, with no coordination failure along the way.

In case (ii), $\bar{\varepsilon}$ is small: automation costs rise slowly, making the firm's automation decision highly sensitive to wages. When the learning rate changes, the stopping condition adjusts the manager premium, which feeds back into worker wages and the automation threshold. This wage channel is now strong enough to sustain two interior configurations for the same $\bar{q}$: one with high learning, a large manager premium, depressed worker wages, and moderate automation, and another with low learning, a small manager premium, high worker wages, and extensive automation. These two interior equilibria coexist with the no-automation equilibrium for intermediate $\bar{q}$. At $\ln \bar{q} = \Lambda$, the two interior equilibria collapse to a single knife-edge equilibrium at $\varphi = \varphi^*$, and for $\ln \bar{q} > \Lambda$ only the no-automation equilibrium remains.

3.2.3. Taxonomy of Stationary Equilibria. Figure 1 summarizes the equilibrium structure in $(\bar{\varphi}, \bar{\varepsilon})$ space. Five regions emerge, separated by three boundary curves: Assumption 1 ($\bar{\varepsilon} = -\ln(1 - \bar{\varphi}/\chi)$, dashed), the vertical line $\bar{\varphi} = \chi$ (dash-dot), and the uniqueness boundary (solid), which traces the locus at which the minimizer $\varphi^* = \chi - \bar{\varphi}/\bar{\varepsilon}$ enters the relevant domain and takes the form $\bar{\varepsilon} = \bar{\varphi}/(\chi - \bar{\varphi})$ on the low-capacity side $\bar{\varphi} < \chi$ and $\bar{\varepsilon} = \bar{\varphi}/\chi$ on the high-capacity side $\bar{\varphi} \geq \chi$. Within each region, the number and type of stationary equilibria vary with the technology price $\bar{q}$ as characterized by the preceding propositions.

Figure 1: Taxonomy of stationary equilibria

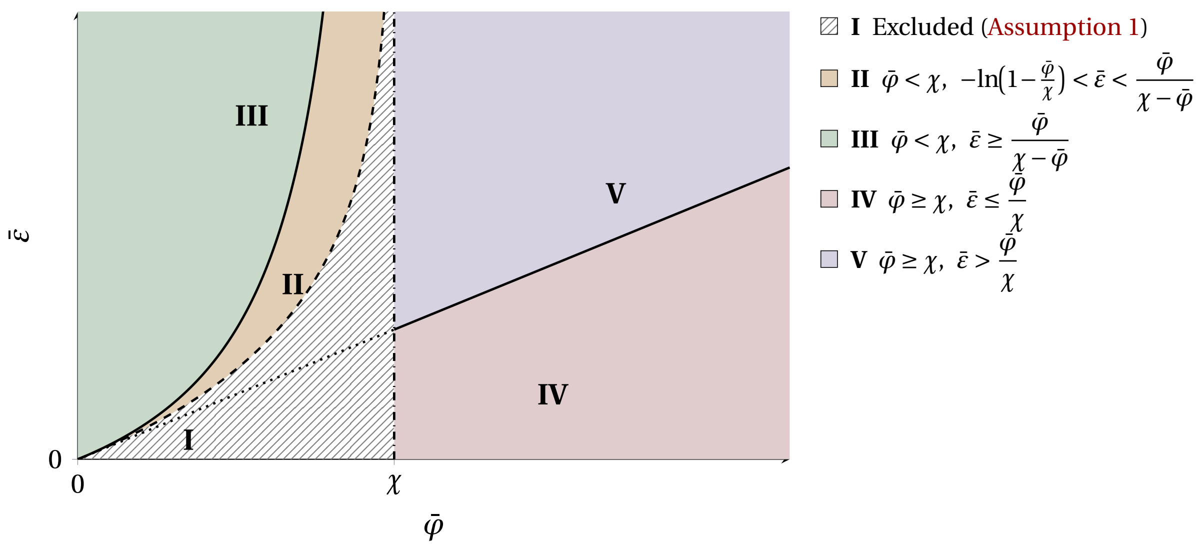

Notes: The axes are $\bar{\varphi} = \delta\bar{\tau}/\eta$ (horizontal) and $\bar{\varepsilon} = \eta\bar{\tau}\varepsilon$ (vertical). Dashed curve: $\bar{\varepsilon} = -\ln(1 - \bar{\varphi}/\chi)$ (Assumption 1). Solid curve (uniqueness boundary): $\bar{\varepsilon} = \bar{\varphi}/(\chi - \bar{\varphi})$ for $\bar{\varphi} < \chi$ and $\bar{\varepsilon} = \bar{\varphi}/\chi$ for $\bar{\varphi} \geq \chi$. Dash-dot line: $\bar{\varphi} = \chi$. Region I is excluded by Assumption 1. Regions II–V are described in the text.

Region I is excluded by Assumption 1: full automation and no automation would be simultaneously admissible, and the analysis of interior equilibria developed above does not apply.

Regions II and III correspond to the low-capacity case $\bar{\varphi} < \chi$ analyzed in Proposition 4. In Region III ($\bar{\varepsilon} \geq \bar{\varphi}/(\chi - \bar{\varphi})$), the stationary equilibrium is unique for each $\bar{q}$: the economy transitions smoothly from no automation to interior automation to full automation as the technology price falls. In Region II ($\bar{\varepsilon} < \bar{\varphi}/(\chi - \bar{\varphi})$), the wage channel operating through the stopping condition is strong enough to sustain multiplicity. For intermediate $\bar{q}$, the no-automation equilibrium coexists with two interior equilibria—one with high learning, a large manager premium, and low automation, another with low learning, a small manager premium, and extensive automation.

Regions IV and V correspond to the high-capacity case $\bar{\varphi} > \chi$ analyzed in Propositions 1 and 2 and Corollary 1. No-automation equilibria do not exist in either region. In Region IV ($\bar{\varepsilon} \le \bar{\varphi}/\chi$), the function g is monotone on $(0, \chi)$: for $\ln \bar{q} \le -1$, the full automation equilibrium coexists with a unique interior equilibrium; for $\ln \bar{q} > -1$, no stationary equilibrium exists. In Region V ($\bar{\varepsilon} > \bar{\varphi}/\chi$), the convexity of g generates pairs of automation equilibria as described by Corollary 1: for $\ln \bar{q} \le -1$, full automation coexists with one interior equilibrium; for $-1 < \ln \bar{q} < \Lambda$, two interior equilibria coexist, one with high learning and one with low learning. Beyond Λ, no stationary equilibrium exists—the economy's learning capacity exceeds what any stationary wage structure can sustain.

3.2.4. Ranking Coexisting Equilibria. Since multiple stationary equilibria can exist at a given value of q, a natural question is whether they can be ranked in terms of welfare. We rank coexisting stationary equilibria on two margins: aggregate real value added; i.e., output net of technology expenditure; and the value of an entering household. The ranking is by the learning rate on both margins.

Value added. We define aggregate value added as total real output net of what the economy spends on production technology:

$$\mathcal{V} \equiv \int_0^{\tau^*} \left(y(\tau)^{1-\eta^{-1}} - q a(\tau) \right) d\tau + \int_{\tau^*}^{f} y(\tau)^{1-\eta^{-1}} d\tau. \tag{30}$$

The first term is the contribution of automated tasks net of their technology cost; the second is the contribution of worker-produced tasks. Substituting firms' optimal input choices, the threshold condition, and the frontier FOC, value added in any automation equilibrium with learning rate $\varphi \in [0, \bar{\varphi}]$ is

$$\mathcal{V}(\varphi) = \bar{\tau}\eta^{-1} \left(\frac{q}{1-\eta^{-1}} \right)^{1-\eta} \left[\frac{1 - \exp\left(-(1-\eta^{-1})\bar{\varepsilon}\left(1 - \frac{\varphi}{\bar{\varphi}}\right)\right)}{(1-\eta^{-1})\bar{\varepsilon}} + \frac{\varphi}{\bar{\varphi}} \exp\left(-(1-\eta^{-1})\bar{\varepsilon}\left(1 - \frac{\varphi}{\bar{\varphi}}\right)\right) \right], \tag{31}$$

and in the no-automation equilibrium it is

$$\mathcal{V}^{NA} = (\bar{\tau}\eta^{-1})^{\eta^{-1}} (1-\eta^{-1})^{1-\eta^{-1}} \left(1 - \frac{\bar{\varphi}}{\chi} \right)^{\eta^{-1}-1}. \tag{32}$$

Full automation is the $\varphi = 0$ case of (31): direct substitution reproduces the value added one would derive from the FA allocation. The no-automation expression does not nest analogously. In any interior automation equilibrium the threshold τ^* is an interior optimum that moves with q, so $\mathcal{V}(\varphi)$ inherits this dependence; in no automation $\tau^* = 0$ is at a corner and does not move with q, so $\mathcal{V}^{NA}$ is independent of q. The two expressions therefore coincide only at the knife-edge price $\ln \bar{q} = \bar{\varepsilon}^{-1} \ln(1 - \bar{\varphi}/\chi)$ at which the interior regime meets no automation.

Proposition 5 ***(Ranking by value added).*** *At any q where multiple stationary equilibria coexist,*

(i) $\mathcal{V}(\varphi)$ is strictly increasing in φ on $[0, \bar{\varphi}]$: Coexisting automation equilibria are strictly ranked by their learning rate, with full automation strictly dominated by any coexisting equilibrium.

(ii) If the no-automation equilibrium coexists with an automation equilibrium, $\mathcal{V}^{NA} > \mathcal{V}(\varphi)$.

Proof. See Appendix D. ■

Economies that sustain broader on-the-job learning produce more real net output at the same technology price. The low-learning equilibrium is not a distributional alternative to the high-learning one — it is a strict aggregate loss.

Value of households. Let $V(1)$ denote the expected discounted lifetime utility of an entering household, obtained as the value function at $x = 1$ in Lemma 1. Using the stopping condition and the frontier FOC,

$$V(1;\varphi) = \frac{\bar{\tau}\eta^{-1}}{\chi} \left(\frac{q}{1-\eta^{-1}} \right)^{1-\eta} \exp\left(-(1-\eta^{-1})\bar{\varepsilon} \left(1 - \frac{\varphi}{\bar{\varphi}} \right) \right) \tag{33}$$

in any automation equilibrium with $\varphi \in [0, \bar{\varphi}]$, and

$$V^{NA}(1) = \frac{(\bar{\tau}\eta^{-1})^{\eta^{-1}} (1-\eta^{-1})^{1-\eta^{-1}}}{\chi} \left(1 - \frac{\bar{\varphi}}{\chi} \right)^{\eta^{-1}-1} \tag{34}$$

in the no-automation equilibrium. As with value added, full automation is recovered by setting $\varphi = 0$ in (33).

Proposition 6 ***(Ranking by the value of households).*** *At any q where multiple stationary equilibria coexist,*

(i) $V(1;\varphi)$ is strictly increasing on $[0, \bar{\varphi}]$: entering households have higher value in higher learning rate equilibria, with full automation strictly dominated by any coexisting equilibrium.

(ii) Whenever the no-automation equilibrium coexists with an automation equilibrium at some

φ, then $V^{NA}(1) > V(1;\varphi)$.

Proof. See Appendix D. ■

Welfare and value added move together across coexisting equilibria. The low-learning configuration is therefore not a redistribution between households and firms: it strictly reduces the value of entering households and aggregate real net output at the same time. The multiplicity of Proposition 1 is a coordination failure with first-order efficiency consequences.

4 A Decline in the Price of the Technology: A Numerical Example

We now solve the full model numerically, relaxing the simplifying assumptions of Section 3. Whereas the theoretical analysis set $\psi = 0$, $\rho \downarrow 0$, and $\xi \downarrow 1$, we evaluate the equilibrium conditions (15), (19)–(18), and (16) with the CES product frontier (1) and positive values of all parameters. The parameter values are reported in Table 1. Figure 2 plots the key equilibrium objects as functions of q. For each value of q in an intermediate range, the system admits two interior solutions, which we call the *high human capital* equilibrium (solid blue) and the *low human capital* equilibrium (solid red). The two branches coexist over a wide interval of q: at high q only the high human capital equilibrium exists; as q falls a second branch appears and the gap between the two widens. This multiplicity is the numerical counterpart of Corollary 1 in the general CES case.

Parameter	Value	Description
ρ	0.05	Discount rate
χ	0.01	Death rate
$1-1/\eta$	0.75	Returns to scale
ξ	1.02	Managerial human capital elasticity
ε	1.70	Automation cost elasticity
ψ	0.05	Technology weight in frontier CES
ω	0.20	Frontier CES elasticity of substitution
$\bar{\tau}$	1.20	Frontier technology parameter
δ	0.02	Learning rate

Table 1: Assigned parameter values.

Technology and learning across equilibria. Figure 2 displays the stationary equilibrium objects that characterize the career ladder. Figure 2c shows the human capital growth rate φ, the variable governing the model's feedback mechanism. In the high human capital equilibrium, φ *rises* as q falls. In the low human capital equilibrium, φ *falls* as q falls. These opposing responses are the numerical counterpart of Proposition 3: a decline in the technology price strengthens whichever feedback loop the economy is already in. In the high human capital equilibrium, cheaper technology reduces the cost of expanding the frontier, which creates more tasks for workers, accelerates learning, and produces more skilled managers who further expand the frontier. In the low human capital equilibrium, cheaper technology instead makes it profitable to automate more existing tasks, shrinking the range of work available for on-the-job learning and pushing φ closer to zero.

Figure 2a displays the automation cutoff τ^*, the threshold below which tasks are produced by technology. In both equilibria τ^* rises as q falls—cheaper technology makes it profitable to automate a wider range of products—but the level of τ^* is higher in the low human capital equilibrium. The reason is a general-equilibrium scarcity effect. In the low human capital equilibrium, workers transition to management quickly and few remain in the labor force; although the smaller task frontier simultaneously reduces labor demand, this supply-side collapse dominates, driving up the worker wage w_e (documented in Figure 3 below). Since $\tau^* = \varepsilon^{-1}\ln(w_e/q)$, a higher w_e raises the automation threshold. As a result, a larger share of the product frontier is automated in the low human capital equilibrium despite the frontier itself being smaller.

Figure 2b shows the worker-manager cutoff x^*, the level of human capital at which a household transitions from production to management. The two branches also diverge. In the high human capital equilibrium, x^* rises steeply as q falls, meaning workers accumulate several times their initial human capital before becoming managers. In the low human capital equilibrium, x^* remains close to 1 throughout the range: workers transition to management almost immediately, arriving with minimal skills. This divergence is tightly linked to the learning rate. The stopping condition in Lemma 1 implies that a higher φ raises the option value of remaining a worker—the capital gain from continued learning—which delays the transition to management and pushes x^* further from unity.

Figure 2d displays the product frontier f, the total measure of tasks in the economy. The frontier is determined jointly by managerial human capital and technology through the CES aggregator (1). In the high human capital equilibrium, highly skilled managers combined with cheap technology expand the frontier substantially as q falls. In the low human capital equilibrium, managers arrive with $x^* \approx 1$ and the frontier responds only modestly to cheaper technology, since low-quality managers cannot effectively leverage it to generate new tasks. The key observation is that even though automation is higher in the low human capital equilibrium—τ^* is larger—the net range of worker-produced tasks $f - \tau^*$ is larger in the high human capital equilibrium because the frontier is sufficiently expanded. Since $\varphi = \delta(f - \tau^*)$, this is the mechanism through which the two equilibria sustain different learning rates for the same technology price.

Figure 2: Technology and learning across equilibria

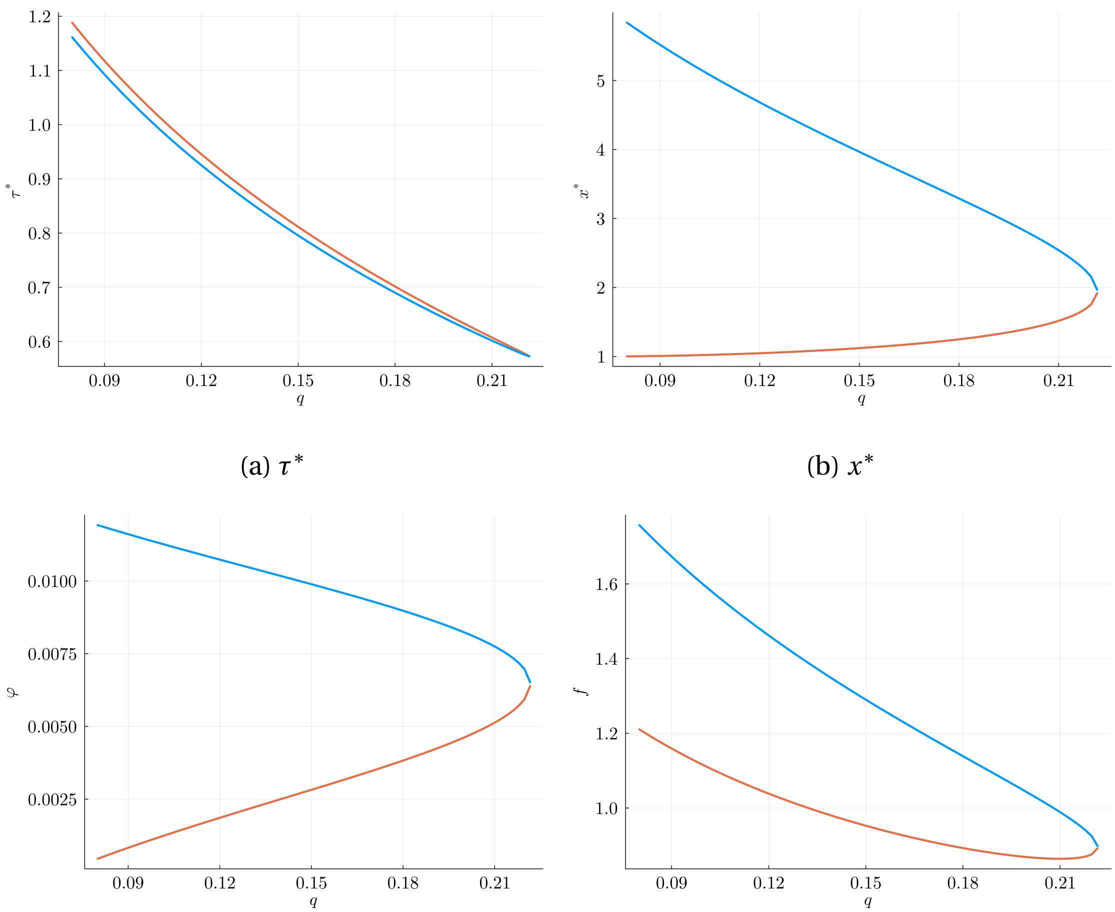

(a) τ^* (b) x^*

(c) φ (d) f

Notes: Each panel plots a stationary equilibrium object against the technology price q. The solid blue line depicts the high human capital equilibrium; the solid red line depicts the low human capital equilibrium. Figure 2a: automation cutoff τ^*, the threshold below which tasks are performed by technology. Figure 2b: worker-manager cutoff x^*, the level of human capital at which a worker transitions to management. Figure 2c: human capital growth rate φ. Figure 2d: product frontier f, the total measure of tasks in the economy.

Wages and manager income across equilibria. Figure 3 displays wages and managerial income. In interior equilibria, the marginal task at the frontier is produced by a worker, so the marginal value of expanding the frontier—and hence the return to managerial input—depends on the worker wage w_e through (15). This link between the two labor markets generates opposing wage patterns across equilibria.

In the low human capital equilibrium, workers transition quickly to management because the learning gains from remaining a worker are small. The rapid outflow leaves few workers in the labor force, driving up w_e (Figure 3a). A higher w_e raises the cost of producing each worker task, which reduces the marginal value of an additional task at the frontier and therefore the marginal product of managerial labor. As a result, w_m is lower (Figure 3b). In the high human capital equilibrium, workers remain in the labor force longer, accumulating human capital and generating a larger effective supply. This keeps w_e lower, which raises the marginal value of frontier tasks and therefore the return to managerial input, pushing w_m higher. As technology becomes cheaper (q falls), these forces amplify: in the low human capital equilibrium, increased automation further reduces the worker supply and narrows the w_e–w_m gap, while in the high human capital equilibrium, frontier expansion raises the demand for managerial input and pushes w_m higher.

Figure 3: Wages and manager income across equilibria

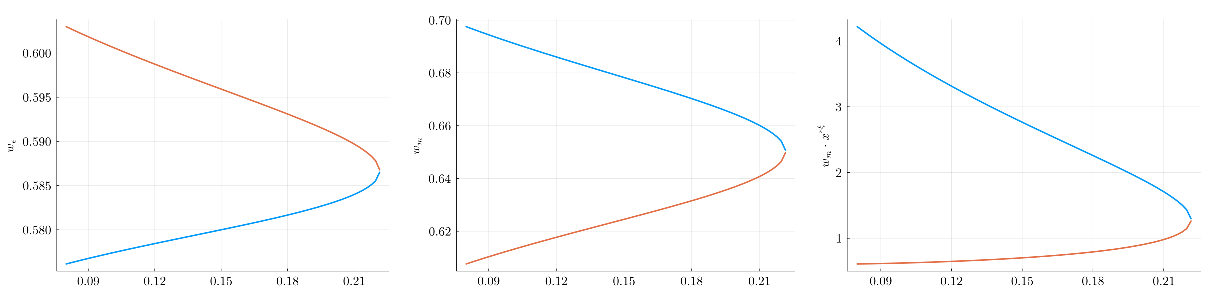

(a) w_e (b) w_m (c) Manager income

Notes: Each panel plots a price or income measure against the technology price q. The solid blue line depicts the high human capital equilibrium; the solid red line depicts the low human capital equilibrium. Figure 3a: worker wage per unit of human capital w_e. Figure 3b: manager wage per unit of managerial human capital w_m. Figure 3c: total manager income $w_m \cdot h(x^*)$, combining the manager wage rate with the human capital level at which workers transition to management.

Welfare, output, and profits across equilibria. Figure 4 displays aggregate outcomes on a logarithmic scale. Figure 4a shows the expected discounted lifetime utility of an entering worker, $V(1)$. The high human capital equilibrium delivers substantially higher welfare, reflecting the present value of higher managerial income earned over a longer career. The temporarily lower worker wage w_e in the high human capital equilibrium is more than offset by the option value of accumulating human capital and becoming a high-quality manager.

Figures 4b and 4c show value added and total firm profits. Both equilibria exhibit rising output as q falls, reflecting direct productivity gains from cheaper technology. The two branches are close at high q and diverge modestly as q falls, with the high human capital equilibrium generating somewhat higher value added. Total firm profits, however, are nearly identical across equilibria on the logarithmic scale.

The disconnect between firm profits and worker welfare is striking. Firms earn similar profits regardless of which equilibrium prevails, because the surplus from automation offsets the smaller frontier in the low human capital equilibrium. Workers, by contrast, fare dramatically worse in the low human capital trap. This disconnect is precisely the externality that motivates the planner's problem in Section 5: atomistic firms, taking the learning rate φ as given, do not internalize how their automation decisions affect the economy's stock of human capital.

Figure 4: Welfare, output, and profits across equilibria

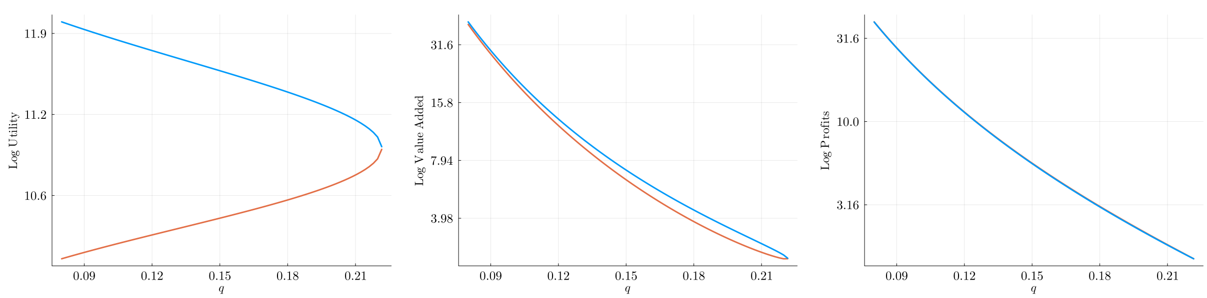

(a) Individual Utility (b) Value Added (c) Profits

Notes: Each panel plots an aggregate outcome against the technology price q on a logarithmic scale. The solid blue line depicts the high human capital equilibrium; the solid red line depicts the low human capital equilibrium. Figure 4a: expected discounted lifetime utility of an entering worker $V(1)$. Figure 4b: value added, defined as total output net of technology costs. Figure 4c: total firm profits.

4.1. Robustness

The theoretical analysis in Section 3 identifies four structural features that determine whether the economy is vulnerable to the low human capital trap as the price of technology falls: the breadth of task automation (governed by ε), the speed of workplace learning (δ), how effectively technology complements managers in expanding the frontier (ω and ψ), and the rate of workforce turnover (χ). Figures 5–9 vary each of these parameters one at a time, holding the others at their baseline values. Each figure plots six equilibrium objects—τ^*, x^*, φ, f, $V(1)$, and the automation share τ^*/f—against the technology price q. Solid and dashed lines distinguish different parameter values; blue lines correspond to the high human capital equilibrium and red lines to the low human capital equilibrium. The qualitative structure—a region of multiplicity flanked by regions with a unique equilibrium—is preserved throughout. We focus on how each parameter shifts the boundaries of the multiplicity region and the gap between the two equilibria.

Sensitivity to ε (Figure 5). The parameter ε governs how quickly the productivity of technology in automating tasks decays across the product space: $z(\tau) = e^{-\varepsilon\tau}$. A higher ε means that technology is productive only for a narrow range of low-indexed tasks, so a decline in q automates fewer products. From the threshold $\tau^* = \varepsilon^{-1}\ln(w_e/q)$, a higher ε compresses τ^* for any given wage-price ratio. Figure 5 confirms this: raising ε shifts the automation cutoff downward in both equilibria and widens the region of q over which the high human capital equilibrium exists. The learning rate φ is lower for larger ε in the high human capital branch, even as the economy retains a larger overall share of worker-produced tasks when automation is less elastic.

Sensitivity to ω (Figure 6). The parameter ω is the elasticity of substitution between technology and managerial human capital in the CES frontier (1). A higher ω means that technology can substitute more easily for managerial quality in generating new tasks. Figure 6 shows that raising ω *narrows the gap* between the two equilibria: φ falls in the high human capital branch and rises in the low human capital branch. The intuition is that when the two inputs are more substitutable, the advantage of having highly skilled managers is diluted—technology can partially compensate for low managerial quality, lifting the low human capital equilibrium. At the same time, the high human capital equilibrium benefits less from its skilled managers because technology is already a close substitute. The region of q over which both equilibria

Figure 5: Sensitivity to ε

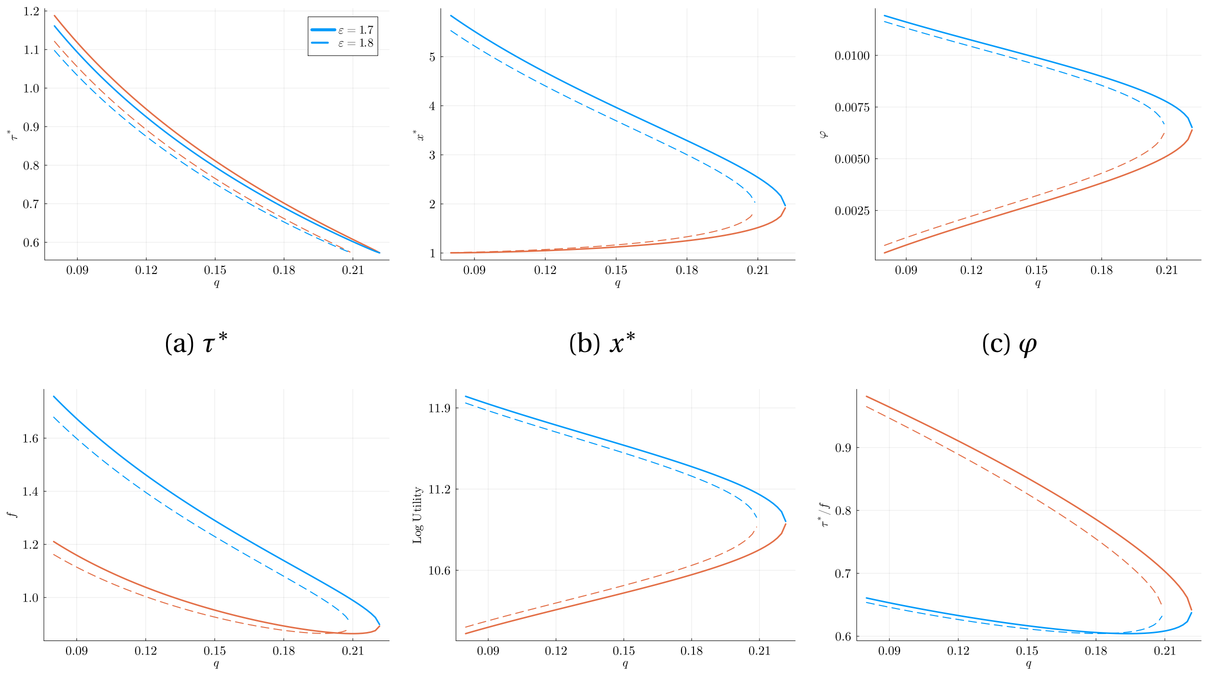

(a) τ^* (b) x^* (c) φ

(d) f (e) Individual Utility (f) τ^*/f

Notes: Each panel plots a stationary equilibrium object against the technology price q for different values of the automation cost elasticity ε. Blue lines correspond to the high human capital equilibrium; red lines correspond to the low human capital equilibrium. Solid and dashed lines distinguish different parameter values. Panels show: (a) automation cutoff τ^*, (b) worker-manager cutoff x^*, (c) human capital growth rate φ, (d) product frontier f, (e) expected lifetime utility of an entering worker $V(1)$, (f) automation share τ^*/f.

coexist widens slightly.

Sensitivity to χ (Figure 7). The parameter χ is the death rate, governing workforce turnover. A higher χ means the economy must constantly replenish its stock of experienced workers from entrants who start at $x = 1$. Figure 7 shows that raising χ substantially widens the gap between the two equilibria and shifts the region of multiplicity toward higher values of q—the low human capital equilibrium appears at technology prices where it previously did not exist. The learning rate φ falls and the worker-manager cutoff x^* declines in the high human capital branch, reflecting that faster turnover prevents workers from accumulating as much human capital before death removes them. From the stationary distribution $N(x) = 1 - x^{-\chi/\varphi}$, a higher ratio χ/φ concentrates more mass at low human capital levels, reducing the aggregate supply of managerial quality and constraining the frontier.

Figure 6: Sensitivity to ω

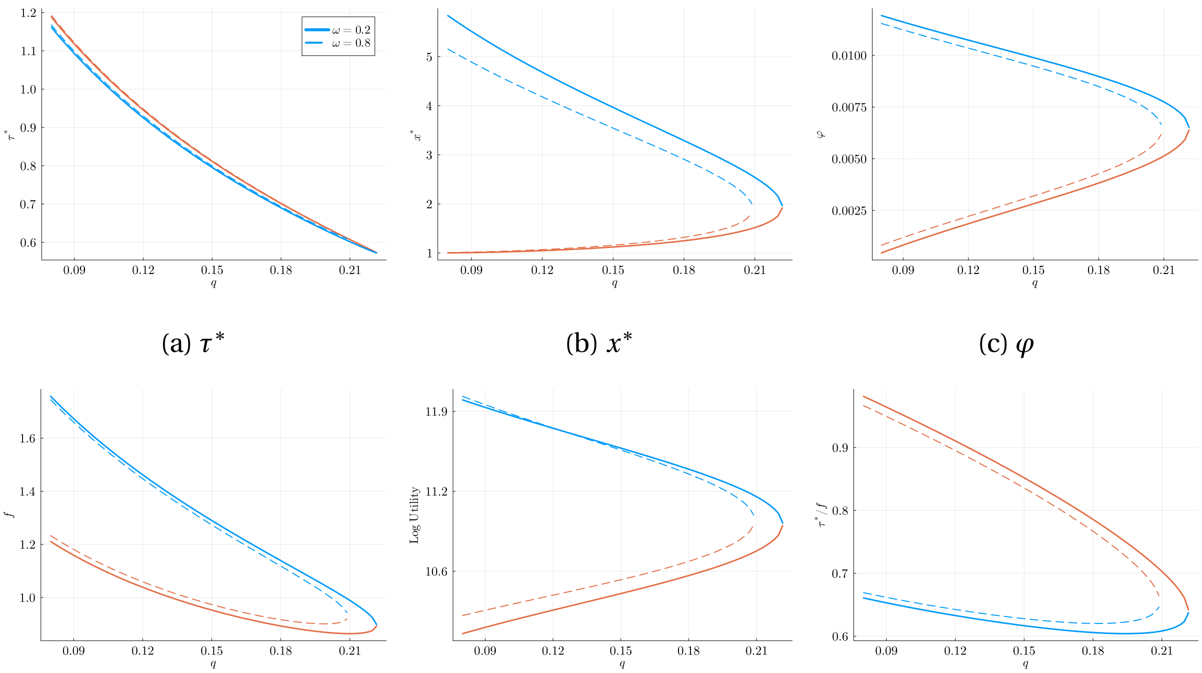

(a) τ^* (b) x^* (c) φ

(d) f (e) Individual Utility (f) τ^*/f

Notes: Each panel plots a stationary equilibrium object against the technology price q for different values of the frontier CES elasticity ω. Blue lines correspond to the high human capital equilibrium; red lines correspond to the low human capital equilibrium. Solid and dashed lines distinguish different parameter values. Panels show: (a) automation cutoff τ^*, (b) worker-manager cutoff x^*, (c) human capital growth rate φ, (d) product frontier f, (e) expected lifetime utility of an entering worker $V(1)$, (f) automation share τ^*/f.

Sensitivity to δ (Figure 8). The parameter δ scales the learning-by-doing function $\varphi = \delta(f - \tau^*)$, governing how much human capital workers accumulate for a given range of tasks. Figure 8 shows that raising δ narrows the gap between the two equilibria: φ falls slightly in the high human capital branch and rises in the low human capital branch. In the low human capital equilibrium, the constraint on learning is the narrow range of worker-produced tasks $f - \tau^*$; a higher δ directly raises φ for a given task range, generating more human capital from the same set of entry-level positions. In the high human capital equilibrium, a higher δ raises the option value of remaining a worker, which delays the transition to management, reduces the supply of managers, and constrains the frontier—which overwhelms the direct effect of higher δ in increasing on the learning rate and leads to a reduction in φ.

Figure 7: Sensitivity to χ

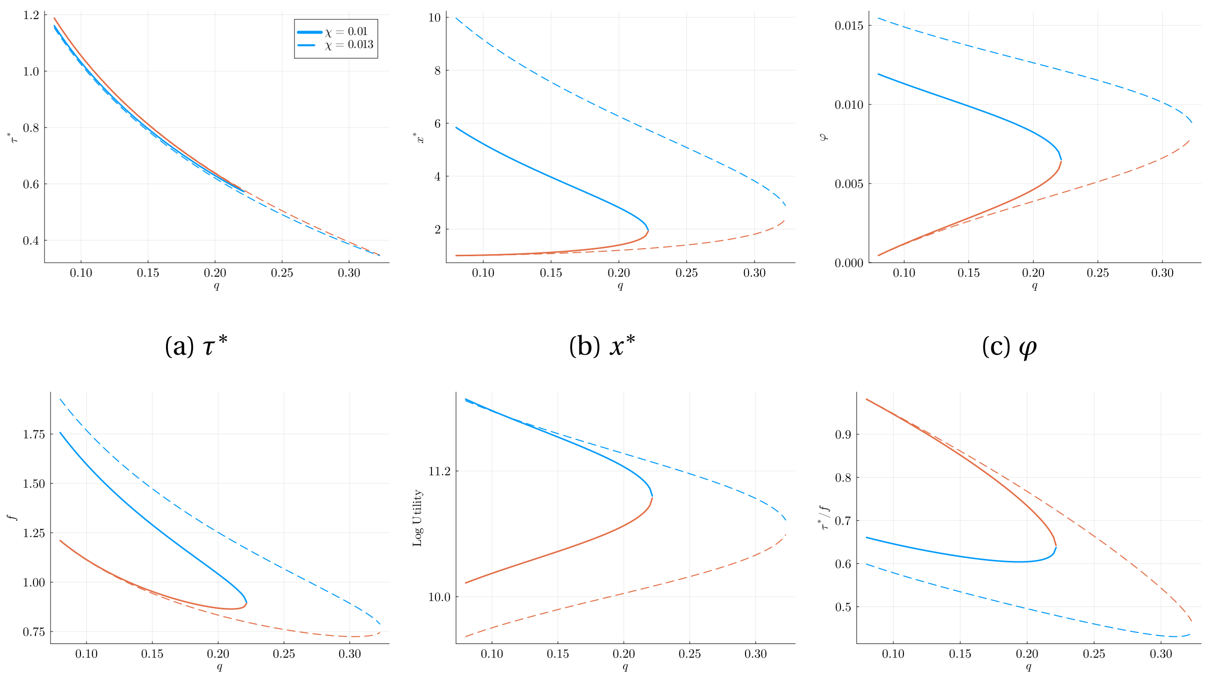

(a) τ^* (b) x^* (c) φ

(d) f (e) Individual Utility (f) τ^*/f

Notes: Each panel plots a stationary equilibrium object against the technology price q for different values of the death rate χ. Blue lines correspond to the high human capital equilibrium; red lines correspond to the low human capital equilibrium. Solid and dashed lines distinguish different parameter values. Panels show: (a) automation cutoff τ^*, (b) worker-manager cutoff x^*, (c) human capital growth rate φ, (d) product frontier f, (e) expected lifetime utility of an entering worker $V(1)$, (f) automation share τ^*/f.

Sensitivity to ψ (Figure 9). The parameter ψ governs the relative weight of technology in the CES frontier (1). The limiting case $\psi = 0$ recovers the model analyzed in Section 3, in which the frontier depends only on managerial human capital. A positive ψ allows cheap technology to contribute directly to frontier expansion. Figure 9 shows that raising ψ narrows the gap between equilibria: the frontier f and the learning rate φ both fall in the high human capital branch and rise in the low human capital branch. The mechanism parallels that of ω: a larger technology weight in the frontier reduces the importance of managerial quality. In the high human capital equilibrium, where managers are highly skilled, shifting weight toward technology dilutes the return to their human capital, compressing the frontier and lowering φ. In the low human capital equilibrium, where managers arrive with minimal skills, the technology channel provides a direct path to frontier expansion that does not depend on managerial quality, raising f and φ.

Figure 8: Sensitivity to δ

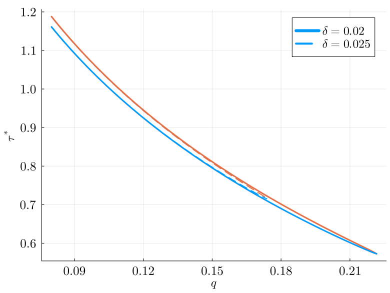

(a) τ^*

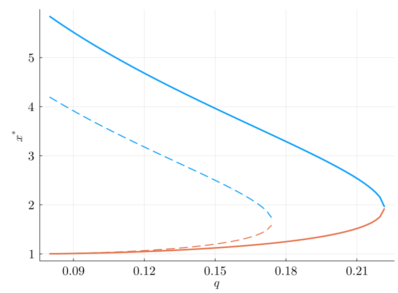

(b) x^*

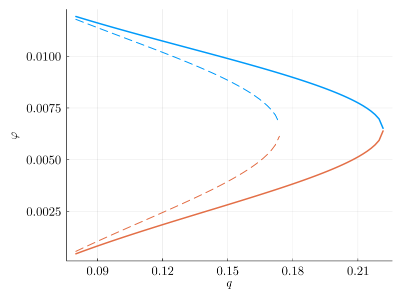

(c) φ

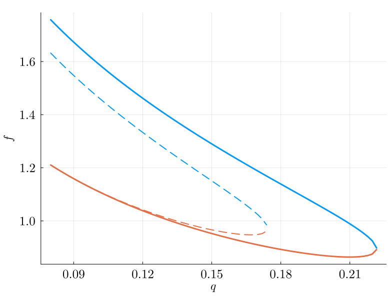

(d) f

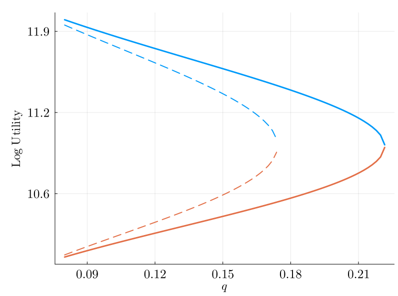

(e) Individual Utility

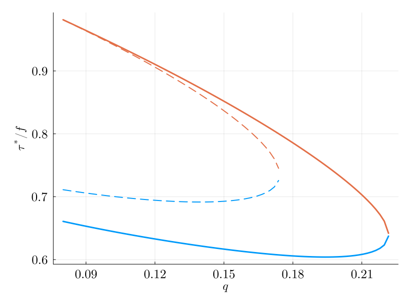

(f) τ^* / f

Notes: Each panel plots a stationary equilibrium object against the technology price q for different values of the learning rate δ. Blue lines correspond to the high human capital equilibrium; red lines correspond to the low human capital equilibrium. Solid and dashed lines distinguish different parameter values. Panels show: (a) automation cutoff τ^*, (b) worker-manager cutoff x^*, (c) human capital growth rate φ, (d) product frontier f, (e) expected lifetime utility of an entering worker $V(1)$, (f) automation share τ^* / f.

Figure 9: Sensitivity to ψ

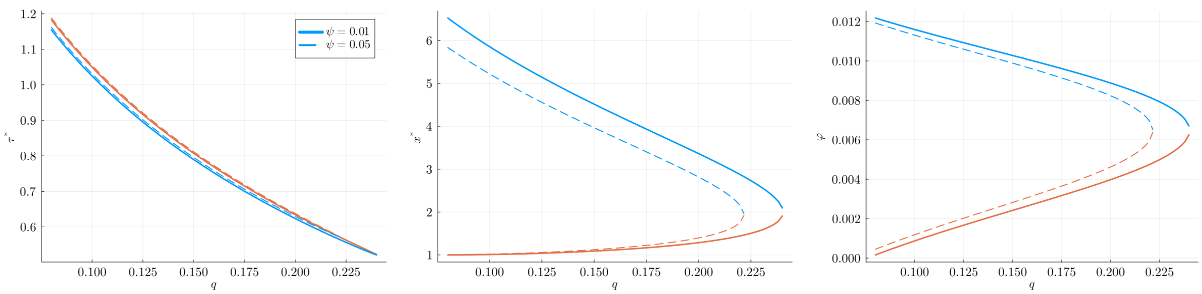

(a) τ^* (b) x^* (c) φ

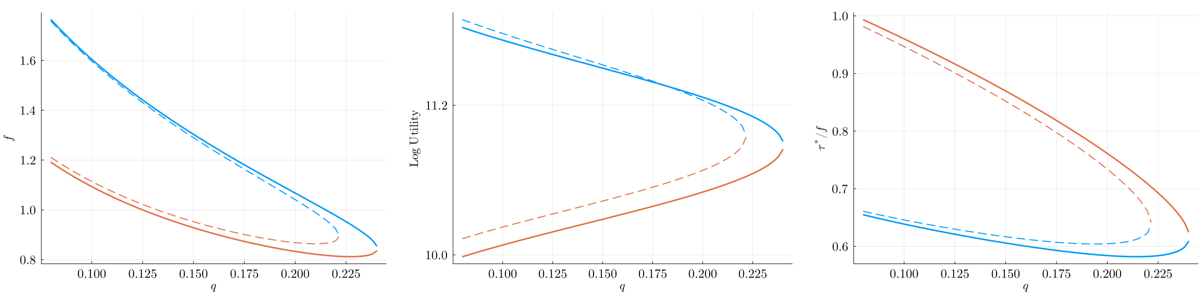

(d) f (e) Individual Utility (f) τ^*/f

Notes: Each panel plots a stationary equilibrium object against the technology price q for different values of the technology weight in frontier production ψ. Blue lines correspond to the high human capital equilibrium; red lines correspond to the low human capital equilibrium. Solid and dashed lines distinguish different parameter values. Panels show: (a) automation cutoff τ^*, (b) worker-manager cutoff x^*, (c) human capital growth rate φ, (d) product frontier f, (e) expected lifetime utility of an entering worker $V(1)$, (f) automation share τ^*/f.

5 Planner's Problem

This section characterizes the planner's problem. For any variable d, we denote with $\tilde{d}$ its value in the planner's allocation. For example, if τ^* is the threshold in the market equilibrium, $\tilde{\tau}^*$ is the threshold in the planner's allocation.

Let

$$U_h^w = \int_h^\infty e^{-(\rho+\chi)(t-h)} \tilde{c}_{t,h} dt \tag{35}$$

be the value of the cohort born in period h, with $\tilde{c}_{t,h}$ denoting the consumption of this cohort at time t. The planner's problem is to choose the allocation to maximize

$$\mathscr{W}_0 = \int_{-\infty}^\infty e^{-\rho h} U_h^w dh \tag{36}$$

where $e^{-\rho h}$ is the Pareto weight across different cohorts. The next Lemma shows that the planner's problem can be rewritten as maximizing the discounted value of aggregate consumption, which is the standard planner's problem in a representative agent model.

Lemma 2. *The planner's objective function is equivalent to maximizing the discounted value of aggregate consumption:*

$$\mathscr{W}_0 = \frac{1}{\chi} \int_0^\infty e^{-\rho t} \tilde{c}_t dt + t.i.p. \tag{37}$$

where c_t is aggregate consumption at time t, $1/\chi$ is a monotonic transformation, and $t.i.p.$ stands for terms that are independent of the allocation after time 0.

With the objective function, we can write the planner's problem as choosing the allocation:

$$\mathscr{A}_p := \{\tilde{y}(\tau,t), \tilde{s}(t), \tilde{a}_m(t), \tilde{f}(t), \tilde{\tau}^*(t), \tilde{N}(x,t), \tilde{N}_e(x,t)\}_{\tau\in[0,f(t)], x\in\mathbb{R}_+, t\geq 0}$$

to maximize

$$\max_{\mathscr{A}_p} \int_0^\infty e^{-\rho t} \left[\int_0^{\tilde{\tau}^*(t)} \left(\tilde{y}(\tau,t)^{1-\eta^{-1}} - \frac{q\tilde{y}(\tau,t)}{\tilde{z}(\tau)} \right) d\tau + \int_{\tilde{\tau}^*(t)}^{\tilde{f}(t)} \left(\tilde{y}(\tau,t)^{1-\eta^{-1}} \right) d\tau - q a_m(t) \right] dt \tag{38}$$

subject to the product frontier technology (1), feasibility in the labor markets workers and managers, which can be written as

$$\int h(x) d\left(\tilde{N}(x,t) - \tilde{N}_e(x,t)\right) = \tilde{s}(t), \tag{39}$$

$$\int x d\tilde{N}_e(x,t) = \int_{\tilde{\tau}^*(t)}^{\tilde{f}(t)} \tilde{y}(\tau,t), \tag{40}$$

$$\tag{41}$$

the learning-by-doing process (5) given by

$$\tilde{\varphi}(t) = \delta(\tilde{f}(t) - \tilde{\tau}^*(t)), \tag{42}$$

the evolution of the distribution of human capital (8), $\tilde{\tau}^*(t) \leq f(t)$, and the non-negativity constraint $d\tilde{N}(x,t) \geq 0$, $d\tilde{N}_e(x,t) \geq 0$, and $d\tilde{N}(x,t) - d\tilde{N}_e(x,t) \geq 0$.

We now characterize the optimal policy in three parts. First, we characterize the optimal automation policy, which is the planner's choice of $\tilde{\tau}^*(t)$ and $\tilde{f}(t)$, which imposes cross-equation restrictions to the Lagrange multipliers in the planner problem. Second, we characterize the law of motion of workers, $\tilde{N}_e(x,t)$ and $\tilde{N}(x,t) - \tilde{N}_e(x,t)$, which solve an optimal stopping problem. Then, we evaluate the equilibrium condition in the steady state and characterize the planner solution as a system of 4 equations and 4 unknowns.

The Lagrange multipliers of the planner's problem have natural interpretations as shadow prices and can be directly compared to their market-equilibrium counterparts. The multiplier on the worker labor-market constraint (40), denoted $\tilde{w}_e(t)$, is the planner's shadow wage per unit of worker human capital: it measures the marginal value of relaxing the constraint that worker human capital supply equals task demand. Analogously, the multiplier on the manager labor-market constraint (39), denoted $\tilde{w}_m(t)$, is the shadow wage per unit of managerial human capital. Finally, there is a multiplier on the learning-by-doing constraint (42) captures the social value of the learning rate $\tilde{\varphi}(t)$, the object that is absent from individual firm and worker optimization and is therefore the source of wedges between the planner's allocation and the competitive equilibrium. We represent a normalized version of this multiplies with $\tilde{\kappa}(t) \in (0,1)$. For simplicity, for the rest of this section, we assume that $\tilde{\tau}^* \in (0, \tilde{f}(t))$ for all t. In the Online Appendix, we characterize also the corner solutions.

Proposition 7 ***(Wedges).*** *In the planner's problem, the optimal automation policy is to choose $\tilde{\tau}^*(t)$ and $\tilde{f}(t)$ to satisfy the following conditions:*

$$q e^{\varepsilon \tilde{\tau}^*(t)} = \tilde{w}_e(t)(1 - \tilde{\kappa}(t)) \tag{43}$$

$$\frac{1}{\bar{\tau}}\left[\psi\, q^{1-\omega} + (1-\psi)\,\tilde{w}_m(t)^{1-\omega}\right]^{\frac{1}{1-\omega}} = \eta^{-1}\left(\frac{(1-\tilde{\kappa}(t))\,\tilde{w}_e(t)}{1-\eta^{-1}}\right)^{1-\eta} \tag{44}$$

output in each task given by

$$y(\tau, t) = \begin{cases} \left(\frac{q e^{\varepsilon\tau}}{1-\eta^{-1}}\right)^{-\eta}, & \tau < \tilde{\tau}^*(t) \\ \left(\frac{\tilde{w}_e(t)}{1-\eta^{-1}}\right)^{-\eta}, & f(t) \geq \tau > \tilde{\tau}^*(t) \end{cases} \tag{45}$$

Proposition 7 reveals two externalities that the firm in the market equilibrium fails to internalize, both captured by the single shadow value $\tilde{\kappa}(t) > 0$ on the learning-by-doing constraint (42). The first externality operates through the automation margin $\tau^*(t)$ as we see in equation (43). When a firm automates one additional task, it reduces the measure of worker-produced tasks $[\tau^*, f]$ and therefore lowers the aggregate learning rate $\varphi = \delta(f - \tau^*)$. The firm captures the full cost saving from automation but ignores this negative spillover on human capital accumulation. The planner corrects this by effectively raising the shadow cost of automation: the relevant comparison price for technology is $\tilde{w}_e(1-\tilde{\kappa})$ rather than $\tilde{w}_e$, so the planner automates fewer tasks whenever $\tilde{\kappa} > 0$. The second externality operates through the frontier expansion margin $f(t)$ as we see in equation (44). A larger frontier raises the measure of worker-produced tasks and therefore the learning rate, but firms expanding the frontier capture only the direct profit from new tasks and ignore this positive spillover. The planner internalizes this benefit, and the same wedge $\tilde{\kappa}$ appears in (44): the planner is willing to incur a higher unit cost of frontier expansion than the market because it also values the associated increase in φ.

A wage subsidy to workers at rate $\tilde{\kappa}(t)$—making the effective wage $(1-\tilde{\kappa})\,w_e$—would correct both externalities simultaneously. By lowering the effective cost of worker human capital, such a subsidy discourages automation (reducing τ^*) and raises the marginal value of frontier expansion (increasing f), correctly moving both margins in the planner's direction. However, the wage subsidy is not a first-best instrument. As equation (45) shows, it would also distort output within every worker-produced task: since task-level output is $(w_e/(1-\eta^{-1}))^{-\eta}$, lowering the perceived worker wage causes the firm to overproduce in each worker task relative to the planner's allocation $(\tilde{w}_e/(1-\eta^{-1}))^{-\eta}$, creating a wedge on a margin that the competitive equilibrium already prices correctly. In general equilibrium, the distortion is compounded since the subsidy raises the equilibrium wage of workers and affects the threshold to become a manager. Note also that a tax on automation at rate $\tilde{\kappa}/(1-\tilde{\kappa})$—which would equivalently raise the threshold τ^*—is not the same instrument as a labor subsidy for the frontier margin: because the unit frontier cost depends on both q and w_m through the CES aggregator, taxing automation alone shifts the cost of frontier expansion differently from a subsidy to labor. Despite these distortions, the wage subsidy generates significant welfare gains by moving both the automation threshold and the task frontier in the planner's preferred direction.

The planner's allocation of workers across occupations is an optimal stopping problem: at each instant, the planner must decide whether an agent with human capital x should continue producing tasks and accumulating human capital, or transition permanently to management. To characterize this decision, let $\tilde{v}(x,t)$ be the Lagrange multiplier on the Kolmogorov forward equation (8) at level x in period t and define

$$\partial_x \tilde{V}(x,t) \equiv \tilde{v}(x,t),$$

as the marginal social value of human capital, so that $\tilde{V}(x,t)$ is the planner's shadow value of a worker at human-capital state x—the planner analog of the household value function evaluated at shadow wages. This object is the only quantity the planner needs to determine the optimal transition rule. On the support of the worker distribution, the structure of the problem is identical to the household's stopping problem in the decentralized economy, with one key difference: the private wages w_e and w_m are replaced by the planner's shadow wages $\tilde{w}_e$ and $\tilde{w}_m$. Proposition 9 formalizes this analogy and characterizes the optimal wedge.

Proposition 8 ***(Human capital accumulation).*** *In the planner's problem, the optimal transition policy for workers solves the Hamilton-Jacobi-Bellman variational inequality*

$$(\rho+\chi)\tilde{V}(x,t) = \max\left\{\tilde{w}_m(t)h(x), \tilde{\varphi}(t)x\tilde{V}_x(x,t) + \tilde{w}_e(t)x\right\} + \partial_t \tilde{V}(x,t) \tag{46}$$

and the wedge satisfies

$$\underbrace{(\tilde{w}_e(t)(1-\tilde{\kappa}(t)))^{1-\eta} - \tilde{w}_e(t)^{1-\eta}}_{wedge} = \underbrace{\delta\eta(1-\eta^{-1})^{1-\eta}\int_1^\infty x\tilde{V}_x(x,t)\,d\tilde{N}_e(x,t)}_{aggregate\ marginal\ value\ of\ human\ capital} \tag{47}$$

With all the equilibrium conditions in hand, we now discuss the implementability of the first best. While there could be several ways to implement the first best, we look for one that is parsimonious and targets each externality with a specific instrument. We consider two instruments. First, to correct the externality arising from task composition on human capital accumulation, we introduce a tax on profits from AI tasks, denoted $T_{AI}(t)$. By raising the effective cost of AI tasks at the margin, this tax induces firms to keep the learning threshold $\tau^*(t)$ at its socially optimal level, internalizing the human capital externality. Second, to correct the externality of product development on aggregate growth, we introduce a subsidy to to frontier task expenditures, denoted $T_F(t)$. By lowering the effective cost of developing new task, this subsidy induces firms to expand the knowledge frontier at the socially optimal rate. Finally, we assume lump-sum taxes to the households to finance the next revenue-expenditures of the government. Under these two instruments, the firm's objective function becomes:

$$\int_0^{f(t)} \max\left\{x(\tau,t)^{1-\eta^{-1}} - w_e(t)x(\tau,t), (1-T_{AI}(t))\underbrace{\left((z(\tau)a(\tau,t))^{1-\eta^{-1}} - qa(\tau,t)\right)}_{\text{profits from AI}}\right\}d\tau \tag{48}$$

$$\cdots - (1-T_F(t))\underbrace{\left(w_m(t)\int h(x)dN_m(x,t) + qa_m(t)\right)}_{\text{managerial expenditure}}$$

The next proposition shows the values these instruments must take for the competitive equilibrium to implement the first best.

Proposition 9 ***(Implementability of First Best).*** *Assume that*

$$T_F(t) = T_{AI}(t) = 1 - (1-\tilde{\kappa}(t))^{\eta-1}. \tag{49}$$

Then, there exists an equilibrium that implements the first best.

Since the same wedge appears in both of the planner's optimality conditions, both instruments take the same value, entirely determined by the planner's optimal wedge $\tilde{\kappa}(t)$. Recall from (47) that $\tilde{\kappa}(t)$ is a proportional wedge on the shadow worker wage: it enters symmetrically in the automation FOC (43) and the frontier FOC (44) through the effective wage $(1-\tilde{\kappa}(t))\tilde{w}_e(t)$. It captures the extent to which the social value of human capital accumulation exceeds its private value: a larger $\tilde{\kappa}$ calls for higher taxes and subsidies. The analysis also clarifies why a natural single-instrument alternative—a uniform tax on the price of the technology—would be incorrect. Such a tax would indeed slow automation, but because the same technology also enables managers to generate new tasks, it would simultaneously discourage the frontier expansion the planner wants to encourage, leaving the under-innovation distortion unaddressed or even worsening it.

6 Conclusion

This paper develops a theory of how automation and task creation jointly shape career dynamics and the long-run accumulation of human capital. The central mechanism operates through the learning rate: workers acquire skills by performing tasks, so the equilibrium stock of human capital depends on the range of tasks that remain worker-produced. Cheaper technology affects this range through two opposing channels. On the automation margin, it displaces entry-level tasks, compressing the range from below and reducing learning opportunities. On the frontier margin, it complements managerial human capital in generating new tasks, expanding the range from above and enhancing learning. The tension between these how two forces interact with on-the-job human capital accumulation is what gives the model its richness.

We highlight how these interactions can lead to an economy with multiple equilibrium. When the frontier effect is strong relative to the automation effect, the economy sustains a high human capital equilibrium: workers remain in the labor force longer, accumulate more skills, and arrive at management with higher quality, which in turn supports further frontier expansion. When the automation effect dominates, the economy can fall into a low human capital trap: reduced learning produces lower-quality managers, which constrains frontier expansion, which further concentrates the economy's tasks at the automated end — a self-reinforcing cycle. Crucially, atomistic firms making individual automation and expansion decisions do not internalize either the negative spillover from automation or the positive spillover from frontier expansion onto the economy's aggregate learning rate. Likewise, atomistic workers at a given firm cannot individually incentivize the firm to defer automation by offering to work for a lower wage. The competitive equilibrium is therefore constrained inefficient on both margins simultaneously.

The planner's solution corrects this by raising the shadow cost of automation and reducing the shadow cost of frontier expansion relative to the market. The planner's first-best combines a tax on automation profits with a subsidy on frontier-maintenance expenditures at a common rate. A wage subsidy to workers can partially decentralize this solution by discouraging automation and encouraging learning, though it is not a first-best instrument because it also distorts task-level output decisions. A subsidy to worker wages is preferred to a tax on automation because the tax on automation will deter the frontier expansion of new tasks. The analysis suggests that evaluating automation policy in isolation — without accounting for the complementary role of task creation — will systematically understate the social value of policies that sustain worker learning.

Several natural extensions follow from the framework. Allowing for worker heterogeneity in learning rates would introduce distributional questions about which workers are most exposed to the human capital trap. Endogenizing technology prices through profit-maximizing innovation would allow the model to speak to the strategic interaction between automation investment and labor market policy. And the growing empirical literature on AI adoption and worker skill trajectories — particularly evidence on how AI affects the composition of tasks performed by early-career workers and the life cycle profile of wages — could discipline the model's key parameters and assess how close real economies are to the multiplicity region. We leave these for future work.

References

ACEMOGLU, D. (1996): "A Microfoundation for Social Increasing Returns in Human Capital Accumulation," *Quarterly Journal of Economics*, 111(3), 779–804.

ACEMOGLU, D., AND D. AUTOR (2011): "Skills, Tasks and Technologies: Implications for Employment and Earnings," in *Handbook of Labor Economics*, ed. by O. Ashenfelter, and D. Card, vol. 4B, pp. 1043–1171. Elsevier.

ACEMOGLU, D., D. AUTOR, AND S. JOHNSON (2026): "Building Pro-Worker Artificial Intelligence," Working Paper 34854, National Bureau of Economic Research.

ACEMOGLU, D., AND S. JOHNSON (2024): "Learning from Ricardo and Thompson: Technology, Wages, and Power, Then and Now," *Annual Review of Economics*, 16, Also available as NBER Working Paper No. 32416.

ACEMOGLU, D., D. KONG, AND A. OZDAGLAR (2026): "AI, Human Cognition and Knowledge Collapse," *Working Paper*.

ACEMOGLU, D., AND P. RESTREPO (2018): "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment," *American Economic Review*, 108(6), 1488–1542.

——— (2019): "Automation and New Tasks: How Technology Displaces and Reinstates Labor," *Journal of Economic Perspectives*, 33(2), 3–30.

——— (2020a): "Robots and Employment: Evidence from US Labor Markets," *Journal of Political Economy*, 128(6), 2188–2244.

——— (2020b): "Robots and Jobs: Evidence from US Labor Markets," *Journal of Political Economy*, 128(6), 2188–2244.

——— (2022a): "Tasks, Automation, and the Rise in U.S. Wage Inequality," *Econometrica*, 90(4), 1973–2016.

——— (2022b): "Tasks, Automation, and the Rise in US Wage Inequality," *Econometrica*, 90(5), 1973–2016.

ANDERSON, J. R. (1982): "Acquisition of Cognitive Skill," *Psychological Review*, 89(4), 369–406.

ARROW, K. J. (1962a): "The Economic Implications of Learning by Doing," *Review of Economic Studies*, 29(3), 155–173.

——— (1962b): "The Economic Implications of Learning by Doing," *Review of Economic Studies*, 29(3), 155–173.

AUTOR, D. H. (2015): "Why Are There Still So Many Jobs? The History and Future of Workplace Automation," *Journal of Economic Perspectives*, 29(3), 3–30.

Autor, D. H., F. Levy, and R. J. Murnane (2003a): "The Skill Content of Recent Technological Change: An Empirical Exploration," *Quarterly Journal of Economics*, 118(4), 1279–1333.

——— (2003b): "The Skill Content of Recent Technological Change: An Empirical Exploration," *Quarterly Journal of Economics*, 118(4), 1279–1333.

Babina, T., A. Fedyk, A. He, and J. Hodson (2024): "Artificial Intelligence, Firm Growth, and Product Innovation," *Journal of Financial Economics*, 151, 103745.

Brynjolfsson, E., B. Chandar, and R. Chen (2025a): "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence," *Stanford Digital Economy Lab Working Paper*.

Brynjolfsson, E., B. Chandar, and R. Chen (2025b): "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence," Discussion paper, Stanford Digitial Economy Lab Working Paper.

Caicedo, S., R. E. Lucas, and E. Rossi-Hansberg (2019): "Learning, Career Paths, and the Distribution of Wages," *American Economic Journal: Macroeconomics*, 11(1), 49–88.

Chari, V. V., and H. Hopenhayn (1991): "Vintage Human Capital, Growth, and the Diffusion of New Technology," *Journal of Political Economy*, 99(6), 1142–1165.

Ciccone, A., and K. Matsuyama (1996): "Start-Up Costs and Pecuniary Externalities as Barriers to Economic Development," *Journal of Development Economics*, 49(1), 33–59.

Cooper, A. (2025): "Why Anthropic CEO Dario Amodei spends so much time warning of AI's potential dangers," 60 Minutes, CBS News, Interview with Dario Amodei. Transcript available at: https://www.cbsnews.com/news/anthropic-ceo-dario-amodei-warning-of-ai-potential-dangers-60-minutes-transcript/.

Dewey, J. (1938): *Experience and Education*. Kappa Delta Pi.

Ericsson, K. A., R. T. Krampe, and C. Tesch-Römer (1993): "The Role of Deliberate Practice in the Acquisition of Expert Performance," *Psychological Review*, 100(3), 363–406.

Fossen, F. M., T. McLemore, and A. Sorgner (2024): "Artificial Intelligence and Entrepreneurship," IZA Discussion Paper 17055, Institute of Labor Economics (IZA), Bonn, Germany.

Friebel, G., Y. Huang, Y. Li, S. Shukla, and T. Zhang (2026): "Pyramids, Diamonds, and Oscillations: AI and the Internal Labor Market of the Firm," *Working Paper*.

Galor, O., and D. Tsiddon (1997): "The Distribution of Human Capital and Economic Growth," *Journal of Economic Growth*, 2(1), 93–124.

Garicano, L. (2000): "Hierarchies and the Organization of Knowledge in Production," *Journal of Political Economy*, 108(5), 874–904.

Garicano, L., and L. Rayo (2017a): "Relational Knowledge Transfers," *American Economic Review*, 107(9), 2695–2730.

——— (2017b): "Relational Knowledge Transfers," *American Economic Review*, 107(9), 2695–2730.

——— (2025a): "Training in the Age of AI," *CEPR Discussion Paper*, (20634).

——— (2025b): "Training in the Age of AI: A Theory of Apprenticeship Viability," Discussion Paper 20634, CEPR.

Garicano, L., and E. Rossi-Hansberg (2006): "Organization and inequality in a knowledge economy," *The Quarterly Journal of Economics*, 121(4), 1383–1435.

Goos, M., A. Manning, and A. Salomons (2014): "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring," *American Economic Review*, 104(8), 2509–2526.

Grant, A. (2025): "Sam Altman on the future of AI and humanity," ReThinking with Adam Grant (Podcast), Available at: https://www.youtube.com/watch?v=c0NqpG--Pzw.

Hassler, J., and J. V. Rodríguez Mora (2000): "Intelligence, Social Mobility, and Growth," *American Economic Review*, 90(4), 888–908.

Ide, E. (2025a): "Automation, AI, and the Intergenerational Transmission of Knowledge," *arXiv preprint*.

——— (2025b): "Automation, AI, and the Intergenerational Transmission of Knowledge," Working paper, IESE Business School, arXiv:2507.16078; CEPR Discussion Paper No. 20940.

Jovanovic, B., and Y. Nyarko (1996): "Learning by Doing and the Choice of Technology," *Econometrica*, 64(6), 1299–1310.

Lucas, Robert E., J. (1988a): "On the Mechanics of Economic Development," *Journal of Monetary Economics*, 22(1), 3–42.

LUCAS, R. E. (1988b): "On the Mechanics of Economic Development," *Journal of Monetary Economics*, 22(1), 3–42.

LUCAS, R. E., AND B. MOLL (2014): "Knowledge Growth and the Allocation of Time," *Journal of Political Economy*, 122(1), 1–51.

MOKYR, J., C. VICKERS, AND N. L. ZIEBARTH (2015): "The History of Technological Anxiety and the Future of Economic Growth: Is This Time Different?," *Journal of Economic Perspectives*, 29(3), 31–50.

REDDING, S. (1996): "The Low-Skill, Low-Quality Trap: Strategic Complementarities between Human Capital and R&D," *Economic Journal*, 106(435), 458–470.

ROMER, P. M. (1986): "Increasing Returns and Long-Run Growth," *Journal of Political Economy*, 94(5), 1002–1037.

ROOSE, K. (2025): "For Some Recent Graduates, the A.I. Apocalypse May Already Be Here," *New York Times*, May 30, 2025.

SHEN, J. H., AND A. TAMKIN (2026): "How AI Impacts Skill Formation," *arXiv preprint*.

STOKEY, N. L. (1988a): "Learning by Doing and the Introduction of New Goods," *Journal of Political Economy*, 96(4), 701–717.

——— (1988b): "Learning by Doing and the Introduction of New Goods," *Journal of Political Economy*, 96(4), 701–717.

TELFORD, T. (2026): "For the First Time in 50 Years, College Grads are Losing Their Edge," *Wasthington Post*, January 31, 2026.

YOUNG, A. (1991a): "Learning by Doing and the Dynamic Effects of International Trade," *Quarterly Journal of Economics*, 106(2), 369–405.

——— (1991b): "Learning by Doing and the Dynamic Effects of International Trade," *Quarterly Journal of Economics*, 106(2), 369–405.

——— (1993a): "Invention and Bounded Learning by Doing," *Journal of Political Economy*, 101(3), 443–472.

——— (1993b): "Invention and Bounded Learning by Doing," *Journal of Political Economy*, 101(3), 443–472.

ZEIRA, J. (1998): "Workers, Machines, and Economic Growth," *Quarterly Journal of Economics*, 113(4), 1091–1117.

A Proof of Lemma 1

Part (i): Irreversibility. In a stationary equilibrium, the HJB (6) reduces to $(\rho+\chi)V(x) = \max\{w_e x + \varphi x V'(x),\ w_m x^{\xi}\}$, which depends only on x. Since human capital is frozen during management, a managing household remains at x and faces the identical problem at every future instant. Therefore, the decision never reverses and the problem is a stopping time.

Part (ii): Threshold in x. By Part (i), the problem is a stopping time: the household works while $x < x^*$ and manages for $x \geq x^*$. In the worker region, the HJB (6) reduces to

$$(\rho+\chi)V(x) = w_e\, x + \varphi\, x\, V'(x), \qquad x \in [1, x^*) \tag{50}$$

In the manager region, $V(x) = w_m x^{\xi}/(\rho+\chi)$. At the boundary x^*, value matching and smooth pasting require

$$V(x^*) = \frac{w_m (x^*)^{\xi}}{\rho+\chi} \tag{51}$$

$$V'(x^*) = \frac{w_m\, \xi\, (x^*)^{\xi-1}}{\rho+\chi} \tag{52}$$

The general solution to (50) is

$$V(x) = C x^{\frac{\rho+\chi}{\varphi}} + \frac{w_e\, x}{\rho+\chi-\varphi} \tag{53}$$

where C is a constant of integration. Value matching (51) gives

$$C = \left[\frac{w_m (x^*)^{\xi}}{\rho+\chi} - \frac{w_e\, x^*}{\rho+\chi-\varphi}\right](x^*)^{-\frac{\rho+\chi}{\varphi}} \tag{54}$$

Differentiating (53):

$$V'(x) = C\frac{\rho+\chi}{\varphi}\, x^{\frac{\rho+\chi}{\varphi}-1} + \frac{w_e}{\rho+\chi-\varphi} \tag{55}$$

Evaluating at $x = x^*$ and substituting (54):

$$V'(x^*) = \left[\frac{w_m (x^*)^{\xi-1}}{\rho+\chi} - \frac{w_e}{\rho+\chi-\varphi}\right]\frac{\rho+\chi}{\varphi} + \frac{w_e}{\rho+\chi-\varphi} \tag{56}$$

Setting (56) equal to the smooth-pasting condition (52):

$$\left[\frac{w_m (x^*)^{\xi-1}}{\rho+\chi} - \frac{w_e}{\rho+\chi-\varphi}\right]\frac{\rho+\chi}{\varphi} + \frac{w_e}{\rho+\chi-\varphi} = \frac{w_m\, \xi\, (x^*)^{\xi-1}}{\rho+\chi}$$

The second and third terms on the left combine to $\frac{w_e}{\rho+\chi-\varphi}\left[1 - \frac{\rho+\chi}{\varphi}\right] = -w_e/\varphi$, leaving $w_m (x^*)^{\xi-1}/\varphi - w_e/\varphi$ on the left. Multiplying through by φ and rearranging:

$$w_m (x^*)^{\xi-1}\left(1 - \frac{\xi\varphi}{\rho+\chi}\right) = w_e \tag{57}$$

which is (16). Substituting (54) back into (53):

$$V(x) = \left[\frac{w_m (x^*)^{\xi}}{\rho+\chi} - \frac{w_e x^*}{\rho+\chi-\varphi}\right]\left(\frac{x}{x^*}\right)^{\frac{\rho+\chi}{\varphi}} + \frac{w_e x}{\rho+\chi-\varphi} \tag{58}$$

■

B Stationary Equilibrium Derivations

This appendix collects the factor supply and market clearing expressions used in Section 3.1. All derivations assume the stationary threshold policy of Lemma 1.

From the stationary distribution (17), the aggregate supply of managerial human capital is

$$s = x^{*\xi-\chi/\varphi} \tag{59}$$

and the aggregate supply of worker human capital is

$$X^s = \frac{\chi}{\varphi-\chi}\left(x^{*1-\chi/\varphi}-1\right) \tag{60}$$

Each worker-produced product demands the same human capital input (from (12)), so the firm's total demand is $(f-\tau^*)(w_e/(1-\eta^{-1}))^{-\eta}$. Using $\varphi = \delta(f-\tau^*)$, market clearing for worker human capital requires

$$\varphi\left(\frac{w_e}{1-\eta^{-1}}\right)^{-\eta} = \frac{\delta\chi}{\varphi-\chi}\left(x^{*1-\chi/\varphi}-1\right) \tag{61}$$

C Derivations Under $\psi = 0$, $\rho \downarrow 0$, $\xi \downarrow 1$

Under these simplifications, the stopping condition becomes $w_m = w_e(1-\varphi/\chi)^{-1}$, the managerial supply is $s = x^{*1-\chi/\varphi}$, and the frontier is $f = \bar{\tau}\, x^{*1-\chi/\varphi}$.

No automation. With $\tau^* = 0$, all products are worker-produced. The frontier FOC gives $\eta^{-1}(w_e/(1-\eta^{-1}))^{1-\eta} = \bar{\tau}^{-1} w_m$. Substituting $w_m = w_e(1-\varphi/\chi)^{-1}$, the growth rate equation $\varphi = \delta\bar{\tau}\, x^{*1-\chi/\varphi}$ and market clearing $\delta^{-1}\varphi(w_e/(1-\eta^{-1}))^{-\eta} = (1-\chi/\varphi)^{-1}(1-x^{*1-\chi/\varphi})$ together imply $\varphi = \delta\bar{\tau}/\eta$. The equilibrium exists provided $\varphi < \chi$ (i.e., $\delta\bar{\tau}/(\eta\chi) < 1$). Solving for w_e:

$$w_e = \left(\bar{\tau}\eta^{-1}\right)^{\eta^{-1}}(1-\eta^{-1})^{1-\eta^{-1}}\left(1-\frac{\bar{\tau}}{\eta}\frac{\delta}{\chi}\right)^{\eta^{-1}}$$

The condition $w_e \le q$ is equivalent to (22).

Interior automation. From the frontier FOC with $\psi = 0$: $qe^{\varepsilon\tau^*} = (\bar{\tau}\eta^{-1})^{\eta^{-1}}(1-\eta^{-1})^{1-\eta^{-1}}(1-\varphi/\chi)^{\eta^{-1}}$, which gives (26). Combining the growth rate equation $\varphi = \delta(f-\tau^*)$ with worker human capital market clearing gives (24). Setting these equal yields (27). ■

D Proofs of Propositions in Section 3

Proof of Proposition 1. The function g is defined on $(-\infty,\chi)$ with $g(\varphi)\to\infty$ as $\varphi\to\chi$. Since g is globally convex, it has at most two roots. An interior equilibrium exists, so g has at least one root in $(0,\chi)$, which requires $g(\varphi^*)<0$. By global convexity, g then has exactly two roots in $(-\infty,\chi)$.

If both roots lie in $(0,\chi)$, we have two interior equilibria with distinct growth rates (except at the knife-edge $g(\varphi^*)=0$). If only one root lies in $(0,\chi)$ and the other is non-positive, then $g(0)<0$, which gives $1+\ln\bar{q}<0$, i.e., $\ln\bar{q}<-1$. Since $\ln\bar{q}\le -1$ is precisely the condition (23), a full automation equilibrium with zero learning exists. ■

Proof of Proposition 2. By hypothesis $\ln\bar{q}<-1$, which gives $g(0)=1+\ln\bar{q}<0$. At the minimizer $\varphi=\varphi^*$,

$$g(\varphi^*)=(1+\ln\bar{q})+\bar{\varepsilon}^{-1}\left(1-\frac{\bar{\varepsilon}\chi}{\bar{\varphi}}+\ln\frac{\bar{\varepsilon}\chi}{\bar{\varphi}}\right).$$

The function $\ell(x)=1-x+\ln x$ is strictly concave on $(0,\infty)$ with unique maximum $\ell(1)=0$, so $\ell(x)\le 0$ for all $x>0$. Applying this to $x=\bar{\varepsilon}\chi/\bar{\varphi}$ shows the second term is non-positive, hence $g(\varphi^*)\le 1+\ln\bar{q}<0$. So g has two distinct roots in $(-\infty,\chi)$.

Since $g(0)<0$, at most one root can be non-positive. Suppose both roots were non-positive; then by convexity $g(0)>0$, a contradiction. Hence exactly one root lies in $(0,\chi)$, constituting the unique interior automation equilibrium. ■

Proof of Proposition 3. The function g in (27) depends on q only through $\ln\bar{q}$, which shifts g vertically. As q falls, g shifts down. Starting from the knife-edge case $g(\varphi^*)=0$ with a single automation equilibrium, as q falls further the two roots of the globally convex g diverge: the larger root (high-learning equilibrium) increases and the smaller root (low-learning equilibrium) decreases. ■

Proof of Proposition 4. Since $\bar{\varphi}<\chi$, the no-automation equilibrium exists if and only if $\ln\bar{q}\ge\bar{\varepsilon}^{-1}\ln(1-\bar{\varphi}/\chi)$ by (22), and the full automation equilibrium exists if and only if $\ln\bar{q}\le -1$ by (23). Under Assumption 1, $-1<\bar{\varepsilon}^{-1}\ln(1-\bar{\varphi}/\chi)$, so these two regions do not overlap.

Interior automation equilibria are roots of g in $(0,\bar{\varphi})$. Evaluating g at the boundaries:

$$g(0)=1+\ln\bar{q},\qquad g(\bar{\varphi})=\ln\bar{q}-\bar{\varepsilon}^{-1}\ln\left(1-\frac{\bar{\varphi}}{\chi}\right)$$

The derivative is $g'(\varphi)=-\bar{\varphi}^{-1}+\bar{\varepsilon}^{-1}(\chi-\varphi)^{-1}$, so $g'(0)=-\bar{\varphi}^{-1}+(\bar{\varepsilon}\chi)^{-1}$. Since $x<-\ln(1-x)$ for all $x\in(0,1)$, Assumption 1 implies $\bar{\varepsilon}\ge-\ln(1-\bar{\varphi}/\chi)>\bar{\varphi}/\chi$, hence $g'(0)<0$.

Case (i). The condition $\bar{\varepsilon}\ge\bar{\varphi}/(\chi-\bar{\varphi})$ is equivalent to $g'(\bar{\varphi})\le 0$. Since $g''>0$, the derivative g' is increasing, so $g'(\bar{\varphi})\le 0$ implies $g'\le 0$ on all of $[0,\bar{\varphi}]$. Thus g is strictly decreasing. A root in $(0,\bar{\varphi})$ exists if and only if $g(0)>0$ and $g(\bar{\varphi})<0$, i.e., $\ln\bar{q}\in(-1,\bar{\varepsilon}^{-1}\ln(1-\bar{\varphi}/\chi))$, and it is unique by monotonicity. Combined with the boundary equilibria, exactly one stationary equilibrium exists for each $\bar{q}$.

Case (ii). The condition $\bar{\varepsilon} < \bar{\varphi}/(\chi - \bar{\varphi})$ gives $g'(\bar{\varphi}) > 0$. Since $g'(0) < 0$ and g' is continuous and increasing, there is a unique $\varphi^* = \chi - \bar{\varphi}/\bar{\varepsilon} \in (0, \bar{\varphi})$ at which $g'(\varphi^*) = 0$, so g attains its minimum at φ^*.

For $\ln \bar{q} \in (-1, \bar{\varepsilon}^{-1} \ln(1 - \bar{\varphi}/\chi))$, we have $g(0) > 0$ and $g(\bar{\varphi}) < 0$. Because g decreases on $(0, \varphi^*)$ and increases on $(\varphi^*, \bar{\varphi})$, and $g(\bar{\varphi}) < 0$, the function remains negative to the right of its first zero crossing. The intermediate value theorem gives exactly one root, so the interior automation equilibrium is again unique.

For $\ln \bar{q} > \bar{\varepsilon}^{-1} \ln(1 - \bar{\varphi}/\chi)$, the no-automation equilibrium exists and both $g(0) > 0$ and $g(\bar{\varphi}) > 0$. Interior roots exist if and only if $g(\varphi^*) < 0$. Setting $g(\varphi^*) = 0$ and solving yields $\ln \bar{q} = \Lambda$ with $\Lambda = \chi/\bar{\varphi} - 1 - \bar{\varepsilon}^{-1}[1 + \ln(\chi\bar{\varepsilon}/\bar{\varphi})]$. For $\ln \bar{q} \in (\bar{\varepsilon}^{-1} \ln(1 - \bar{\varphi}/\chi), \Lambda)$, $g(\varphi^*) < 0$ and convexity gives exactly two interior roots, which coexist with the no-automation equilibrium. At $\ln \bar{q} = \Lambda$ the two roots collide at φ^*, and for $\ln \bar{q} > \Lambda$ only the no-automation equilibrium remains. ■

Proof of Proposition 5. Write $\alpha \equiv 1 - \eta^{-1}$, $u \equiv \varphi/\bar{\varphi} \in [0,1]$, $z \equiv 1 - \varphi/\chi$, and $\bar{z} \equiv 1 - \bar{\varphi}/\chi$.

Part (i). The bracket in (31) is

$$h(u) \equiv \frac{1 - e^{-\alpha\bar{\varepsilon}(1-u)}}{\alpha\bar{\varepsilon}} + u\, e^{-\alpha\bar{\varepsilon}(1-u)},$$

with derivative $h'(u) = u\,\alpha\bar{\varepsilon}\, e^{-\alpha\bar{\varepsilon}(1-u)} \geq 0$ on $[0,1]$, vanishing only at $u = 0$. For any $0 \leq u_1 < u_2 \leq 1$, $h(u_2) - h(u_1) = \int_{u_1}^{u_2} h'(s)\, ds > 0$ since h' is continuous and strictly positive on $(0,1]$. The prefactor of h in (31) does not depend on φ, so $\mathcal{V}$ is strictly increasing in φ on $[0, \bar{\varphi}]$.

Part (ii). From (20), $(q/(1-\eta^{-1}))^{1-\eta} = (\bar{\tau}\eta^{-1})^{-\alpha}(1-\eta^{-1})^{\alpha}\bar{q}^{-\alpha\bar{\varepsilon}}$. Substituting into (31) and dividing by (32),

$$\frac{\mathcal{V}(\varphi)}{\mathcal{V}^{NA}} = \bar{z}^{\alpha}\, \bar{q}^{-\alpha\bar{\varepsilon}}\, h(u).$$

The interior equilibrium condition $g(\varphi) = 0$ gives $\ln \bar{q} = u - 1 + \bar{\varepsilon}^{-1} \ln z$, equivalently $\bar{q}^{-\alpha\bar{\varepsilon}} = e^{\alpha\bar{\varepsilon}(1-u)} z^{-\alpha}$. Substituting and simplifying,

$$\frac{\mathcal{V}(\varphi)}{\mathcal{V}^{NA}} = \left(\frac{\bar{z}}{z}\right)^{\alpha} \left[1 + \frac{e^{y} - 1 - y}{\alpha\bar{\varepsilon}}\right], \qquad y \equiv \alpha\bar{\varepsilon}(1-u) \geq 0. \tag{62}$$

The no-automation equilibrium coexists with an automation equilibrium only when $\ln \bar{q} \geq \bar{\varepsilon}^{-1} \ln \bar{z}$, by (22). Combined with $g(\varphi) = 0$, this gives $y \leq \alpha \ln(z/\bar{z})$, so $(\bar{z}/z)^{\alpha} \leq e^{-y}$. Applying this bound to (62), it suffices to show

$$(1 - \alpha\bar{\varepsilon})(e^{y} - 1) < y \qquad \text{for all } y \in (0, \alpha\bar{\varepsilon}]. \tag{63}$$

If $\alpha\bar{\varepsilon} \geq 1$, the left-hand side is non-positive and (63) is immediate. If $\alpha\bar{\varepsilon} < 1$, the function $y \mapsto y - (1 - \alpha\bar{\varepsilon})(e^{y} - 1)$ vanishes at $y = 0$ and has derivative $1 - (1 - \alpha\bar{\varepsilon})e^{y} > 0$ on $[0, -\ln(1 - \alpha\bar{\varepsilon}))$.

Since $x < -\ln(1-x)$ for all $x \in (0,1)$, the derivative is positive throughout $[0, \alpha\bar{\varepsilon}]$ and (63) follows. Finally, Assumption 1 gives $\bar{\varepsilon} > -\ln\bar{z}$, so $y \leq \alpha\ln(z/\bar{z}) \leq -\alpha\ln\bar{z} < \alpha\bar{\varepsilon}$, placing y in the required range. ■

Proof of Proposition 6. Retain the notation of Appendix D.

Part (i). Differentiating (33) with respect to φ, $\partial V(1;\varphi)/\partial\varphi = (\alpha\bar{\varepsilon}/\bar{\varphi})\, V(1;\varphi) > 0$ on $[0,\bar{\varphi}]$.

Part (ii). Applying the substitutions from the proof of Proposition 5 to (33) and (34),

$$\frac{V(1;\varphi)}{V^{NA}(1)} = \bar{z}^{\alpha}\,\bar{q}^{-\alpha\bar{\varepsilon}}\,e^{-\alpha\bar{\varepsilon}(1-u)} = \left(\frac{\bar{z}}{z}\right)^{\alpha}.$$

For $\varphi < \bar{\varphi}$, $z > \bar{z}$, so the ratio is strictly less than unity. ■

E Proofs for Section 5

Proof of Lemma 2. By definition,

$$\mathcal{W}_0 = \int_{-\infty}^{\infty} e^{-\rho x} U_x^w dx \tag{64}$$

$$= \int_{-\infty}^{\infty} e^{-\rho x}\left(\int_x^{\infty} e^{-(\rho+\chi)(t-x)} c_{t,x} dt\right) dx \tag{65}$$

$$= \int_{-\infty}^{\infty}\int_x^{\infty} e^{-\rho t} e^{-\chi(t-x)} c_{t,x} dt dx \tag{66}$$

$$= \int_{-\infty}^{\infty}\int_x^{\infty} e^{-\rho t} e^{-\chi(t-x)} c_{t,x}\mathbb{1}[t \geq 0] dt dx + \underbrace{\int_{-\infty}^{\infty}\int_x^{\infty} e^{-\rho t} e^{-\chi(t-x)} c_{t,x}\mathbb{1}[t < 0] dt dx}_{t.i.p.} \tag{67}$$

$$= \int_0^{\infty} e^{-\rho t}\int_{-\infty}^{t} e^{-\chi(t-x)} c_{t,x} dx dt + t.i.p \tag{68}$$

With a unit-measure of workers and birth rate χ, the mass of cohort x alive at time t is $\chi e^{-\chi(t-x)}$, so the aggregate consumption of workers at time t is $\int_{-\infty}^{t} e^{-\chi(t-x)} c_{t,x} dx = \frac{C_t^w}{\chi}$, where C_t^w is the agregate consumption of workers. Since the consumption of firms' owner is zero, we have that $C_t^w = C_t$. Using this result and defining $C_t = C_t^w + C_t^o$ as aggregate consumption, we have

$$\mathcal{W} = \int_0^{\infty} e^{-\rho t}\frac{C_t}{\chi} dt + t.i.p \tag{69}$$

$$= \frac{1}{\chi}\int_0^{\infty} e^{-\rho t} C_t dt + t.i.p. \tag{70}$$

since $\chi > 0$ is a constant, the planner's problem can be rewritten as maximizing $\int_0^{\infty} e^{-\rho t} C_t dt$. ■

Given an initial distribition of human capital $N_0(x)$ with $\int N_0(x)dx = 1$, the planner problem with its Lagrange multiplies is given by:

$$\max_{\mathscr{A}_p} \int_0^\infty e^{-\rho t} \left[\int_0^{\tilde{\tau}^*(t)} \left(\tilde{y}(\tau,t)^{1-\eta^{-1}} - \frac{q y(\tau,t)}{z(\tau)} \right) d\tau + \int_{\tilde{\tau}^*(t)}^{f(t)} \left(\tilde{y}(\tau,t)^{1-\eta^{-1}} \right) d\tau - q \tilde{a}_m(t) \right]$$

$s.t.$

$$-\chi\left(1-\tilde{N}(x,t)\right) + \tilde{\varphi}(t) x \partial_x \tilde{N}_e(x,t) + \partial_t \tilde{N}(x,t) = 0, \qquad \forall t > 0 \qquad (e^{-\rho t}\tilde{v}(x,t))$$

$$\int h(x) d\left(\tilde{N}(x,t) - \tilde{N}_e(x,t)\right) - s(t) = 0, \qquad \forall t \geq 0 \qquad (e^{-\rho t}\tilde{w}_m(t))$$

$$\int x d\tilde{N}_e(x,t) - \int_{\tilde{\tau}^*(t)}^{f(t)} \tilde{y}(\tau,t) d\tau = 0, \qquad \forall t \geq 0 \qquad (e^{-\rho t}\tilde{w}_e(t))$$

$$\delta(f(t) - \tilde{\tau}^*(t)) - \varphi(t) = 0, \qquad \forall t \geq 0 \qquad \left(e^{-\rho t} \frac{\tilde{v}_\varphi(t)}{\eta} \left(\frac{\tilde{w}_e(t)}{1-\eta^{-1}}\right)^{1-\eta}\right)$$

$$d\tilde{N}(x,t) \geq 0, \qquad \forall t \geq 0 \qquad (e^{-\rho t}\tilde{v}_N(x,t))$$

$$d\tilde{N}_e(x,t) \geq 0, \qquad \forall t \geq 0 \qquad (e^{-\rho t}\tilde{v}_e(x,t))$$

$$d\tilde{N}(x,t) - d\tilde{N}_e(x,t) \geq 0, \qquad \forall t \geq 0 \qquad (e^{-\rho t}\tilde{v}_m(x,t))$$

$$\bar{\tau}\left[\psi^{\omega^{-1}} \tilde{a}_m(t)^{1-\omega^{-1}} + (1-\psi)^{\omega^{-1}} \tilde{s}(t)^{1-\omega^{-1}}\right]^{\frac{1}{1-\omega^{-1}}} - f(t) = 0, \qquad \forall t \geq 0, \qquad (e^{-\rho t}\tilde{c}(t))$$

$$\tilde{N}(x,0) = N_0(x), \qquad \forall x \qquad (\tilde{v}_{N_0}(x))$$

Together with transversality conditions

$$\lim_{T\to\infty} e^{-\rho T} v(x,T)\tilde{N}(x,T)dx = 0 \quad \text{for all } x \tag{71}$$

We now derive the optimality condition in two blocks. The first lemma characterizes the first order conditions for $\tilde{y}(\tau,t)$, $\tilde{\tau}^*$, $\tilde{f}(t)$, $\tilde{a}_m(t)$, and $\tilde{s}(t)$. The second proposition characterizes the first order conditions for $\tilde{N}(x,t)$ and $\tilde{N}_e(x,t)$. For the next two lemmas, we focus in the optimal policy such that $\tilde{N}_e(x,t)$ has bounded support for all $t \geq 0$, and $v(x,t)$ is locally integrable, continues in x, and differentiable in t.

Lemma 3. *Define $\tilde{\kappa}(t) := 1 - \left(1 + \delta\tilde{v}_\varphi(t)\right)^{\frac{1}{1-\eta}}$. The first order conditions for $\tilde{y}(\tau,t)$, $\tilde{\tau}^*$, $\tilde{f}(t)$, $\tilde{a}_m(t)$, and $\tilde{s}(t)$ are given by*

$$y(\tau,t) = \begin{cases} \left(\frac{qe^{\varepsilon\tau}}{1-\eta^{-1}}\right)^{-\eta}, & \tau < \tilde{\tau}^*(t) \\ \left(\frac{\tilde{w}_e(t)}{1-\eta^{-1}}\right)^{-\eta}, & f(t) \geq \tau > \tilde{\tau}^*(t) \end{cases}$$

$$qe^{\varepsilon\tilde{\tau}^*(t)} \begin{cases} \geq \tilde{w}_e(t)(1-\tilde{\kappa}(t)) & \textit{if } \tilde{\tau}^* \geq 0 \\ \leq \tilde{w}_e(t)(1-\tilde{\kappa}(t)) & \textit{if } \tilde{\tau}^* \leq \tilde{f}(t) \end{cases}$$

$$\tilde{p}_f(t) = \begin{cases} \eta^{-1}\left(\frac{(1-\tilde{\kappa}(t))\tilde{w}_e(t)}{1-\eta^{-1}}\right)^{1-\eta}, & \textit{if } \tilde{\tau}^* < \tilde{f}(t) \\ \eta^{-1}\left(\frac{\tilde{w}_e(t)}{1-\eta^{-1}}\right)^{1-\eta}\left[\left(\frac{qe^{\varepsilon\tilde{\tau}^*(t)}}{\tilde{w}_e(t)}\right)^{1-\eta} + (1-\tilde{\kappa}(t))^{1-\eta} - 1\right], & \textit{if } \tilde{\tau}^* = \tilde{f}(t) \end{cases}$$

where

$$\tilde{a}_m = \left(\frac{\tilde{p}_f(t)\,\bar{\tau}\,\psi^{\omega^{-1}}}{q}\right)^{\omega} \frac{f(t)}{\bar{\tau}}, \tag{72}$$

$$\tilde{s} = \left(\frac{\tilde{p}_f(t)\,\bar{\tau}(1-\psi)^{\omega^{-1}}}{\tilde{w}_m}\right)^{\omega} \frac{f(t)}{\bar{\tau}}, \tag{73}$$

$$\tilde{p}_f(t) = \frac{1}{\bar{\tau}}\left[\psi\, q^{1-\omega} + (1-\psi)\,\tilde{w}_m(t)^{1-\omega}\right]^{\frac{1}{1-\omega}} \tag{74}$$

Proof. We now characterize the first order conditions for $\tilde{y}(\tau,t)$, $\tilde{\tau}^*$, $\tilde{f}(t)$, $\tilde{a}_m(t)$, and $\tilde{s}(t)$.

- FOCs $\tilde{y}(\tau,t)$: For each τ, the optimality condition is given by

$$(1-\eta^{-1})\,\tilde{y}(\tau,t)^{-\eta^{-1}} = \frac{q}{z(\tau)}, \forall \tau < \tilde{\tau}^*$$
$$(1-\eta^{-1})\,\tilde{y}(\tau,t)^{-\eta^{-1}} = \tilde{w}_e(t), \forall \tau > \tilde{\tau}^*$$

  Observe that output is discontinuous at $\tilde{\tau}^*$ (see the condition for $\tilde{\tau}^*$ below). Output is given by

$$y(\tau,t) = \begin{cases} \left(\frac{qe^{\varepsilon\tau}}{1-\eta^{-1}}\right)^{-\eta}, & \tau < \tilde{\tau}^*(t) \\ \left(\frac{\tilde{w}_e(t)}{1-\eta^{-1}}\right)^{-\eta}, & f(t) \geq \tau > \tilde{\tau}^*(t) \end{cases}$$

  and net output for task $\tau < \tilde{\tau}^*$ is given by

$$\tilde{y}(\tau,t)^{1-\eta^{-1}} - \frac{q\,y(\tau,t)}{z(\tau)} = \eta^{-1}\left(\frac{qe^{\varepsilon\tau}}{1-\eta^{-1}}\right)^{1-\eta} \tag{75}$$

  and for task $\tau > \tilde{\tau}^*$ is given by

$$\tilde{y}(\tau,t)^{1-\eta^{-1}} - \tilde{w}_e\tilde{y}(\tau,t) = \eta^{-1}\left(\frac{\tilde{w}_e(t)}{1-\eta^{-1}}\right)^{1-\eta} \tag{76}$$

- FOC $\tilde{\tau}^*$: If we write the relevant part for $\tilde{\tau}^*$ of the Lagrangian as

$$\mathscr{L} = \ldots \int_0^{\infty} e^{-\rho t}\Big[\int_0^{\tilde{\tau}^*(t)}\left(\tilde{y}(\tau,t)^{1-\eta^{-1}} - \frac{q\,y(\tau,t)}{z(\tau)}\right)d\tau + \int_{\tilde{\tau}^*(t)}^{f(t)}\left(\tilde{y}(\tau,t)^{1-\eta^{-1}}\right)d\tau$$
$$- w_e(t)\int_{\tilde{\tau}^*(t)}^{f(t)}\tilde{y}(\tau,t)d\tau - \delta\tilde{\tau}^*(t)\frac{\tilde{v}_\varphi(t)}{\eta}\left(\frac{\tilde{w}_e(t)}{1-\eta^{-1}}\right)^{1-\eta}\Big]dt$$

  Assume that $\tilde{\tau}^*(t) \geq 0$, the FOC for $\tilde{\tau}^*$ is given by

$$\lim_{\tau\uparrow\tau^*(t)}\left(\tilde{y}(\tau,t)^{1-\eta^{-1}} - \frac{q\tilde{y}(\tau,t)}{z(\tau)}\right)$$
$$- \lim_{\tau\downarrow\tau^*(t)}\left(\tilde{y}(\tau,t)^{1-\eta^{-1}} - \tilde{w}_e(t)\tilde{y}(\tau,t)\right) - \delta\frac{\tilde{v}_\varphi(t)}{\eta}\left(\frac{\tilde{w}_e(t)}{1-\eta^{-1}}\right)^{1-\eta} \leq 0$$

and using the net output from the FOC for $\tilde{y}(\tau,t)$, we can rewrite the FOC for $\tilde{\tau}^*$ as

$$\lim_{\tau\uparrow\tau^*(t)}\left(\eta^{-1}\left(\frac{qe^{\varepsilon\tau}}{1-\eta^{-1}}\right)^{1-\eta}\right) - \lim_{\tau\downarrow\tau^*(t)}\left(\eta^{-1}\left(\frac{\tilde{w}_e(t)}{1-\eta^{-1}}\right)^{1-\eta}\right) - \delta\frac{\tilde{v}_\varphi(t)}{\eta}\left(\frac{\tilde{w}_e(t)}{1-\eta^{-1}}\right)^{1-\eta} \le 0$$

or equivalently

$$\begin{aligned}
\eta^{-1}\left(\frac{qe^{\varepsilon\tilde{\tau}^*(t)}}{1-\eta^{-1}}\right)^{1-\eta} &\le \eta^{-1}\left(\frac{\tilde{w}_e(t)}{1-\eta^{-1}}\right)^{1-\eta} + \delta\eta^{-1}\tilde{v}_\varphi(t)\left(\frac{\tilde{w}_e(t)}{1-\eta^{-1}}\right)^{1-\eta} && \iff \\
qe^{\varepsilon\tilde{\tau}^*(t)} &\ge \tilde{w}_e(t)\left(1+\delta\tilde{v}_\varphi(t)\right)^{\frac{1}{1-\eta}} \quad \text{(since } \eta>1) && \iff \\
qe^{\varepsilon\tilde{\tau}^*(t)} &\ge \tilde{w}_e(t)(1-\tilde{\kappa}(t)) \text{ if } \tilde{\tau}^*(t)\ge 0
\end{aligned}$$

Applying similar steps, we have that

$$qe^{\varepsilon\tilde{\tau}^*(t)} \le \tilde{w}_e(t)(1-\tilde{\kappa}(t)) \text{ if } \tilde{\tau}^*(t)\le \tilde{f}(t) \tag{77}$$

- FOCs $\tilde{f}(t)$, $\tilde{a}_m(t)$, and $\tilde{s}(t)$: If we write the relevant part of the Lagrangian

$$\mathscr{L} = \cdots\int_0^\infty e^{-\rho t}\Big[\int_{\tilde{\tau}^*(t)}^{\tilde{f}(t)}\left(\tilde{y}(\tau,t)^{1-\eta^{-1}}\right)d\tau - w_e(t)\int_{\tilde{\tau}^*(t)}^{\tilde{f}(t)}\tilde{y}(\tau,t)d\tau + \frac{\tilde{v}_\varphi(t)}{\eta}\left(\frac{\tilde{w}_e(t)}{1-\eta^{-1}}\right)^{1-\eta}\delta\tilde{f}(t)$$
$$\cdots + \tilde{p}_f(t)\left(\bar{\tau}\left[\psi^{\omega^{-1}}\tilde{a}_m(t)^{1-\omega^{-1}} + (1-\psi)^{\omega^{-1}}\tilde{s}(t)^{1-\omega^{-1}}\right]^{\frac{1}{1-\omega^{-1}}} - \tilde{f}(t)\right) - \tilde{w}_m\tilde{s}(t) - q\tilde{a}_m\Big]dt$$

The FOCs for these variables are given by

$$\begin{aligned}
\tilde{f}(t): &\qquad \tilde{p}_f(t) = \eta^{-1}\left(\frac{\tilde{w}_e(t)}{1-\eta^{-1}}\right)^{1-\eta} + \frac{\tilde{v}_\varphi(t)}{\eta}\left(\frac{\tilde{w}_e(t)}{1-\eta^{-1}}\right)^{1-\eta}\delta \\
\tilde{a}_m(t): &\qquad q = \tilde{p}_f(t)\,\bar{\tau}\left[\psi^{\omega^{-1}}\tilde{a}_m^{1-\omega^{-1}} + (1-\psi)^{\omega^{-1}}\tilde{s}^{1-\omega^{-1}}\right]^{\frac{\omega^{-1}}{1-\omega^{-1}}}\psi^{\omega^{-1}}\tilde{a}_m^{-\omega^{-1}} \\
\tilde{s}(t): &\qquad \tilde{w}_m = \tilde{p}_f(t)\,\bar{\tau}\left[\psi^{\omega^{-1}}\tilde{a}_m^{1-\omega^{-1}} + (1-\psi)^{\omega^{-1}}\tilde{s}^{1-\omega^{-1}}\right]^{\frac{\omega^{-1}}{1-\omega^{-1}}}(1-\psi)^{\omega^{-1}}\tilde{s}^{-\omega^{-1}}
\end{aligned}$$

Let $\gamma \equiv 1-\omega^{-1} = (\omega-1)/\omega$ and $g \equiv \tilde{f}(t)/\bar{\tau} = \left[\psi^{1-\gamma}\tilde{a}_m^\gamma + (1-\psi)^{1-\gamma}\tilde{s}^\gamma\right]^{1/\gamma}$, so the FOCs for $\tilde{a}_m$ and $\tilde{s}$ become

$$q = \tilde{p}_f(t)\,\bar{\tau}\,g^{1-\gamma}\psi^{1-\gamma}\tilde{a}_m^{\gamma-1}, \qquad \tilde{w}_m = \tilde{p}_f(t)\,\bar{\tau}\,g^{1-\gamma}(1-\psi)^{1-\gamma}\tilde{s}^{\gamma-1}$$

Multiplying through by $\tilde{a}_m$ and $\tilde{s}$ respectively and summing (using $g^\gamma = \psi^{1-\gamma}\tilde{a}_m^\gamma + (1-\psi)^{1-\gamma}\tilde{s}^\gamma$, i.e. Euler's theorem for the homogeneous-of-degree-one CES):

$$q\tilde{a}_m + \tilde{w}_m\tilde{s} = \tilde{p}_f(t)\,\bar{\tau}\,g^{1-\gamma}\cdot g^\gamma = \tilde{p}_f(t)\,\tilde{f}(t)$$

so the total cost equals $\tilde{p}_f(t)\,\tilde{f}(t)$ and $\tilde{p}_f(t)$ is the unit cost of the frontier. From each FOC, the cost-minimizing inputs are

$$\tilde{a}_m = \left(\frac{\tilde{p}_f(t)\,\bar{\tau}\,\psi^{1-\gamma}}{q}\right)^{1/(1-\gamma)} g, \qquad \tilde{s} = \left(\frac{\tilde{p}_f(t)\,\bar{\tau}(1-\psi)^{1-\gamma}}{\tilde{w}_m}\right)^{1/(1-\gamma)} g.$$

Substituting back into the definition of g and cancelling g:

$$1 = \left(\tilde{p}_f(t)\,\bar{\tau}\right)^{\gamma/(1-\gamma)} \left[\psi\, q^{-\gamma/(1-\gamma)} + (1-\psi)\,\tilde{w}_m^{-\gamma/(1-\gamma)}\right].$$

Using $\gamma/(1-\gamma) = \omega - 1$:

$$\left(\tilde{p}_f(t)\,\bar{\tau}\right)^{\omega-1} = \left[\psi\, q^{1-\omega} + (1-\psi)\,\tilde{w}_m^{1-\omega}\right]^{-1},$$

and therefore

$$\tilde{p}_f(t) = \frac{1}{\bar{\tau}}\left[\psi\, q^{1-\omega} + (1-\psi)\,\tilde{w}_m(t)^{1-\omega}\right]^{\frac{1}{1-\omega}},$$

which is the standard CES unit-cost function (14). Using the first FOC for $\tilde{\tau}^* < \tilde{f}(t)$, the frontier price is given by

$$\begin{aligned}
\frac{1}{\bar{\tau}}\left[\psi\, q^{1-\omega} + (1-\psi)\,\tilde{w}_m(t)^{1-\omega}\right]^{\frac{1}{1-\omega}} = \tilde{p}_f(t) = \eta^{-1}\left(\frac{\tilde{w}_e(t)}{1-\eta^{-1}}\right)^{1-\eta} + \frac{\tilde{v}_\varphi(t)}{\eta}\left(\frac{\tilde{w}_e(t)}{1-\eta^{-1}}\right)^{1-\eta}\delta, \quad &\iff \\
\tilde{p}_f(t) = \eta^{-1}\left(\frac{\left(1+\delta\tilde{v}_\varphi(t)\right)^{\frac{1}{1-\eta}}\,\tilde{w}_e(t)}{1-\eta^{-1}}\right)^{1-\eta} \quad &\iff \\
\tilde{p}_f(t) = \eta^{-1}\left(\frac{(1-\tilde{\kappa}(t))\,\tilde{w}_e(t)}{1-\eta^{-1}}\right)^{1-\eta}
\end{aligned}$$

and if $\tilde{\tau}^* = \tilde{f}(t)$

$$\begin{aligned}
\tilde{p}_f(t) &= \eta^{-1}\left(\frac{qe^{\varepsilon\tilde{\tau}^*(t)}}{1-\eta^{-1}}\right)^{1-\eta} + \frac{\tilde{v}_\varphi(t)}{\eta}\left(\frac{\tilde{w}_e(t)}{1-\eta^{-1}}\right)^{1-\eta}\delta \\
&= \eta^{-1}\left(\frac{qe^{\varepsilon\tilde{\tau}^*(t)}\,\tilde{w}_e(t)}{\tilde{w}_e(t)(1-\eta^{-1})}\right)^{1-\eta} + \frac{(1-\tilde{\kappa}(t))^{1-\eta}-1}{\eta}\left(\frac{\tilde{w}_e(t)}{1-\eta^{-1}}\right)^{1-\eta}\delta \\
&= \eta^{-1}\left(\frac{\tilde{w}_e(t)}{1-\eta^{-1}}\right)^{1-\eta}\left[\left(\frac{qe^{\varepsilon\tilde{\tau}^*(t)}}{\tilde{w}_e(t)}\right)^{1-\eta} + (1-\tilde{\kappa}(t))^{1-\eta} - 1\right]
\end{aligned}$$

■

Lemma 4. *Let $\tilde{V}(x,t) := -\int_x^\infty \tilde{v}(x',t)\,dx'$. Then the optimal stopping time for a worker with human capital x at time t in the planner's problem is given by*

$$(\rho+\chi)\tilde{V}(x,t) = \max\left\{\tilde{w}_m(t)h(x), \tilde{\varphi}(t)x\tilde{V}_x(x,t) + \tilde{w}_e(t)x\right\} + \partial_t\tilde{V}(x,t) \tag{78}$$

(x,t) in the support of $\tilde{N}(x,t)$ and

$$\frac{(\tilde{w}_e(t)(1-\tilde{\kappa}(t)))^{1-\eta} - \tilde{w}_e(t)^{1-\eta}}{\delta} = \eta(1-\eta^{-1})^{1-\eta}\int_1^\infty x\tilde{V}_x(x,t)\,d\tilde{N}_e(x,t) \tag{79}$$

Proof. We now characterize the first order conditions

- FOCs for $dN(x,t)$: The Lagrangian terms involving $N(x,t)$ are:

$$\mathscr{L} = \cdots \int_0^\infty e^{-\rho t} \int_0^\infty \left(\chi \tilde{N}(x,t) + \partial_t \tilde{N}(x,t)\right) \tilde{v}(x,t)\, dx\, dt$$
$$+ \int_0^\infty \int_0^\infty e^{-\rho t} \left[\tilde{w}_m(t) h(x) + \tilde{v}_N(x,t) + \tilde{v}_m(x,t)\right] d\tilde{N}(x,t)\, dt$$

Step 1: IBP in t to remove $\partial_t \tilde{N}$. Focus on the $\partial_t \tilde{N}$ term. Switching the order of integration and integrating by parts in t:

$$\begin{aligned}
\int_0^\infty \int_0^\infty e^{-\rho t} \partial_t \tilde{N}(x,t)\, \tilde{v}(x,t)\, dx\, dt &= \int_0^\infty \left(\int_0^\infty e^{-\rho t} \partial_t \tilde{N}(x,t)\, \tilde{v}(x,t)\, dt\right) dx \\
&= \int_0^\infty \left(\left[e^{-\rho t} \tilde{v}(x,t) \tilde{N}(x,t)\right]_{t=0}^{t=\infty} - \int_0^\infty \partial_t\left(e^{-\rho t} \tilde{v}(x,t)\right) \tilde{N}(x,t)\, dt\right) dx \\
&= -\int_0^\infty \int_0^\infty e^{-\rho t} \left(-\rho \tilde{v}(x,t) + \partial_t \tilde{v}(x,t)\right) \tilde{N}(x,t)\, dt\, dx \\
&= \int_0^\infty \int_0^\infty e^{-\rho t} \left(\rho \tilde{v}(x,t) - \partial_t \tilde{v}(x,t)\right) \tilde{N}(x,t)\, dt\, dx
\end{aligned}$$

where the boundary term $\left[e^{-\rho t} \tilde{v}(x,t) \tilde{N}(x,t)\right]_{t=0}^{t=\infty} = 0$ by the transversality condition (upper limit) and because $\tilde{N}(x,0)$ is predetermined (lower limit). Combining with the $\chi \tilde{N}$ term:

$$\int_0^\infty \int_0^\infty e^{-\rho t} \left(\chi \tilde{N}(x,t) + \partial_t \tilde{N}(x,t)\right) \tilde{v}(x,t)\, dx\, dt = \int_0^\infty \int_0^\infty e^{-\rho t} \left[(\rho + \chi) \tilde{v}(x,t) - \partial_t \tilde{v}(x,t)\right] \tilde{N}(x,t)\, dx\, dt$$

Step 2: IBP in x to convert $N(x,t)\,dx$ to $dN(x,t)$. We have $N(x,t) = \int_0^x dN(x',t)$, so switching the order of integration:

$$\int_0^\infty g(x,t)\, N(x,t)\, dx = \int_0^\infty g(x,t) \int_0^x dN(x',t)\, dx = \int_0^\infty \underbrace{\left(\int_x^\infty g(x',t)\, dx'\right)}_{\equiv G(x,t)} dN(x,t)$$

Applying this with $g(x,t) = \left[(\rho + \chi) \tilde{v}(x,t) - \partial_t \tilde{v}(x,t)\right]$ and defining

$$\tilde{V}(x,t) := -\int_x^\infty \tilde{v}(x',t)\, dx'$$

we get $G(x,t) = \left[-(\rho + \chi) \tilde{V}(x,t) + \partial_t \tilde{V}(x,t)\right]$.[11] Therefore:

$$\int_0^\infty \int_0^\infty e^{-\rho t} \left[(\rho + \chi) \tilde{v} - \partial_t \tilde{v}\right] N\, dx\, dt = \int_0^\infty \int_0^\infty e^{-\rho t} \left[-(\rho + \chi) \tilde{V}(x,t) + \partial_t \tilde{V}(x,t)\right] dN(x,t)\, dt$$

Step 3: Collect the coefficient of $dN(x,t)$. Putting Steps 1–2 together with the remaining constraint terms:

$$\mathscr{L} = \cdots \int_0^\infty \int_0^\infty e^{-\rho t} \Big[-(\rho + \chi) \tilde{V}(x,t) + \partial_t \tilde{V}(x,t) + \tilde{w}_m(t) h(x) + \tilde{v}_N(x,t) + \tilde{v}_m(x,t)\Big] dN(x,t)\, dt$$

Since $dN(x,t) \geq 0$, optimality requires the coefficient to be ≤ 0, with equality when

[11] Mechanically, we are imposing that $\lim_{x \to \infty} \tilde{V}(x,t) = 0$. Notice that for all $x > \tilde{x}^*(t)$ the KFE left hand side of is reduce to $0 = 0$ so we choose the lagrange multiplier accordingly to have this condition satisfied.

$dN(x,t) > 0$:

$$-(\rho+\chi)\tilde{V}(x,t) + \partial_t \tilde{V}(x,t) + \tilde{w}_m(t)h(x) + \tilde{v}_m(x,t) = -\tilde{v}_N(x,t) \leq 0$$

with $\nu_N(x,t) = 0$ if $dN(x,t) > 0$.

- FOCs for $d\tilde{N}_e(x,t)$: The Lagrangian terms involving $\tilde{N}_e(x,t)$ are

$$\begin{aligned}\mathscr{L} = \cdots \int_0^\infty e^{-\rho t} \int_0^\infty (\tilde{v}(x,t))\tilde{\varphi}(t)x\partial_x \tilde{N}_e(x,t)\,dx\,dt \\ + \int_0^\infty \int_0^\infty e^{-\rho t}[-\tilde{w}_m(t)h(x) + \tilde{w}_e(t)x + \tilde{v}_e(x,t) + \tilde{v}_m(x,t)]\,d\tilde{N}_e(x,t)\,dt\end{aligned}$$

We repeat the same steps as for $dN(x,t)$. Applying integration by parts in x to

$$\int_0^\infty (\tilde{v}(x,t))\,\tilde{\varphi}(t)\,x\partial_x \tilde{N}_e(x,t)\,dx \tag{80}$$

converts $\partial_x \tilde{N}_e(x,t)\,dx$ into $d\tilde{N}_e(x,t)$:

$$\int_0^\infty (\tilde{v}(x,t))\,\tilde{\varphi}(t)\,x\partial_x \tilde{N}_e(x,t)\,dx = \Big[\tilde{\varphi}(t)\,x\,\tilde{v}(x,t)\,\tilde{N}_e(x,t)\Big]_{x=0}^{\infty} - \int_0^\infty \tilde{\varphi}(t)\,\frac{\partial(x\tilde{v}(x,t))}{\partial x}\,\tilde{N}_e(x,t)\,dx.$$

Assuming bounded support for $\tilde{N}_e(x,t)$ with $\lim_{x\to\infty}\tilde{v}(x,t) = 0$, the boundary term vanishes. Writing $\tilde{N}_e(x,t) = \int_0^x d\tilde{N}_e(x',t)$ and swapping the order of integration,

$$\begin{aligned}\int_0^\infty \tilde{\varphi}(t)\,\frac{\partial(x\tilde{v}(x,t))}{\partial x}\,\tilde{N}_e(x,t)\,dx &= \int_0^\infty \tilde{\varphi}(t)\int_{x'}^\infty \frac{\partial(x\tilde{v}(x,t))}{\partial x}\,dx\,d\tilde{N}_e(x',t) \\ &= \int_0^\infty \tilde{\varphi}(t)\,x\,\tilde{v}(x,t)\,d\tilde{N}_e(x,t),\end{aligned}$$

where the last equality uses $\int_{x'}^\infty \partial_x(x\tilde{v})\,dx = [x\tilde{v}]_{x'}^\infty = -x'\tilde{v}(x',t)$ under the boundary assumption above. Substituting $\partial_x \tilde{V}(x,t) = \tilde{v}(x,t)$,

$$\int_0^\infty (\tilde{v})\,\tilde{\varphi}\,x\partial_x \tilde{N}_e\,dx = \int_0^\infty \tilde{\varphi}(t)\,x\,\tilde{v}(x,t)\,d\tilde{N}_e(x,t) = \int_0^\infty \tilde{\varphi}(t)\,x\partial_x \tilde{V}(x,t)\,d\tilde{N}_e(x,t).$$

Therefore,

$$\mathscr{L} = \cdots + \int_0^\infty \int_0^\infty e^{-\rho t}\left[\tilde{\varphi}(t)x\partial_x \tilde{V}(x,t) - \tilde{w}_m(t)h(x) + \tilde{w}_e(t)x + \tilde{v}_e(x,t) - \tilde{v}_m(x,t)\right]d\tilde{N}_e(x,t)\,dt$$

so

$$\tilde{\varphi}(t)x\tilde{V}_x(x,t) - \tilde{w}_m h(x) + \tilde{w}_e(t)x - \tilde{v}_m(x,t) = -\tilde{v}_e(x,t) \leq 0$$

with $\tilde{v}_e(x,t) = 0$ if $d\tilde{N}_e(x,t) > 0$.

- HBJVI for $\tilde{V}(x,t)$: Assume that the planner choose optimally to $dN_e(x,t) > 0$ for some x and t. Then $\nu_e(x,t) = \nu_N(x,t) = 0$ and we have

$$-(\rho+\chi)\tilde{V}(x,t) + \partial_t \tilde{V}(x,t) + \tilde{w}_m(t)h(x) + \tilde{v}_m(x,t) = 0 \tag{81}$$

$$\tilde{\varphi}(t)x\tilde{V}_x(x,t) - \tilde{w}_m(t)h(x) + x\tilde{w}_e(t) - \tilde{v}_m(x,t) = 0 \tag{82}$$

or

$$(\rho+\chi)\tilde{V}(x,t) = x\tilde{w}_e(t) + \tilde{\varphi}(t)x\tilde{V}_x(x,t) + \partial_t\tilde{V}(x,t) \tag{83}$$

with

$$-(\rho+\chi)\tilde{V}(x,t) + \partial_t\tilde{V}(x,t) + \tilde{w}_m(t)h(x) = -\tilde{v}_m(x,t) \leq 0. \tag{84}$$

If $d\tilde{N}_e(x,t) = 0$ and $d\tilde{N}(x,t) - d\tilde{N}_e(x,t) > 0$ so $\tilde{v}_m(x,t) = \tilde{v}_N(x,t) = 0$ and we have

$$(\rho+\chi)\tilde{V}(x,t) = \tilde{w}_m(t)h(x) + \partial_t\tilde{V}(x,t) \tag{85}$$

and

$$\tilde{\varphi}(t)x\tilde{V}_x(x,t) - \tilde{w}_m h(x) + \tilde{w}_e(t)x = -\tilde{v}_e(x,t) \leq 0$$

or

$$-(\rho+\chi)\tilde{V}(x,t) + \tilde{\varphi}(t)x\tilde{V}_x(x,t) + \partial_t\tilde{V}(x,t) + \tilde{w}_e(t)x = -\tilde{v}_e(x,t) \leq 0$$

Notice that we can write it as a standard HBJ variation inequality

$$(\rho+\chi)\tilde{V}(x,t) = \max\left\{\tilde{w}_m(t)h(x), \tilde{\varphi}(t)x\tilde{V}_x(x,t) + \tilde{w}_e(t)x\right\} + \partial_t\tilde{V}(x,t) \tag{86}$$

- FOCs for $\tilde{\varphi}(t)$: The Lagrangean for $\varphi(t)$ is

$$\mathscr{L} = \cdots + \int_0^\infty e^{-\rho t}\left[\int_0^\infty x\tilde{v}(x,t)\partial_x\tilde{N}_e(x,t)dx - \frac{\tilde{v}_\varphi(t)}{\eta}\left(\frac{\tilde{w}_e(t)}{1-\eta^{-1}}\right)^{1-\eta}\right]\tilde{\varphi}(t)dt$$

Taking the FOC for $\varphi(t)$ gives

$$\frac{\tilde{v}_\varphi(t)}{\eta}\left(\frac{\tilde{w}_e(t)}{1-\eta^{-1}}\right)^{1-\eta} = \int_0^\infty x\tilde{v}(x,t)\partial_x\tilde{N}_e(x,t)dx \tag{87}$$

and given that $\partial_x\tilde{V}(x,t) = \tilde{v}(x,t)$, we can rewrite this as

$$\frac{\tilde{v}_\varphi(t)}{\eta}\left(\frac{\tilde{w}_e(t)}{1-\eta^{-1}}\right)^{1-\eta} = \int_1^\infty x\tilde{V}_x(x,t)d\tilde{N}_e(x,t) \tag{88}$$

Observe that since $(1-\tilde{\kappa}(t)) = \left(1+\delta\tilde{v}_\varphi(t)\right)^{\frac{1}{1-\eta}}$, we have that $\frac{(1-\tilde{\kappa}(t))^{1-\eta}-1}{\delta} = \tilde{v}_\varphi(t)$ and using this result

$$\frac{(\tilde{w}_e(t)(1-\tilde{\kappa}(t)))^{1-\eta} - \tilde{w}_e(t)^{1-\eta}}{\delta} = \eta(1-\eta^{-1})^{1-\eta}\int_1^\infty x\tilde{V}_x(x,t)d\tilde{N}_e(x,t) \tag{89}$$

■

Proposition 10. *Let $T_{AI}(t)$ be a tax on the firms profits for AI and $T_F(t)$ be a subsidy on the firms' task frontier. For simplcity, assume that there is first best such that $\tilde{\tau} \in (0, \tilde{f}(t))$ for all t. The firms* *profit function in the competitive equilibrium is given by*

$$\max_{\mathscr{A}_f}\Bigg\{\int_0^{f(t)}\max\Big\{x(\tau,t)^{1-\eta^{-1}}-w_e(t)x(\tau,t),(1-T_{AI}(t))\underbrace{\Big((z(\tau)a(\tau,t))^{1-\eta^{-1}}-qa(\tau,t)\Big)}_{\text{profits from AI}}\Big\}d\tau \tag{90}$$

$$-(1-T_F(t))\underbrace{\Big(w_m(t)\int h(x)dN_m(x,t)+qa_m(t)\Big)}_{\text{subsidy frontier}}\Bigg\}$$

$$s.t.\quad x(\tau,t)\equiv\int x dN_e(x,t;\tau)$$

$$s(t)=\int h(x)dN_m(x,t),\quad f(t)=\bar{\tau}\left[\psi^{\omega^{-1}}a_m(t)^{1-\omega^{-1}}+(1-\psi)^{\omega^{-1}}s(t)^{1-\omega^{-1}}\right]^{\frac{1}{1-\omega^{-1}}}$$

Then if

$$T_F(t)=T_{AI}(t)=1-(1-\tilde{\kappa}(t))^{\eta-1} \tag{91}$$

there exist a competitive equilibrium that implements the first best allocation described by the planner's solution.

Proof. The proof consists in showing that the optimal choice of the firm in (90) coincides with the planner's solution. Let us guess and verify that $w_e=\tilde{w}_e$ and $w_m=\tilde{w}_m$. Here $y(\tau)$ denotes $z(\tau)a(\tau,t)$ for AI tasks ($\tau<\tilde{\tau}^*$) and $x(\tau,t)$ for labor tasks ($\tau>\tilde{\tau}^*$). Since the tax $(1-T_{AI})$ multiplies profits but not the FOC for $a(\tau,t)$, it cancels and the intensive-margin choices are undistorted. The firm's optimal choice is given by

$$y(\tau)=\begin{cases}\left(\frac{qe^{\varepsilon\tau}}{1-\eta^{-1}}\right)^{-\eta}, & \tau<\tilde{\tau}^*\\ \left(\frac{\tilde{w}_e}{1-\eta^{-1}}\right)^{-\eta}, & f\geq\tau>\tilde{\tau}^*\end{cases} \tag{92}$$

$$\tag{93}$$

Profits are given by

$$\tilde{\pi}(\tau,t)=\begin{cases}(1-T_{AI}(t))\eta^{-1}\left(\frac{qe^{\varepsilon\tau}}{1-\eta^{-1}}\right)^{1-\eta}, & \tau<\tilde{\tau}^*\\ \eta^{-1}\left(\frac{\tilde{w}_e}{1-\eta^{-1}}\right)^{1-\eta}, & f\geq\tau>\tilde{\tau}^*\end{cases} \tag{94}$$

So, the the optimality condition for τ^*

$$(1-T_{AI}(t))\left(\frac{qe^{\varepsilon\tau^*}}{1-\eta^{-1}}\right)^{1-\eta}=\left(\frac{\tilde{w}_e}{1-\eta^{-1}}\right)^{1-\eta}\iff \tag{95}$$

$$qe^{\varepsilon\tau^*}=\tilde{w}_e(1-T_{AI}(t))^{\frac{1}{\eta-1}}=\tilde{w}_e(1-\tilde{\kappa}(t)) \tag{96}$$

$$\tag{97}$$

Therefore, $y(\tau,t)=\tilde{y}(\tau,t)$ and $\tau^*=\tilde{\tau}^*$. Notice that the optimal choice of the firm for $a_m(t)$ and

$s(t)$ is given by

$$p_f(t) = \frac{(1 - T_F(t))}{\bar{\tau}} \left[\psi\, q^{1-\omega} + (1-\psi)\, \tilde{w}_m(t)^{1-\omega}\right]^{\frac{1}{1-\omega}} \tag{98}$$

and since, $\tau^* = \tilde{\tau}^* < \tilde{f}^*$,

$$\frac{(1 - T_F(t))}{\bar{\tau}} \left[\psi\, q^{1-\omega} + (1-\psi)\, \tilde{w}_m(t)^{1-\omega}\right]^{\frac{1}{1-\omega}} = \eta^{-1}\left(\frac{\tilde{w}_e(t)}{1-\eta^{-1}}\right)^{1-\eta} \iff \tag{99}$$

$$\frac{1}{\bar{\tau}} \left[\psi\, q^{1-\omega} + (1-\psi)\, \tilde{w}_m(t)^{1-\omega}\right]^{\frac{1}{1-\omega}} = \eta^{-1}\left(\frac{((1 - T_{R\&D}(t)))^{\frac{1}{(\eta-1)}}\, \tilde{w}_e(t)}{1-\eta^{-1}}\right)^{1-\eta} \iff \tag{100}$$

$$\frac{1}{\bar{\tau}} \left[\psi\, q^{1-\omega} + (1-\psi)\, \tilde{w}_m(t)^{1-\omega}\right]^{\frac{1}{1-\omega}} = \eta^{-1}\left(\frac{\tilde{w}_e(t)(1-\tilde{\kappa}(t))}{1-\eta^{-1}}\right)^{1-\eta} \tag{101}$$

Since the right hand side of the previous equation is equal to $\tilde{p}_f$, the cost-minimizing inputs that achieve unit cost $\tilde{p}_f$ at factor prices $(q, \tilde{w}_m)$ are uniquely pinned down by the CES demand system, so $a_m(t) = \tilde{a}_m(t)$ and $s(t) = \tilde{s}(t)$. Since $f(t) = \bar{\tau}[\psi^{\omega^{-1}} a_m^{1-\omega^{-1}} + (1-\psi)^{\omega^{-1}} s^{1-\omega^{-1}}]^{1/(1-\omega^{-1})}$ by the frontier constraint, it follows immediately that $f(t) = \tilde{f}(t)$. Therefore, all firm choices coincide with the planner's solution, so the planner's allocation is implementable. To complete the verification of the guess $w_e = \tilde{w}_e$ and $w_m = \tilde{w}_m$: since the firm demands $\tilde{a}_m(t)$, $\tilde{s}(t)$, and $\tilde{y}(\tau, t)$, market clearing for labor and AI capital is satisfied at the planner's wages by Lemma 4, which characterizes the workers' optimal stopping rule under the same wages. Thus, there is a competitive equilibrium that implements the planner's solution. Since the planner's solution satisfies feasibility, the competitive equilibrium is efficient. ■