Artificial Intelligence Investments and Expertise Erosion*

Pierre Chaigneau† Evan Jo‡

Queen's University Queen's University

First version: September 2025

This version: August 2026

Abstract

We study the feedback between automation and expertise accumulation. Building on Autor and Thompson's (2025) "expertise framework", we develop an overlapping-generations model with endogenous on-the-job learning. Automation reduces learning in entry-level tasks, slowing expertise acquisition and worsening the economy's expertise distribution. Yet we show that pre-existing automation decreases firms' willingness to pay for improvements in automation: such improvements benefit occupations requiring high human expertise, while the supply of high-expertise workers declines with prior automation. This weakens the dynamic complementarity between existing automation and the subsequent adoption of expertise-substituting technologies.

Keywords: automation, expertise, human capital, investments, learning.

JEL classification: J24, O33, D62, L13.

*This research received no external funding. The authors have no financial relationships or other potential conflicts of interest to disclose.

†Smith School of Business, Goodes Hall, 143 Union Street West, K7L 2P3, Kingston, ON, Canada. Email: pierre.chaigneau@queensu.ca.

‡Smith School of Business, Goodes Hall, 143 Union Street West, K7L 2P3, Kingston, ON, Canada. Email: evan.jo@queensu.ca.

1 Introduction

Recent advances in Artificial Intelligence (AI) and automation have renewed concerns about the long-run consequences of technological change for the workforce. A growing literature studies how automation substitutes for human labor by performing tasks in each occupation (Berger et al. 2024; Brynjolfsson, Chandar, and Chen 2025; Hosseini Maasoum and Lichtinger 2025; Hampole et al. 2025).[1] However, automation may also affect the process through which workers accumulate expertise, with important implications for the future supply of skilled labor and the incentives to develop new technologies. This paper studies the dynamic interaction between automation and human expertise.

To understand how automation affects expertise accumulation, it is important to consider how jobs are organized around tasks. Professional services workers, such as lawyers, consultants, or financial advisors, perform a wide range of tasks, so that each occupation can be viewed as a bundle of tasks rather than a single activity (Autor and Handel 2013; Brynjolfsson, Mitchell, and Rock 2018). Performing these tasks requires broad expertise, which workers acquire partly through formal education but also through on-the-job learning, including both structured training and informal experience (Loewenstein and Spletzer 1994).

Autor and Thompson (2025) formalize these features in an "expertise framework" based on a model of occupational task bundling. Workers are vertically differentiated by expertise ("hierarchical expertise") and occupations involve multiple tasks ("indivisibility"). Technological progress expands the set of tasks that can be automated, but the extent of automation differs across occupations: in some occupations human workers remain necessary for certain expert tasks (occupations above the automation threshold), whereas in others all expert tasks can be automated (occupations below the threshold).

In Autor and Thompson (2025), workers' expertise is exogenous. To better understand the effect of automation on expertise generation, we extend this framework to allow for endogenous expertise acquisition with on-the-job learning in an overlapping generations (OLG) model with young and old workers. Specifically, workers in occupations above the automation threshold accumulate expertise according to a learning process that first-order stochastically dominates that of workers below the threshold. Intuitively, workers below the threshold perform fewer expert tasks and thus face fewer learning opportunities, making large improvements in expertise less likely. This is consistent with evidence that the automation

[1] Anecdotal evidence points to the same pattern in specific occupations: "AI tools are doing the types of simple coding and debugging tasks that junior software developers did to gain experience. AI is also doing work that young employees in the legal and retail sectors once did. And Wall Street firms are reportedly considering steep cuts to entry-level hiring." (*Fortune*, May 2025)

of cognitive tasks induces "cognitive offloading" that reduces learning and cognitive performance (Grinschgl and Neubauer 2022; Zhai, Wibowo, and Li 2024; Kosmyna et al. 2025).

Our first main result shows analytically that a firm's willingness to pay for a marginal improvement in automation is decreasing in the existing level of automation.[2] This result contrasts with the presumption that automation should become more valuable as human expertise deteriorates. Intuitively, since firms take wages as given, the value of marginal automation depends on how it changes the productivity of workers across occupations. Higher past automation reduces the number of workers with expertise above the automation threshold, whose productivity would increase with further automation. By contrast, additional automation does not raise the productivity of workers below the threshold, because all the tasks requiring their expertise have already been automated. The scarcity of high-expertise workers also raises the relative price of output produced by occupations requiring higher expertise, which increases the value of automation; because outputs associated with different occupations are substitutes, this price effect is dominated by the decline in the number of workers augmented by automation.

This result has implications for the legitimate concern that the widespread deployment of powerful AI might erode the expertise distribution to such an extent that it would require additional advances in AI, that is, that advances in AI would be self-reinforcing (Jones 2024, 2026; Acemoglu and Lensman 2024). On the contrary, we show that reaching a higher level of automation does not increase the value of further advances in automation.

This result relies on our assumption that automation destroys learning opportunities: workers in occupations whose expert tasks are all automated are less likely to reach any given range of higher expertise via learning. If automation instead redirected learning, so that workers in these occupations learned with the same probability but reached lower expertise levels, the density of workers could increase over part of the range above the automation threshold. In that case, the willingness to pay for automation could increase with prior automation, and advances in AI would be self-reinforcing. Within the task-bundling framework, destruction is the natural case: occupations whose expert tasks are all automated retain only generic tasks, and thus offer no expert-task exposure toward which learning could be redirected.[3]

Next, we study the dynamic interaction between investment in automation and the evolving expertise distribution. In every period, the level of automation is chosen optimally by

[2] This result takes into account the effect of the expertise distribution on prices and wages, i.e., the willingness to pay is evaluated in general equilibrium.

[3] Machine assistance may also reduce learning in occupations where expert tasks are still performed by workers, as the evidence on cognitive offloading suggests; our main result is robust to this additional form of impairment.

an incumbent AI producer that faces a displacement risk measured by the number of alternative AI producers. Automation depreciates over time, and the depreciated level is freely available to firms. Firms take wages and output prices as given, so the AI producer's revenue is the amount that firms are willing to pay for incremental automation, which depends on the current expertise distribution and therefore on past automation.

We show that no market structure eliminates the erosion externality. Expertise erosion lowers output at every level of automation, but the AI producer derives revenues only from the incremental automation that it sells. Thus, regardless of the market structure, it internalizes erosion only through the decline in the willingness to pay established by our first result, and only over incremental automation, whereas the social planner internalizes the full reduction in future output.

This asymmetry raises the equilibrium level of automation above the first-best level, but there are two countervailing effects. First, the AI producer cannot charge for the part of the increase in output that raises workers' income and profits in already automated occupations (automating an occupation releases its expert workers into the pool of generic labor). Second, it values its inherited level of automation less than the planner does: that level lowers the cost of future investment only for a producer that retains the market, and it raises the freely available baseline level of automation, which reduces its future revenue. The second effect weakens as the depreciation rate rises and vanishes at full depreciation.

In numerical computations, the asymmetry in the internalization of erosion dominates when the adverse effect of automation on learning is sufficiently severe, automation depreciates a lot over one period (about twenty years in our OLG model), and the planner discounts the future at a low rate. In this case, the equilibrium level of automation exceeds the first-best level, otherwise the two countervailing effects dominate. The direction of the inefficiency determines the optimal corrective policy: a tax on the AI producer's investment in the former case, and a subsidy otherwise (neither competition policy nor intellectual property regulation substitutes for this instrument). The two regimes are close at the depreciation rates that we argue are empirically relevant, so the direction of the inefficiency is ultimately an empirical question, which the model reduces to two measurable quantities: the durability of automation and the severity of the learning losses.

Most economic models of technological change predict a positive dynamic complementarity between automation and human capital. In learning-by-doing frameworks, performing tasks raises future productivity (Arrow 1962; Lucas 1988). In models of directed technical change, a declining proportion of skilled workers raises the profitability of skill-substituting innovations, which makes automation self-reinforcing (Acemoglu 1998, 2002). By contrast, we identify a negative dynamic feedback between automation and expertise. The automation of entry-level expert tasks that generate on-the-job learning worsens the future distribution of expertise. While automation augments high-expertise workers, their endogenous scarcity reduces the private value of further advances in automation. The condition separating the two predictions is the elasticity of substitution between outputs associated with different occupations: when it is higher than one, the market-size effect of a thinner pool of high-expertise workers dominates the price effect of their scarcity.

Our model also has implications for the distribution of income. Workers whose expertise when young is above the automation threshold still have learning opportunities, whereas workers below the threshold do not: they become interchangeable and earn the same wage. Automation raises the pay of the former because it augments their productivity and because it increases the scarcity of their high-level expertise when they are old. Thus, it shifts income toward the highest-expertise workers.[4]

Related literature. In standard task models, workers specialize in tasks according to their comparative advantage (Autor, Levy, and Murnane 2003; Acemoglu and Autor 2011; Acemoglu and Restrepo 2019; Acemoglu et al. 2024), and a given task is performed by a single skill group. By contrast, in the model of occupational task bundling on which we rely, a worker offers a set of skills (Heckman and Scheinkman 1987) and a given task is performed in multiple occupations by different workers (Autor and Thompson 2025). Agrawal, Gans, and Goldfarb (2019) also emphasize that the automation of some tasks raises the value of the human inputs that remain, and argue that automating prediction raises the value of human judgment.

Empirically, the effects of AI are heterogeneous: Berger et al. (2024) find that generative AI complements high-level white-collar jobs but substitutes for low-level ones, consistent with our framework, and Chen and Wang (2024) document that the workforce effects of AI innovation combine displacement and augmentation. At the firm level, AI adoption is associated with growth and product innovation (Babina et al. 2024) and with changes in firm value that track workforce exposure (Eisfeldt, Schubert, and Zhang 2023).

Ide (2025), Ide and Talamàs (2025), Garicano and Rayo (2025), and Acemoglu, Kong, and Ozdaglar (2026) study related learning and knowledge-collapse channels. These papers study the erosion of learning and knowledge transmission itself and its consequences for output, organization, and welfare, taking the path of automation as given. Our contribution is the reverse feedback: how erosion changes the private incentive to automate further, and whether the market structure of AI supply mitigates the resulting externality. We also characterize

[4]The effect of automation on wages within a period is as in Autor and Thompson (2025), where the distribution of expertise is exogenous. We focus on a different channel: automation makes high-level expertise endogenously scarcer, which raises the relative pay of the workers who retain expert tasks.

the equilibrium willingness to pay for automation and the AI producer's investment problem, and we show that no market structure eliminates the externality.

In a closely related paper, Ide (2025) develops a task-based overlapping-generations model in which novices acquire tacit knowledge by working alongside experts, and shows that improvements in entry-level automation raise output on adoption but can lower growth and welfare by reallocating novices away from the most knowledgeable experts. The mechanism is thus the composition of novice-expert matches and the transmission of best practices, and the level of automation is exogenous. Instead, we characterize the price of automation in equilibrium, and the investment decision of the AI producer.

In a model where the "masters" choose between apprentices and machines, Garicano and Rayo (2025) study the viability of the apprenticeship contract within the firm with an exogenous automation technology. Baccara, Lee, and Yariv (2023) study the allocation of tasks when assignments generate on-the-job training, which is the trade-off that our threshold creates across occupations, with the level of automation again taken as given. Acemoglu, Kong, and Ozdaglar (2026) study a learning externality in the accumulation of public general knowledge and derive information-design remedies, whereas our externality arises in the market for automation and is corrected by a Pigouvian instrument on the producer's investment.

In our model, workers rather than firms possess expertise. The resulting appropriability problem has implications for the value of organization capital (Eisfeldt and Papanikolaou 2013), and it connects our analysis to the literature on labor and corporate finance (Nishesh, Ouimet, and Simintzi 2024; Maug 2025).

2 The model

We build on the occupational task bundling model of Autor and Thompson (2025). Since our focus is on professional occupations which do not require much capital, such as coding, consulting, legal services, and financial advice, we abstract from capital allocation.

2.1 Occupations, tasks, and automation

Each worker i has expertise e_i. This expertise is transferable, i.e., it is not firm-specific. Each worker is assigned to some "occupation". An occupation involves a bundle of tasks that each require a different level of expertise. A model of task bundling captures important features of professional firms, in which even the most expert workers must also perform some tasks that only require a basic level of expertise, as well as more generic tasks such as meeting clients or presenting findings.[5]

Occupation $\phi \in [0,1]$ involves a bundle of tasks t in the continuum $[-g(\phi), \phi]$, where $g(\phi) \equiv \theta(1-\phi)$. The parameter θ controls the range of generic tasks. The least expert occupation ($\phi = 0$) has generic-task range θ, while the most expert occupation ($\phi = 1$) has no generic tasks: $g(1) = 0$. Following Autor and Thompson (2025), we assume $\theta > 1$ (this implies that the productivity of workers whose expert tasks are partly automated is higher than the productivity of workers who only perform generic tasks).

In occupation ϕ, tasks $t \in [-g(\phi), 0]$ are generic tasks that require labor input but no expertise, whereas tasks $t \in (0, \phi]$ are expert tasks. Generic tasks include managing a team and interacting with computer programs. They are therefore not necessarily menial, and many involve general managerial skills. An expert task only requires labor input if it is not automated. A non-automated expert task t can only be completed by a worker with expertise $e_i \geq t$.

Given a state of automation $I \in [0,1]$, all expert tasks $t \in (0, I]$ can be automated. More basic expert tasks are automated first, and those requiring the highest level of expertise are automated last. We assume that any task that can be automated will be automated. Thus there are three types of tasks: generic tasks $t \in [-g(\phi), 0]$ (require labor, no minimum expertise); automated tasks $t \in (0, I]$ (no labor); and expert tasks $t \in (I, \phi]$ (require labor and minimum expertise $e_i \geq t$). Figure 1 illustrates.

[5] Examples of occupations include product manager, engineer, financial advisor, and lawyer. All these occupations involve many different tasks. Even though some tasks can be delegated or automated, a single worker still needs to perform many different tasks. For example, a financial advisor needs to understand the circumstances of a client, to either generate expert advice or make sense of automated advice, and to be available to meet with the client.

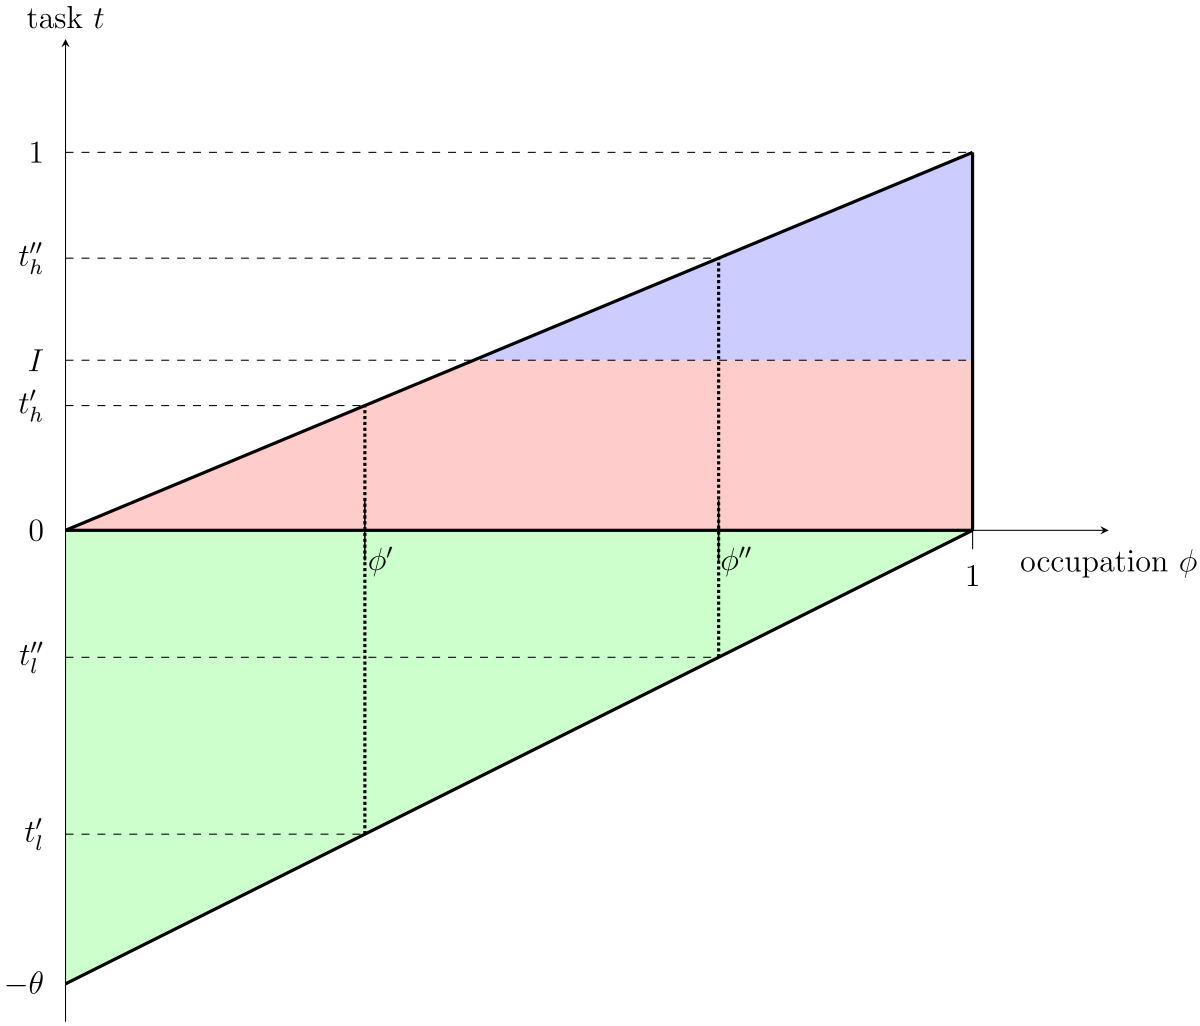

Figure 1: Expertise level required in each task t (y-axis), for all occupations ϕ (x-axis). In occupation ϕ', tasks range from t_l' to t_h'; in occupation ϕ'', tasks range from t_l'' to t_h''. Green area: generic tasks. Red area: automated expert tasks. Blue area: expert tasks requiring human expertise.

The tasks involved in an occupation are perfect complements: all tasks must be completed for output to be produced. Output is thus constrained by the tasks that are not automated, as in the weak-link framework of Jones (2026). A worker divides her unit of labor across the non-automated tasks of her occupation, allocating $\ell_i(t) \geq 0$ to task t; automated tasks are completed at no labor cost, and output per task equals the labor allocated to it. The contribution $y_i(\phi)$ of worker i in occupation ϕ is the measure of tasks in the occupation times the minimum task-level output,

$$y_i(\phi) = (\phi + g(\phi)) \min_t \ \ell_i(t), \tag{1}$$

where the minimum is taken over non-automated tasks, so the worker optimally allocates her unit of labor uniformly across non-automated tasks. With $L(\phi)$ units of labor allocated to occupation ϕ, the output of this occupation is

$$Y(\phi) = \left(L(\phi) y_i(\phi)\right)^{\frac{\eta}{\eta+1}}, \tag{2}$$

where $\eta > 0$ and $\frac{\eta}{\eta+1} < 1$ captures decreasing returns to scale in each occupation.

The economy is populated by a continuum $[0, 1]$ of firms. In each firm there is a continuum $\phi \in [0, 1]$ of occupations. Firms thus combine challenging work that requires the highest level of expertise and less challenging work that requires only limited expertise. For example, in a law firm, some contracts are standard and involve boilerplate language (low ϕ), whereas others are highly complex and involve rare legal issues (high ϕ). Likewise, a financial intermediary can offer different services and serve different clienteles that require different levels of expertise. Final output is

$$Y = \left(\int_0^1 Y(\phi)^{\frac{\lambda-1}{\lambda}} d\phi\right)^{\frac{\lambda}{\lambda-1}}, \tag{3}$$

where $\lambda > 1$ is the elasticity of substitution between occupation-level outputs.[6]

2.2 Overlapping generations and learning

We extend the framework to an OLG model in which each cohort lives for two periods: an individual is young in the first period and old in the second. Each cohort has unit measure. Firms compete for workers by offering spot contracts.[7] The expertise of young workers is exogenous (its density f_y does not depend on automation).[8] The expertise of old workers is endogenous: it depends on their learning opportunities when young, which Assumption 1 relates to the level of automation I_0 prevailing at that time.

[6] Autor and Thompson (2025) make the same assumption because, with an elasticity below one, demand for an occupation's output would be inelastic: a productivity improvement would lower the price of that output proportionally more than it would raise the quantity, so the occupation's revenue and relative employment would fall. $\lambda = 1$ is also the threshold above which the market-size effect dominates the price effect in the theory of directed technical change (Acemoglu 2002).

[7] Long-term contracts would not eliminate the erosion externality. Under the usual assumption of one-sided commitment, a worker is free to leave a firm, so her second-period wage must match the market value of her expertise, and the firm cannot appropriate the return to the learning that it provides. This return can only be bought by the worker herself via a lower first-period wage, and what a young worker can forgo is bounded by limited liability and subsistence constraints.

[8] Expertise erosion therefore does not cumulate across cohorts: automation lowers the expertise of the affected cohort when it is old, whereas the next cohort of young workers is unaffected. Allowing erosion to cumulate across cohorts is left for future research.

Assumption 1. For a young worker with expertise e_y:

- If $e_y > I_0$: expertise when old is drawn from $G_a(\cdot|e_y)$ with support contained in $[e_y, 1]$.
- If $e_y \leq I_0$: expertise when old is drawn from $G_b(\cdot|e_y)$ with support contained in $[e_y, 1]$.

Both distributions have a (possibly zero) density on $(e_y, 1]$, denoted respectively $g_a(\cdot|e_y)$ and $g_b(\cdot|e_y)$, and an atom at e_y. For every $e_y \in [0, 1)$,

$$g_b(x|e_y) \;\leq\; g_a(x|e_y) \quad \text{for almost every } x \in (e_y, 1], \tag{4}$$

with $\int_c^1 g_b(x|e_y)dx < \int_c^1 g_a(x|e_y)dx$ for every $c \in [e_y, 1)$.

Assumption 1 requires that automation censor learning: below the automation threshold, the density of workers reaching any expertise level above their initial level is weakly lower, and the atom at e_y is correspondingly larger. Equivalently, a worker below the threshold draws the same potential expertise gain as a worker above the threshold, but with positive probability that gain is not realized and expertise remains at e_y. Expertise is acquired by performing expert tasks, and on-the-job human capital is specific to the tasks performed (Gibbons and Waldman 2004), so workers in occupations where expert tasks are all automated accumulate less of it. This condition implies that $G_a(\cdot|e_y)$ first-order stochastically dominates $G_b(\cdot|e_y)$, but it is stronger. The discussion following Proposition 1 explains why the distinction matters for our results.

Three strands of evidence support this condition. First, on-the-job human capital is acquired by performing tasks and is specific to the tasks performed (Loewenstein and Spletzer 1994; Gibbons and Waldman 2004), so a worker whose expert tasks are performed by machines does not accumulate the expertise that performing them would have produced. Second, automation removes the assignments through which that accumulation occurs, and it removes them at the bottom of the expertise distribution.[9] Third, the delegation of thinking tasks to technology is a recognized risk to skill acquisition (Grinschgl and Neubauer 2022), and participants who wrote with an AI assistant showed weaker neural engagement and poorer recall of their own work than participants who wrote without one (Kosmyna et al. 2025). The censoring form of the condition follows from the same mechanism: learning requires performing an expert task, which either occurs or does not, so automation removes the opportunity rather than reducing the amount learned when the opportunity arises.

[9] Employment of workers aged 22 to 25 declines relative to that of older workers in the occupations most exposed to generative AI, and the decline is concentrated in the occupations where AI automates rather than augments (Brynjolfsson, Chandar, and Chen 2025). Generative AI also substitutes for low-level white-collar work while complementing high-level work (Berger et al. 2024).

3 Analysis

3.1 Productivity multiplier

Consider a given expertise distribution and a given level of automation $I \in [0,1]$. Define

$$\varphi(\phi; I) \equiv \frac{\phi + g(\phi)}{\max\{\phi - I, 0\} + g(\phi)}. \tag{5}$$

The function $\varphi(\phi; I)$ measures worker productivity in occupation ϕ. With $I = 0$, $\varphi(\phi; 0) = 1$ for all ϕ. With $I > 0$, automation increases the productivity of all workers in occupations $\phi > 0$. Since $\theta > 1$, the increase is larger for occupations requiring more expertise: $\frac{\partial \varphi}{\partial \phi} > 0$ when $I > 0$. Since $g(1) = 0$, the multiplier diverges as $I \to 1$: $\varphi(1; I) = \frac{1}{1-I} \to \infty$.

3.2 Equilibrium allocation

Lemma 1. *Output in occupation ϕ is $Y(\phi) = (L(\phi)\varphi(\phi; I))^{\frac{\eta}{\eta+1}}$.*

All proofs are in Appendix A.

With the final good as numéraire, the price of occupation ϕ output implied by equation (3) is $P_y(\phi) = \left(\frac{Y}{Y(\phi)}\right)^{\frac{1}{\lambda}}$. Let f denote the density of expertise in the population of workers, young and old, and F the corresponding distribution: $F(x)$ is the mass of workers with expertise at most x, and the total mass is 2.

Lemma 2. *A worker with expertise $e_i \leq I$ is assigned to an occupation $\phi \leq I$, with*

$$L(\phi) = F(I) \frac{\varphi(\phi; I)^{\frac{\eta(\lambda-1)}{\lambda+\eta}}}{\int_0^I \varphi(z; I)^{\frac{\eta(\lambda-1)}{\lambda+\eta}} dz},$$

and all such workers receive the common wage $\bar{w}$. Provided that f is non-increasing, a worker with expertise $e_i > I$ is assigned to occupation $\phi = e_i$, with $L(\phi) = f(\phi)$ and wage

$$w(\phi) = \frac{\eta}{\eta+1} Y^{\frac{1}{\lambda}} \left(\frac{\varphi(\phi; I)^{\frac{\eta(\lambda-1)}{\lambda+\eta}}}{f(\phi)} \right)^{\frac{\lambda+\eta}{\lambda(\eta+1)}}.$$

The assignment $\phi = e_i$ is incentive-compatible: a worker cannot be assigned to an occupation whose expert tasks exceed her expertise, and she prefers the highest-wage occupation available, which is $\phi = e_i$ whenever $w(\phi)$ is increasing in ϕ (this holds whenever f is non-increasing, as in the benchmark distribution of Appendix C). Under the learning process of Appendix C, the condition holds for every level of inherited automation: the differential equation displayed in that Appendix makes the population density constant above the inherited threshold, strictly decreasing below it, and downward-jumping at the threshold, so the assignment is incentive-compatible.

Lemma 3. *In any occupation ϕ, the labor share is $\frac{\eta}{\eta+1}$ and the profit share is $\frac{1}{\eta+1}$.*

The wage schedule already shows the redistributive consequences of automation. The within-occupation labor share is always $\frac{\eta}{\eta+1}$, but workers with $e_i > I$ are in occupations where automation augments productivity, and their wages are increasing in ϕ and in I, whereas workers with $e_i \leq I$ are perfect substitutes in fully automated occupations and receive the common wage $\bar{w}$. Automation therefore shifts income toward the highest-expertise workers.

Firms are atomistic and take output prices $P_y(\phi)$ and the wage schedule $w(\phi)$ as given; because outputs associated with different occupations enter equation (3) separably, each firm's problem decomposes into an independent problem for each occupation. For a skilled occupation $\phi > I$, labor input is determined by the inelastic supply of workers of expertise ϕ, $L(\phi) = f(\phi)$ (Lemma 2). The occupation-level profit is

$$\pi(\phi) = P_y(\phi)Y(\phi) - w(\phi)L(\phi). \tag{6}$$

A marginal advance in the automation used by the firm raises its output via the effect on the productivity multiplier $\varphi(\phi; I)$ (it does not change the firm's labor input or the wage it pays). Holding $L(\phi)$, $w(\phi)$, and $P_y(\phi)$ fixed,

$$\frac{\partial \pi(\phi)}{\partial I} = P_y(\phi)\frac{\partial Y(\phi)}{\partial I} - \underbrace{\frac{\partial \left[w(\phi)L(\phi)\right]}{\partial I}}_{=0} = \frac{\partial \left[P_y(\phi)Y(\phi)\right]}{\partial I}. \tag{7}$$

At the margin, the wage cost is fixed, so the firm's marginal profit from automation is simply its marginal revenue from automation, $\frac{\partial (P_y Y)}{\partial I}$, which is also its willingness to pay for a marginal increase in automation.

The income-share identity of Lemma 3, $w(\phi)L(\phi) = \frac{\eta}{\eta+1}P_y(\phi)Y(\phi)$, gives the division of the revenue level once wages have cleared; it does not reduce the firm's marginal willingness to pay. What determines willingness to pay is the marginal wage cost, which is zero: the firm prices automation taking its wage and its labor input as given, since labor supply at each expertise level is inelastic, so the firm earns the entire marginal gain in output as profit. We henceforth refer to $\frac{\partial (P_y Y)}{\partial I}$ as the firm's marginal willingness to pay for automation.

3.3 Marginal willingness to pay and prior automation

Proposition 1 describes how the marginal willingness to pay for automation depends on the level of automation previously reached.

Proposition 1. *Suppose that Assumption 1 holds, that $f_y > 0$ on $[0, 1)$, and that f is non-increasing. Then the equilibrium marginal willingness to pay for an improvement in automation is non-increasing in the level I_0 of automation reached last period at each occupation $\phi > I$, and strictly decreasing in the aggregate, for any $I_0 \leq I < 1$.*

The higher the pre-existing level of automation, the lower the value of additional improvements in automation. A marginal increase in automation does not affect occupations below the automation threshold, whose expert tasks are already automated; it raises the productivity of workers above the threshold, and the measure of these workers is decreasing in prior automation. Proposition 1 establishes that the decline in the value of automation holds in general equilibrium, taking into account adjustments of output prices and wages.

The proof proceeds in three steps. First, under Assumption 1, a higher I_0 lowers the population density of expertise above the threshold and worsens the expertise distribution in a first-order stochastic dominance sense. Second, aggregate output falls: the competitive allocation maximizes output given the expertise distribution, and any assignment feasible under a stochastically worse distribution can be replicated under the better one. Third, equilibrium occupation revenue is $Y^{\frac{1}{\lambda}} \left(f(\phi)\varphi(\phi; I)\right)^{\frac{\eta(\lambda-1)}{\lambda(\eta+1)}}$, and with $\lambda > 1$ both factors are decreasing in I_0, so the willingness to pay falls at every occupation above the automation threshold and remains zero below it.

The economic intuition is the balance between a price effect and a market-size effect. Scarcer high-expertise workers reduce the output of expert occupations and raise their relative prices; the higher prices raise the value of automating expert tasks, while the thinner pool of workers augmented by automation lowers it. With $\lambda > 1$, demand for outputs associated with different occupations is elastic: quantities fall proportionally more than prices rise, so revenue falls (the market-size effect dominates), and the decline in aggregate output reinforces the result by scaling down all revenues. With $\lambda < 1$, the price effect would dominate and the willingness to pay could rise with prior automation, which is the self-reinforcing configuration of the directed-technical-change literature (Acemoglu 1998, 2002). Elastic substitution across occupations is thus the condition under which expertise erosion is self-limiting rather than self-reinforcing. This condition is testable: the feedback from erosion to automation incentives should be self-limiting in sectors where outputs associated with different occupations are readily substitutable, and can be self-reinforcing in sectors where outputs associated with different occupations are strong complements in final production.

First-order stochastic dominance of G_a over G_b would not suffice for Proposition 1: it constrains distribution functions rather than densities, and is consistent with the density of old workers rising on part of the skilled range, in which case the willingness to pay can increase with I_0. Assumption 1 rules this out by requiring that automation destroy learning events rather than redirect them: under learning redirection a worker would still learn but would reach a lower level of expertise, whereas under learning destruction she would not learn with positive probability and would remain at her initial expertise. The distinction is economic rather than technical. If automation instead redirects learning, keeping its frequency but compressing its range toward lower expertise levels, the density of old workers rises on part of the skilled range, and the willingness to pay can increase with prior automation, which is the self-reinforcing configuration.

Within the task-bundling framework, destruction is the natural case below the threshold: occupations where expert tasks are all automated retain only generic tasks, so they offer no expert-task exposure toward which learning could be redirected. This reflects the task-specificity of on-the-job human capital (Gibbons and Waldman 2004): experience accumulated on generic tasks does not build expertise, so learning that is no longer supported by expert-task exposure is lost rather than diverted to lower expertise levels.[10] Jones (2026) makes the same observation about the existence of such occupations, noting that some jobs may consist of bundles of tasks that are all automated, in which case their wages fall to the marginal cost of performing those tasks with machines. Redirection instead requires that workers continue to perform expert tasks, with machine assistance, which in our model occurs only in occupations above the threshold.[11]

Figure 2 shows the erosion mechanism itself: a higher inherited automation level shifts the population expertise density toward lower expertise levels, raising mass below the threshold and reducing it above, a deterioration in the first-order stochastic dominance sense (Proposition 3 in Appendix C).

At the occupation level, the willingness to pay is positive only in occupations above the automation threshold, and it declines with prior automation at all these occupations.

[10] Automation removes the practice through which skill is acquired and maintained: operators of automated systems lose the skills the system performs for them (Bainbridge 1983), and pilots' manual flying skills decline measurably as automation reduces hand-flying practice (Haslbeck and Hoermann 2016).

[11] Assumption 1 leaves learning unimpaired in occupations above the automation threshold, although the evidence on cognitive offloading suggests that machine assistance may also reduce learning in expert tasks that remain human. This form of impairment is again a destruction of learning events. Suppose that the learning probability of a worker in occupation $\phi > I_0$ is proportional to the share of her occupation's tasks that are not automated, with the distribution of expertise conditional on learning unchanged. A higher I_0 then removes learning events throughout the skilled range, and the workers concerned remain at their current expertise instead of advancing. The conclusion of Proposition 1 continues to hold under this extension, which we verify numerically over the parameter range of section 5.

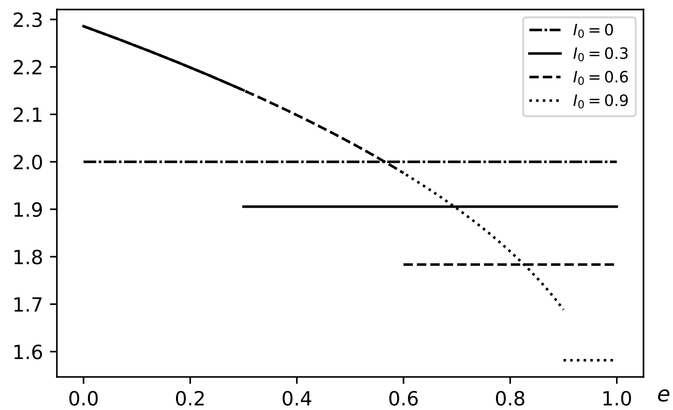

Figure 2: Population expertise density $f(e; I_0)$ for inherited automation $I_0 \in \{0, 0.3, 0.6, 0.9\}$, when a young worker learns with probability $p = 0.60$ above the automation threshold and $q = 0.40$ below it.

4 Endogenous automation dynamics

Section 3 established that the static value of marginal automation declines with the inherited level I_0: expertise erosion is an externality that decreases the willingness to pay for automation. We now extend the model to a dynamic environment in which automation is chosen each period by an incumbent AI producer who faces a displacement risk increasing in the number $N \geq 1$ of AI producers. In this section, we study whether the market mitigates this externality, and whether the equilibrium provides too much or too little automation relative to the social planner's path. The direction of the inefficiency determines the direction of the corrective policy.

4.1 Setup

Time is discrete, indexed by $t = 0, 1, 2, \ldots$, and the level of automation at the beginning of period 0 is zero. We write I_{t-1} for the level of automation at time $t-1$, which is inherited at the start of period t and plays the role of the level denoted I_0 in sections 2 and 3. The discount factor is $\beta \in (0, 1)$. Within each period, the equilibrium is as in section 3: output prices and wages adjust to the current level of automation and to the inherited expertise distribution.

Automation depreciates at rate $\delta \in (0, 1]$, where $\delta = 1$ corresponds to a fully depreciating technology. The parameter δ summarizes two features of software-associated automation: existing models and code become obsolete as tasks, tools, and work processes change, and firms obtain the part that remains usable at no cost, by imitation or open replication. The freely available automation at the start of period t is $\bar{I}_t \equiv (1-\delta)I_{t-1}$, and

$$I_t = (1-\delta)I_{t-1} + \Delta_t, \tag{8}$$

where $\Delta_t \geq 0$ is net investment by the AI producer. The model uses only the product of the two rates that determine δ: the proportion of existing automation that becomes obsolete, and the proportion of the remainder that firms obtain at no cost. Since a period is about twenty years, an annual obsolescence rate of 15 percent, the benchmark rate for knowledge capital in the empirical literature (Peters and Taylor 2017), renders all but approximately 4 percent of existing automation obsolete, and if firms obtain half of that remainder at no cost, the resulting δ exceeds 0.97. Both rates are high for software-associated automation, which depreciates quickly (Li and Hall 2020). The empirically relevant range for δ is therefore close to full depreciation.

The cost of net investment is convex in Δ_t and increasing in the level of automation:

$$C(\Delta_t, I_t) = \frac{c}{2}\frac{\Delta_t^2}{(1-I_t)^2}, \qquad \frac{dC}{dI_t} = \frac{c\Delta_t(1-\bar{I}_t)}{(1-I_t)^3}, \tag{9}$$

where $\frac{dC}{dI_t}$ is a total derivative with $\Delta_t = I_t - \bar{I}_t$. The $(1-I_t)^{-2}$ scaling makes full automation infinitely expensive.

Let $\Omega(I_t, I_{t-1})$ denote equilibrium aggregate output in a period with automation I_t and the expertise density $f(\cdot\,; I_{t-1})$ that depends on prior automation, as characterized in section 3. Higher prior automation worsens the inherited expertise distribution in a FOSD sense (Proposition 3 in Appendix C), which lowers equilibrium output at every level of automation.

The aggregate willingness to pay of firms for a marginal increase in automation is given by aggregating equation (7) across occupations, at the output prices and wages of the equilibrium with automation I_t:

$$W(I_t, I_{t-1}) \equiv \int_{I_t}^{1} \frac{\partial\left(P_y(\phi)Y(\phi)\right)}{\partial I} d\phi > 0. \tag{10}$$

The aggregate willingness to pay is strictly decreasing in the prior level of automation I_{t-1}, because expertise erosion reduces the share of the high-expertise workers whose productivity automation raises (Proposition 1).[12] The integral starts at I_t because occupations at or below the current threshold contribute nothing to a firm's willingness to pay for further

[12]Proposition 1 is stated for $I_0 \leq I < 1$. The dynamic analysis below also evaluates W at levels of automation below the inherited level, where the discontinuity of the population density at the inherited threshold lies inside the integration range and the proof of Proposition 1 does not apply. The monotonicity of W in its second argument over this range is verified numerically, as stated below.

automation: all expert tasks in those occupations are performed by machines, so a marginal improvement in automation does not change their output. Firms nonetheless benefit from an increase in the threshold, because the workers released by the occupation at the threshold produce more in the pool than that occupation produced under expert production. Each firm is atomistic and takes the pooled labor supply as given, whether or not it purchases automation, so a firm's willingness to pay does not take this into account. The gain from reassigning the released workers is the unappropriated surplus

$$\Gamma(I_t, I_{t-1}) \equiv \frac{\partial \Omega}{\partial I_t} - W(I_t, I_{t-1}), \tag{11}$$

which is the difference between the social and the private marginal value of automation.[13] This appropriability failure is standard for innovation under competitive adoption.

The analysis requires four conditions. First, prior automation strictly reduces the demand for automation: $\frac{\partial W}{\partial I_{t-1}} < 0$. Second, it strictly reduces equilibrium output: $\frac{\partial \Omega}{\partial I_{t-1}} < 0$. Third, the unappropriated reassignment surplus is strictly positive: $\Gamma > 0$. Fourth, the first-best steady state is unique.[14] Proposition 1 establishes the first two conditions, over a restricted range and with a weak inequality. We verify the strengthenings and the last two numerically over the parameter range of section 5. We also assume that the value functions are differentiable and that the optimum is interior. We do not repeat these conditions below.

The depreciated level $\bar{I}_t$ is freely available to firms, so the AI producer can only charge for improvements relative to this baseline. In period t, the producer earns the following revenue:[15]

$$\int_{\bar{I}_t}^{I_t} W(u, I_{t-1}) du. \tag{12}$$

4.2 The incumbent's problem

There are $N \geq 1$ symmetric, risk-neutral potential AI producers. At the start of period 0, a producer is drawn uniformly as the incumbent. In every period t, the incumbent chooses $I_t \geq \bar{I}_t$ and sells the resulting improvements to firms,[16] and earns the period payoff

[13]Output price adjustments contribute no first-order term to $\frac{\partial \Omega}{\partial I_t}$, so the willingness to pay of price-taking firms equals the marginal product of automation net of the value of reassigning the released workers. Appendix A gives Γ in closed form, equation (16).

[14]We verify this by checking that the planner's steady-state first-order condition crosses zero once, from above, which is the form used in the proof of Lemma 4.

[15]This expression is the entire area under the willingness-to-pay curve over incremental automation, so the producer is assumed to price discriminate perfectly over $[\bar{I}_t, I_t]$. Under a single price per unit of automation, it would obtain a smaller revenue and would invest less, which would widen the range of depreciation rates at which the equilibrium level of automation is below the first-best level.

[16]Firms are atomistic and take prices and wages as given, so each firm's willingness to pay for a marginal improvement is its marginal revenue product, by equation (7), evaluated at the equilibrium prices and wages.

$\int_{\bar{I}_t}^{I_t} W(u, I_{t-1})du - C(\Delta_t, I_t)$. At the end of the period, the incumbent retains its position with probability $\frac{1}{N}$; with probability $\frac{N-1}{N}$ it is permanently displaced by an entrant and earns zero thereafter.[17] We refer to $1 - \frac{1}{N}$ as the displacement risk: the case $N = 1$ is a permanent incumbent, and the limit $N \to \infty$ an industry in which the incumbent is displaced every period.

We restrict attention to symmetric Markov-perfect equilibria. Appendix B states the incumbent's Bellman equation, its first-order condition (FOC), and the envelope condition for its value V_N. The envelope condition has three terms, which constitute the incumbent's entire exposure to its inherited level of automation. Two of them lower the incumbent's value. First, prior automation reduces the willingness to pay (Proposition 1), which lowers the producer's revenue on each improvement that it sells. Second, a higher level in one period raises the freely available baseline next period, which cannibalizes future sales by reducing the improvements for which the incumbent can charge. The third term raises the incumbent's value: a higher inherited level reduces the investment required to reach any given I_t.

4.3 The planner's problem

The social planner maximizes the discounted sum of aggregate output net of investment cost:

$$V_{FB}(I_{t-1}) = \max_{I_t \geq \bar{I}_t} \left\{ \Omega(I_t, I_{t-1}) - C(\Delta_t, I_t) + \beta V_{FB}(I_t) \right\}, \tag{13}$$

Appendix B states its FOC and envelope condition, together with the equilibrium and planner Euler equations, equations (22) and (23), which follow by substituting the envelope conditions into the respective FOCs.

Comparing the two Euler equations term by term, the equilibrium condition differs from the planner's in four respects. The marginal benefit is the demand W rather than the marginal product $\frac{\partial \Omega}{\partial I_t} = W + \Gamma$. The continuation is discounted at rate $\frac{\beta}{N}$ rather than β. The equilibrium condition contains the cannibalization term $-(1-\delta)W(\bar{I}_{t+1}, I_t)$, which is absent from the planner's condition. Finally, erosion enters the equilibrium condition as the reduction in demand integrated over incremental automation, $\int_{\bar{I}_{t+1}}^{I_{t+1}} \frac{\partial W}{\partial I_{t-1}}(u, I_t)du$, and the planner's condition as the reduction in output, $\frac{\partial \Omega}{\partial I_{t-1}}(I_{t+1}, I_t)$. The reduction in future

A producer posting a price marginally below the aggregate willingness to pay for each marginal unit sells to all firms. We take this outcome as given rather than modeling the price-setting protocol.

[17] The displacement hazard is reduced-form. Only the effective continuation weight $\frac{\beta}{N}$ enters the equilibrium conditions, so the displacement risk operates as a survival-adjusted discount factor on the incumbent's continuation value, and any process that delivers the same expected discounted continuation yields identical predictions.

investment costs is common to both conditions.

These wedges have two sources. The first is expiring appropriability: the automation sold in earlier periods depreciates into the freely available baseline, which raises the level above which the producer can charge. The second is incomplete appropriability: the producer does not obtain the part Γ of the surplus from automation.

As $N \to \infty$, the continuation weight converges to zero and the equilibrium condition reduces to the static condition $W = \frac{dC}{dI_t}$: the incumbent will be displaced, and makes its investment decision by considering only the current period. For $1 < N < \infty$, the equilibrium condition lies between the permanent-incumbent condition and this static limit.

4.4 Direction of the inefficiency

At a steady state I, where $I_t = I_{t-1} = I$ and $\Delta_t = \delta I$, the marginal value of the inherited level of automation is $V'_{FB}(I)$ to the planner and $V'_N(I)$ to the incumbent. Both are derived in Appendix B, which also shows that $V'_N(I)$ does not depend on N at a steady state, since the retention probability does not enter the envelope condition. Define the wedge function:

$$\Phi_N(I) \equiv \Gamma(I, I) + \beta \left[V'_{FB}(I) - \frac{1}{N} V'_N(I) \right]. \tag{14}$$

Lemma 4. *Let $I^*(N)$ be an interior steady state of the equilibrium. The sign of $I^{FB} - I^*(N)$ is equal to the sign of $\Phi_N(I^*(N))$.*

The equilibrium provides too little automation when $\Phi_N(I^*(N)) > 0$ and too much when $\Phi_N(I^*(N)) < 0$. Equation (14) separates the components of the wedge according to the direction in which they operate. Two appropriability wedges reduce the equilibrium level of automation relative to the first best, and neither involves expertise erosion. The first is the unappropriated surplus: the price of a marginal increase in automation is the aggregate willingness to pay W rather than the social marginal product $W + \Gamma$, so the producer obtains less than the social gain generated by each increment of automation that it sells. The second is the valuation of the continuation. The planner values the inherited level at $\beta V'_{FB}$, which at low depreciation rates is dominated by the reduction in future investment costs. The producer weights its continuation by the retention probability $\frac{1}{N}$, so it obtains that reduction only if it retains the market, and it sets that reduction against cannibalization, since every unit sold in one period becomes a freely available substitute for the units it would sell next period.

Expertise erosion operates in the opposite direction. The planner internalizes the reduction in output and therefore restrains investment. The producer internalizes only the reduction in the demand for the incremental automation that it sells, which is one component of that reduction and is weighted by the retention probability $\frac{1}{N}$ (Proposition 2). The producer therefore restrains investment less than the planner does, and this component raises the equilibrium level of automation above the first-best level. In sum, the direction of the inefficiency depends on the durability of automation and on the severity of expertise erosion.[18]

Proposition 2 isolates the part of expertise erosion that the market internalizes.

Proposition 2. *(i) For every N, the AI producer internalizes expertise erosion only through the decline in the willingness to pay over incremental automation ($\int_{\bar{I}_{t+1}}^{I_{t+1}} \frac{\partial W}{\partial I_{t-1}}(u, I_t)du$).*
(ii) If automation fully depreciates ($\delta = 1$), the equilibrium provides too much automation if and only if the discounted reduction in output caused by expertise erosion exceeds the unappropriated surplus plus the discounted reduction in demand that the producer internalizes, all evaluated at $I^*(N)$.

Part (i) explains why no market structure fully prices the externality: the reductions in the output produced with the freely available baseline level of automation and in the unappropriated surplus caused by expertise erosion do not enter the AI producer's objective function, for any N. Part (ii) removes the durable-goods channels: with a fully depreciating technology, there is no cannibalization and no inherited automation to reduce investment costs. Among the three remaining terms, only the reduction in output caused by expertise erosion pushes toward over-automation, and it is discounted while the unappropriated surplus is not. As a result, over-automation requires severe erosion and a high discount factor (section 5).

These results rely on the conditions stated in section 4.1, and they hold for any learning process that satisfies Assumption 1, with no functional form imposed on that process (the functional form determines the magnitude of the components of Φ_N).

4.5 Corrective policy

The direction of the inefficiency determines the corrective instrument. Consider a proportional subsidy at rate s per unit of net investment Δ_t, paid to the incumbent. A negative subsidy is a tax on the incumbent's net investment.

[18]The wedge is the convex combination $\Phi_N = \left(1 - \frac{1}{N}\right)\Phi_\infty + \frac{1}{N}\Phi_1$ of its values at the two extreme market structures, so varying N changes the size of the inefficiency but neither its direction nor its existence whenever the two extremes share the same sign, as they do throughout the numerical computations. Appendix A shows how the displacement risk moves the steady state.

Corollary 1. *The rate*

$$s^* = \frac{\Phi_N\big(I^{FB}\big)}{1 - \frac{\beta}{N}(1-\delta)} \tag{15}$$

implements I^{FB} as a steady state of the equilibrium.

The denominator of s^* is positive, so the instrument is a subsidy when $\Phi_N(I^{FB})$ is positive and a tax when it is negative. It corrects all three components of Φ_N simultaneously, because all three enter the same FOC for investment. Provided that Φ_N does not change sign between $I^*(N)$ and I^{FB}, which holds throughout the numerical computations, the instrument is a subsidy if and only if the equilibrium level of automation is below the first-best level.

The presence of a dynamic externality does not by itself determine the direction of the intervention. A tax on investment in automation is optimal only when the reduction in output caused by expertise erosion exceeds the unappropriated surplus and the reduction in demand that the producer internalizes, and a subsidy is optimal otherwise. This contrasts with the presumption that dynamic externalities from transformative technologies justify slowing their adoption (Jones 2024, 2026; Acemoglu and Lensman 2024).

5 Numerical illustration

We illustrate the analysis using the specific learning process of Appendix C and two examples that involve a different learning process: $(p, q) = (0.60, 0.40)$ in the first, and $(p, q) = (0.90, 0)$ in the second, where workers below the automation threshold never learn, so that the adverse effect of automation on learning is more severe in the second. In both examples, $\eta = 2$, $\lambda = 2$, $\theta = 1.5$, $c = 1.5$, and, unless otherwise noted, $\delta = 0.97$. The discount factor is $\beta = 0.67$, which is implied by an annual rate of 2 percent over a period of about twenty years. The parameter values are illustrative, so we report signs, orderings, and comparative statics rather than magnitudes. See Appendix B for analytical and computational details, including the treatment of the discontinuity of the population density at the inherited threshold and the robustness of the reported signs to that discontinuity.

There is under-automation in the first example: $I^{FB} = 0.271$ against $I^*(1) = 0.268$ and $I^*(\infty) = 0.269$. There is over-automation in the second example: $I^{FB} = 0.253$ against $I^*(1) = 0.254$ and $I^*(\infty) = 0.258$. According to Corollary 1, the corrective instrument is a subsidy in the first example, at a rate of 0.022 for $N = 1$ and 0.017 for $N \to \infty$, and a tax in the second, at 0.007 and 0.032. In both examples, a permanent incumbent ($N = 1$) is associated with the lowest steady-state level of automation: it internalizes that current automation increases the freely available baseline in the next period, which cannibalizes its future sales, and that expertise erosion reduces future demand for automation. By contrast, a producer exposed to displacement discounts these effects.

Figure 3 plots the gap $I^{FB} - I^*(N)$ as a function of the depreciation rate. In the first example the gap is positive throughout, falling from 0.097 at $\delta = 0.10$ to 0.003 at $\delta = 0.97$ for $N = 1$. In the second it changes sign, between $\delta = 0.90$ and $\delta = 0.95$ for $N \to \infty$ and between $\delta = 0.95$ and $\delta = 0.97$ for $N = 1$. The decomposition of Φ_N accounts for this profile. When δ is close to 100 percent, the reduction in future investment costs and cannibalization are negligible (Proposition 2), so the sign is determined by the comparison between the reduction in output caused by expertise erosion and the unappropriated surplus, and the former exceeds the latter only when expertise erosion is sufficiently severe. As δ falls, the planner continuation $\beta V'_{FB}$ becomes the largest positive component, because the inherited level reduces future investment costs, and cannibalization grows as well, so both examples provide too little automation.

At low depreciation rates, the direction of the inefficiency is robust to the remaining parameters, including a range of labor shares from $\frac{1}{3}$ to $\frac{8}{9}$, and steady-state automation is increasing in N in every configuration that we compute. The level of automation does not determine the direction either: reducing c so that the first-best level of automation rises from 0.26 to 0.65 leaves the first example below that level and widens the distance.

Neither the depreciation rate nor the severity of the learning losses determines the direction of the inefficiency on its own. The discount factor matters as well. A period of the model spans a working generation, so β and δ refer to the same horizon: over twenty years, an annual rate of 5 percent gives $\beta = 0.36$ and an annual rate of 2 percent gives $\beta = 0.67$. At $\beta = 0.36$, the equilibrium provides too little automation for any depreciation rate in both examples, so that the corrective instrument is a subsidy. Thus, over-automation requires severe expertise erosion, fast depreciation, and a discount rate low enough for the discounted reduction in output from expertise erosion to be higher than the unappropriated surplus.

6 Regulatory and managerial implications

Competition policy is not a substitute for the Pigouvian instrument of Corollary 1. Indeed, Φ_N cannot be set to zero by changing N in any configuration that we compute. Neither a permanent incumbent nor an incumbent displaced every period restores efficiency. Since the demand for future automation depends on the current expertise distribution (Proposition 1), an incumbent facing a lower displacement risk has a greater private interest in preserving learning opportunities. Its objective is nonetheless the revenue from the automation that it sells rather than aggregate output, so even a permanent incumbent does not take into ac-

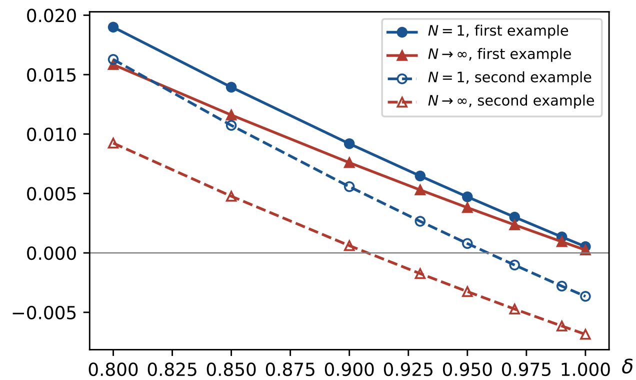

Figure 3: Gap to the first best $I^{FB} - I^*(N)$ as a function of δ, for $N = 1$ and $N \to \infty$, in the two examples. The range of δ is restricted to the values that section 4 identifies as empirically relevant. The gap narrows as δ rises, and changes sign in the second example.

count the reduction in aggregate output caused by expertise erosion, or the unappropriated surplus.

Intellectual property protection does not substitute for it either. Both change the supply of automation by affecting the producer's continuation weight and the share of the surplus that it captures, whereas both externalities operate through the demand for automation. Firms' willingness to pay excludes the unappropriated surplus Γ, which they obtain whether or not they purchase automation, and it excludes the erosion of future output, which they do not internalize because expertise is possessed by workers and because a firm's automation decision does not change the economy-wide distribution of expertise.

The inefficiency stems from expiring appropriability rather than from market power (section 4), which connects the analysis to the theory of durable-goods monopoly (Coase 1972). Since the depreciated level becomes freely available, the AI producer competes against its own past sales, and faster depreciation increases the amount it can charge, which is analogous to planned obsolescence (Bulow 1986). By contrast, expertise erosion reduces the willingness to pay at any level of automation (Proposition 1), which makes the AI producer worse off. Strengthening intellectual property over automation therefore transfers surplus to the AI producer, and it eliminates the appropriability wedges in the limit of perpetual appropriability with a permanent incumbent.

The direction of the intervention depends on two variables that can be measured empirically: the depreciation rate of automation, discussed in section 4, and the strength of the learning losses from automation – for which the evidence on cognitive offloading (Grinschgl and Neubauer 2022; Kosmyna et al. 2025) and on declining entry-level hiring (Hosseini Maasoum and Lichtinger 2025) is directly relevant. At the depreciation rates that section 4 identifies as empirically relevant, the two examples of section 5 involve the opposite inefficiency, although they differ only in the severity of the learning losses.

The model also delivers managerial implications. It implies that assigning entry-level expert tasks to workers rather than to machines is an investment in the future expertise distribution. Since expertise is possessed by workers who are hired in spot markets, its return is transferable across firms, so no firm has a claim on the expertise that it creates. Moreover, in the model, learning depends on the economy-wide level of automation, so no individual firm can preserve it unilaterally. Long-term contracts would not change this: with one-sided commitment, the return to learning can only be bought by the worker herself, as studied by Garicano and Rayo (2025), and financing constraints bound what she can pay (Maug 2025). Finally, the appropriability wedges of the AI producer cannot be addressed by a labor contract.

7 Conclusion

We study the potentially self-reinforcing nature of advances in automation because of their adverse effect on expertise accumulation. In our model, additional automation changes the expertise distribution in two ways. First, it redistributes expertise below the automation threshold, which is inconsequential, because no expertise is needed in these occupations where expert tasks are already fully automated. Second, it decreases expertise above the threshold. As a result, the private value for a firm of further advances in automation, and therefore the aggregate demand faced by suppliers of automation, is negatively related to the level of automation previously reached: further advances improve the productivity of high-expertise workers, whose share is decreasing in prior automation. This slows the feedback loop between pre-existing automation and further advances in expertise-substituting technologies.

We then extend the analysis to a dynamic environment in which automation is chosen by an incumbent AI producer that faces a displacement risk. No market structure eliminates the erosion externality: the producer internalizes erosion only through the decline in the willingness to pay for incremental automation. The equilibrium provides too little automation when the unappropriated surplus from worker reassignment, the cannibalization of future sales, and the reduction in future investment costs dominate, and too much automation when the reduction in output caused by expertise erosion dominates. The inefficiency can be corrected by a subsidy on the AI producer's investment in the former case and by a tax in the latter case.

A Proofs

Proof of Lemma 1: The total task measure in occupation ϕ is $\phi + g(\phi)$, and the non-automated task measure is $\max\{\phi - I, 0\} + g(\phi)$. Uniform allocation of labor across non-automated tasks yields per-worker output $\varphi(\phi; I)$, so $Y(\phi) = (L(\phi)\varphi(\phi; I))^{\frac{\eta}{\eta+1}}$. ■

Proof of Lemma 2: The firm chooses $L(\phi)$ to maximize occupation-level profit $\pi(\phi) = P_y(\phi)Y(\phi) - L(\phi)w(\phi)$. The FOC gives $L(\phi) = \varphi(\phi; I)^{\frac{\eta(\lambda-1)}{\lambda+\eta}} \left[\frac{\eta}{\eta+1} Y^{\frac{1}{\lambda}} w(\phi)^{-1}\right]^{\frac{\lambda(\eta+1)}{\lambda+\eta}}$. For $\phi > I$, labor supply equals $f(\phi)$, giving the skilled wage. For $\phi \leq I$, integrating labor demand and equating to $F(I)$ gives $\bar{w}$ and $L(\phi)$. ■

Proof of Lemma 3: From the FOC, $L(\phi)w(\phi) = \frac{\eta}{\eta+1} P_y(\phi)Y(\phi)$. ■

Proof of Proposition 1: Consider a marginal increase $dI_0 > 0$ with $I \geq I_0$ fixed. Young workers with expertise $z \in (I_0, I_0 + dI_0]$ switch from the learning distribution $G_a(\cdot|z)$ to $G_b(\cdot|z)$. By Assumption 1, $g_b(x|z) \leq g_a(x|z)$ for almost every $x > z$, so the density of old workers satisfies

$$\frac{\partial f_o(e_i; I_0)}{\partial I_0} = f_y(I_0)\left[g_b(e_i|I_0) - g_a(e_i|I_0)\right] \leq 0 \quad \text{for all } e_i > I_0,$$

with strict inequality on a set of positive measure in every interval $(c, 1]$ by the strict tail condition in Assumption 1. Since f_y does not depend on I_0, the population density $f(e_i) = f_y(e_i) + f_o(e_i)$ is non-increasing in I_0 for every $e_i > I_0$, and strictly decreasing on a set of positive measure in $(I, 1]$. Because the switching workers' probability mass moves only downward, from expertise levels above I_0 toward I_0, the population distribution also deteriorates in the first-order stochastic dominance sense.

The equilibrium allocation maximizes aggregate output given the expertise distribution: the program of assigning workers to occupations to maximize Y, subject to the constraint that a worker of expertise e can be assigned only to occupations $\phi \leq e$ or fully automated occupations, is concave, and its FOCs coincide with the equilibrium conditions of Lemma 2 with the multipliers equal to wages. A first-order stochastically dominated distribution shrinks the feasible set of this program: matching workers by quantile, each worker in the better economy weakly exceeds her counterpart in expertise, so any assignment feasible under the worse distribution can be replicated under the better one. It follows that Y is non-increasing in I_0.

Equilibrium occupation revenue is $P_y(\phi)Y(\phi) = Y^{\frac{1}{\lambda}} Y(\phi)^{\frac{\lambda-1}{\lambda}} = Y^{\frac{1}{\lambda}} \left(f(\phi)\varphi(\phi; I)\right)^{\frac{\eta(\lambda-1)}{\lambda(\eta+1)}}$ for $\phi > I$. Since the atomistic firm takes prices as given, its marginal willingness to pay at occupation $\phi > I$ is, by equation (7),

$$m(\phi; I, I_0) = \frac{\partial \left(P_y(\phi)Y(\phi)\right)}{\partial I} = \frac{\eta}{\eta+1}\frac{P_y(\phi)Y(\phi)}{\phi - I + g(\phi)} > 0,$$

evaluated at equilibrium prices, and is zero for $\phi \leq I$ (all expert tasks already automated). Since $\lambda > 1$, the exponent on $f(\phi)$ is positive, so both factors of $P_y(\phi)Y(\phi)$ are non-increasing in I_0: the aggregate factor $Y^{\frac{1}{\lambda}}$ falls, and the density factor falls at every $\phi > I$, strictly on a set of positive measure; $\varphi(\phi; I)$ and $\phi - I + g(\phi)$ do not depend on I_0. Therefore, $m(\phi; I, I_0)$ is non-increasing in I_0 at every $\phi > I$, and the aggregate $\int_I^1 m(\phi; I, I_0)d\phi$ is strictly decreasing in I_0. ■

The unappropriated surplus Γ (section 4): Let $A_0(I_t; I_{t-1}) \equiv \int_0^{I_t} \left(L(\phi)\varphi(\phi; I_t)\right)^{\frac{\eta(\lambda-1)}{\lambda(\eta+1)}} d\phi$ denote the pooled-sector aggregate, with L the assignment of Lemma 2. By Lemmas 1 to 3, equilibrium aggregate output is $\Omega = A^{\frac{\lambda}{\lambda-1}}$ with $A = A_0(I_t; I_{t-1}) + \int_{I_t}^1 \left(f(\phi)\varphi(\phi; I_t)\right)^{\frac{\eta(\lambda-1)}{\lambda(\eta+1)}} d\phi$, and $A^{\frac{1}{\lambda-1}} = \Omega^{\frac{1}{\lambda}}$, so that $\frac{\partial\Omega}{\partial I_t} = \frac{\lambda}{\lambda-1}\Omega^{\frac{1}{\lambda}}\frac{\partial A}{\partial I_t}$. For $\phi > I_t$, $\frac{\partial\varphi}{\partial I} = \frac{\varphi}{\phi - I + g(\phi)}$ and $\varphi(\phi; I_t) \to \varphi(I_t; I_t)$ as $\phi \downarrow I_t$, so the derivative of the skilled integral is $-\left(f(I_t)\varphi(I_t; I_t)\right)^{\frac{\eta(\lambda-1)}{\lambda(\eta+1)}} + \frac{\eta(\lambda-1)}{\lambda(\eta+1)}\int_{I_t}^1 \frac{(f\varphi)^{\frac{\eta(\lambda-1)}{\lambda(\eta+1)}}}{\phi - I_t + g(\phi)} d\phi$. Using $P_y(\phi)Y(\phi) = \Omega^{\frac{1}{\lambda}}\left(f\varphi\right)^{\frac{\eta(\lambda-1)}{\lambda(\eta+1)}}$, the identity $\frac{\lambda}{\lambda-1}\frac{\eta(\lambda-1)}{\lambda(\eta+1)} = \frac{\eta}{\eta+1}$, and the fixed-price marginal willingness to pay $\frac{\partial(P_yY)}{\partial I} = \frac{\eta}{\eta+1}\frac{P_yY}{\phi - I + g(\phi)}$ from the proof of Proposition 1, the integral term equals $W(I_t, I_{t-1})$ by equation (10). The remaining terms give the unappropriated reassignment surplus in closed form:

$$\Gamma(I_t, I_{t-1}) = \frac{\lambda}{\lambda-1}\Omega^{\frac{1}{\lambda}}\left(\frac{\partial A_0}{\partial I_t} - \left(f(I_t; I_{t-1})\varphi(I_t; I_t)\right)^{\frac{\eta(\lambda-1)}{\lambda(\eta+1)}}\right). \tag{16}$$

Output price adjustments contribute no first-order term to $\frac{\partial\Omega}{\partial I_t}$, so the willingness to pay of price-taking firms accounts for the entire marginal product except for the unappropriated reassignment surplus. When $I_t = I_{t-1}$, the density has a jump at the threshold, and the formula gives the one-sided derivatives of Ω with the corresponding one-sided limits of f. ■

Proof of Lemma 4: At a steady state the equilibrium policy maps $I^*(N)$ to itself, so equation (21) can be evaluated at $I_t = I_{t-1} = I^*(N)$, which gives equation (24). Substituting into equation (20) yields $W - \frac{dC}{dI_t} = -\frac{\beta}{N}V_N'(I^*(N))$ at $I^*(N)$. Since $\frac{\partial\Omega}{\partial I_t} = W + \Gamma$, the planner's steady-state FOC residual E^{FB}, defined in equation (25), satisfies $E^{FB}(I^*(N)) = \Gamma(I^*(N), I^*(N)) + \beta V_{FB}'(I^*(N)) - \frac{\beta}{N}V_N'(I^*(N)) = \Phi_N(I^*(N))$. Since E^{FB} is strictly positive below I^{FB} and strictly negative above it, the residual at $I^*(N)$ is positive if and only if $I^*(N) < I^{FB}$. ■

Comparative statics in the displacement risk (section 4.4): Fix δ and suppose that, for each N, the steady-state form of equation (22) crosses zero once from above in I. We show that $\operatorname{sign}(I^*(N') - I^*(N)) = -\operatorname{sign}(V_N'(I^*(N)))$ for $N' > N$. At a steady state, equation (22) reads $G_N(I) \equiv W(I,I) - \frac{dC}{dI_t} + \frac{\beta}{N}V_N'(I) = 0$. For $N' > N$, $G_{N'}(I^*(N)) = \beta\left(\frac{1}{N'} - \frac{1}{N}\right)V_N'(I^*(N))$, whose sign is the opposite of that of $V_N'(I^*(N))$ since $\frac{1}{N'} < \frac{1}{N}$. If $G_{N'}(I^*(N)) > 0$, single crossing from above implies $I^*(N') > I^*(N)$; the case $G_{N'}(I^*(N)) < 0$ is symmetric, and $G_{N'}(I^*(N)) = 0$ gives $I^*(N') = I^*(N)$. ■

Proof of Proposition 2: (i) The producer's period revenue is $\int_{\bar{I}_t}^{I_t} W(u, I_{t-1})du$ by equation (12), and I_{t-1} enters it only through the integrand and the lower limit. Differentiating under the integral sign gives the demand-erosion term of equation (21), which, evaluated at $t+1$, is $\int_{\bar{I}_{t+1}}^{I_{t+1}} \frac{\partial W}{\partial I_{t-1}}(u, I_t)du$. The cannibalization term comes from the lower limit and does not involve erosion. The cost function depends on I_{t-1} only through $\bar{I}_t$ and not through the expertise distribution, so it carries no erosion either. Aggregate output does not enter the AI producer's objective function, so the reductions in the output produced with the freely available baseline level of automation and in Γ caused by expertise erosion appear nowhere in equation (22). (ii) At $\delta = 1$, $\bar{I}_{t+1} = 0$, so the cost-relief term in V_{FB}' and V_N' and the cannibalization term in V_N' carry the factor $1-\delta$ and vanish, leaving $V_N'(I) = \int_0^I \frac{\partial W}{\partial I_{t-1}}(u, I)du$. It follows that $\Phi_N(I) = \Gamma(I,I) + \beta\frac{\partial \Omega}{\partial I_{t-1}}(I,I) - \frac{\beta}{N}\int_0^I \frac{\partial W}{\partial I_{t-1}}(u, I)du$, and the sign characterization follows from Lemma 4. ■

Proof of Corollary 1: With the subsidy, the term $+s$ is added to the left-hand side of equation (20), and the term $-(1-\delta)s$ is added to the right-hand side of equation (21). Imposing a steady state at I^{FB} and substituting the envelope into the FOC gives $E^{FB}(I^{FB}) - \Phi_N(I^{FB}) + s\left(1 - \frac{\beta}{N}(1-\delta)\right) = 0$; since $E^{FB}(I^{FB}) = 0$, the result follows. The denominator is positive because $\frac{\beta}{N}(1-\delta) \leq \beta < 1$. ■

B Equilibrium and computation

The population density $f(e; I_{t-1})$ is available in closed form: with $f_y(e) = (1 + \frac{p}{2-p})(1-e)^{\frac{p}{2-p}}$ normalized to unit mass, the density of old workers follows equation (27), and $f = f_y + f_o$. Equilibrium aggregate output follows from Lemmas 1 to 3:

$$\Omega(I, J) = \left(\int_0^1 \left(L(\phi)\varphi(\phi; I)\right)^{\frac{\eta(\lambda-1)}{\lambda(\eta+1)}} d\phi\right)^{\frac{\lambda}{\lambda-1}}, \tag{17}$$

where J denotes the inherited level of automation and $L(\phi)$ is the equilibrium assignment of Lemma 2 given the density $f(\cdot\,; J)$: $L(\phi) = f(\phi; J)$ above the automation threshold and the pooled allocation below it. The demand of equation (10) is

$$W(I,J)=\frac{\eta}{\eta+1}\Omega(I,J)^{\frac{1}{\lambda}}\int_{I}^{1}\frac{\left(f(\phi;J)\varphi(\phi;I)\right)^{\frac{\eta(\lambda-1)}{\lambda(\eta+1)}}}{\phi-I+g(\phi)}d\phi, \tag{18}$$

and the producer's revenue follows by integration, equation (12). The marginal product $\frac{\partial\Omega}{\partial I}$ is available in closed form as well: its smooth part equals W, as established below equation (11), and the residual Γ is given by equation (16). The erosion derivatives entering the Euler equations are computed numerically. We compute steady states as roots of the steady-state Euler equations, and we verify them by value-function iteration on a grid, whose fixed point coincides with the Euler root to within the resolution of the grid. **The incumbent's problem.** The incumbent's Bellman equation is

$$V_N(I_{t-1})=\max_{I_t\geq\bar{I}_t}\left\{\int_{\bar{I}_t}^{I_t}W(u,I_{t-1})du-C(\Delta_t,I_t)+\frac{\beta}{N}V_N(I_t)\right\}. \tag{19}$$

The function V_N is the expected present value of the incumbent's profits, the expectation being taken over displacement; under risk neutrality, $V_N(I_{t-1})$ is the market value of the incumbent AI producer at the start of a period with inherited automation I_{t-1}. The FOC for an interior optimum is

$$W(I_t,I_{t-1})+\frac{\beta}{N}V_N'(I_t)=\frac{c\Delta_t(1-\bar{I}_t)}{(1-I_t)^3}, \tag{20}$$

and, at an interior optimum, the envelope theorem gives

$$V_N'(I_{t-1})=\underbrace{\int_{\bar{I}_t}^{I_t}\frac{\partial W}{\partial I_{t-1}}(u,I_{t-1})du}_{<0\text{ (demand erosion)}}-\underbrace{(1-\delta)W\left(\bar{I}_t,I_{t-1}\right)}_{>0\text{ (cannibalization)}}+\underbrace{(1-\delta)\frac{c\Delta_t}{(1-I_t)^2}}_{>0\text{ (cost relief)}}. \tag{21}$$

The first term is the erosion of the demand for incremental automation, the second the cannibalization of future sales by the freely available baseline, and the third the relief of future investment costs.

The planner's problem. The FOC of equation (13) is $\frac{\partial\Omega}{\partial I_t}+\beta V_{FB}'(I_t)=\frac{dC}{dI_t}$, and the envelope theorem gives $V_{FB}'(I_{t-1})=\frac{\partial\Omega}{\partial I_{t-1}}+(1-\delta)\frac{c\Delta_t}{(1-I_t)^2}$.

Euler equations. Evaluating equation (21) at $t+1$ and substituting it into equation (20), the equilibrium path satisfies

$$W(I_t, I_{t-1}) - \frac{dC}{dI_t} + \frac{\beta}{N}\left[\int_{\bar{I}_{t+1}}^{I_{t+1}} \frac{\partial W}{\partial I_{t-1}}(u, I_t)du - (1-\delta)W(\bar{I}_{t+1}, I_t) + (1-\delta)\frac{c\Delta_{t+1}}{(1-I_{t+1})^2}\right] = 0. \tag{22}$$

The planner's Euler equation follows in the same way from the envelope of equation (13):

$$\frac{\partial \Omega}{\partial I_t}(I_t, I_{t-1}) - \frac{dC}{dI_t} + \beta\left[\frac{\partial \Omega}{\partial I_{t-1}}(I_{t+1}, I_t) + (1-\delta)\frac{c\Delta_{t+1}}{(1-I_{t+1})^2}\right] = 0. \tag{23}$$

Steady-state marginal values. At a steady state I, where $I_t = I_{t-1} = I$ and $\Delta_t = \delta I$, the marginal value of the inherited level of automation to the planner and to the incumbent are

$$\begin{aligned} V'_{FB}(I) &= \frac{\partial \Omega}{\partial I_{t-1}}(I, I) + (1-\delta)\frac{c\delta I}{(1-I)^2}, \\ V'_N(I) &= \int_{(1-\delta)I}^{I} \frac{\partial W}{\partial I_{t-1}}(u, I)du - (1-\delta)W((1-\delta)I, I) + (1-\delta)\frac{c\delta I}{(1-I)^2}. \end{aligned} \tag{24}$$

The expression for V'_N does not depend on N, since the retention probability does not enter the envelope condition. The planner's steady-state FOC residual is

$$E^{FB}(I) \equiv \frac{\partial \Omega}{\partial I_t}(I, I) - \frac{c\delta I\,(1-(1-\delta)I)}{(1-I)^3} + \beta V'_{FB}(I), \tag{25}$$

which vanishes at the first-best steady state I^{FB}. The fourth condition of section 4.1 takes the form $E^{FB}(I) > 0$ for $I < I^{FB}$ and $E^{FB}(I) < 0$ for $I > I^{FB}$, which is the form used in the proof of Lemma 4.

The density discontinuity. The population density jumps at $e = I_{t-1}$, because the learning probability switches discretely from q to p at the inherited threshold, so $\frac{\partial \Omega}{\partial I_t}$ is one-sided where steady states are evaluated. The jump is an artifact of the two-point learning process, not of the economics. When the switch is replaced by a continuous increase from q to p across an interval containing the threshold, the derivative is two-sided, and as that interval narrows the steady states converge to the values reported in section 5, which are the values obtained with the right derivative. We therefore use the right derivative throughout. The reported gaps lie between the values obtained under the two conventions, and the two conventions give the same sign in either example, although they differ more in the second, where the jump is larger. The same discontinuity makes the planner's problem non-concave at high depreciation rates, so its optimal policy alternates between one level above I^{FB} and one level below it. Smoothing the learning probability yields a stationary optimum at the midpoint, and replacing I^{FB} by that midpoint leaves the sign of the gap unchanged in both examples.

C Specific learning process

We put more structure on the learning process to derive the expertise distributions used in the numerical illustrations of section 5.

Assumption 2. A young worker with expertise e_i^y in occupation $\phi > I_0$ has expertise when old: (a) $e_i^o = e_i^y$ with probability $1-p$; (b) e_i^o uniform on $[e_i^y, 1]$ with probability p.

Assumption 3. A young worker with expertise e_i^y in occupation $\phi \leq I_0$ has expertise when old: (a) $e_i^o = e_i^y$ with probability $1-q$, where $q \in [0, p)$; (b) e_i^o uniform on $[e_i^y, 1]$ with probability q.

Assumptions 2 and 3 satisfy Assumption 1 with censoring probability $1 - \frac{q}{p}$: below the threshold, each learning event is lost with probability $1 - \frac{q}{p}$ and otherwise leads to the same uniform draw.

Without automation ($I_0 = 0$), the density of old workers is

$$f_o(e_i) = p \int_0^{e_i} f_y(z) \frac{1}{1-z} dz + (1-p) f_y(e_i). \tag{26}$$

With automation $I_0 \in [0, 1]$,

$$f_o(e_i) = \begin{cases} q \displaystyle\int_0^{e_i} f_y(z) \frac{1}{1-z} dz + (1-q) f_y(e_i) & \text{if } e_i \leq I_0, \\ q \displaystyle\int_0^{I_0} f_y(z) \frac{1}{1-z} dz + p \int_{I_0}^{e_i} f_y(z) \frac{1}{1-z} dz + (1-p) f_y(e_i) & \text{if } e_i > I_0. \end{cases} \tag{27}$$

The population density is $f(e_i) = f_y(e_i) + f_o(e_i)$. When $I_0 = 0$, expertise is uniformly distributed on $[0, 1]$, which is what pins down f_y.

If $I_0 = 0$, the density of the expertise distribution of all workers at $e_i \in [0, 1]$ is, up to a scaling constant,

$$f_y(e_i) + f_o(e_i) = f_y(e_i) + p \int_0^{e_i} f_y(z) \frac{1}{1-z} dz + (1-p) f_y(e_i).$$

Requiring $f_y(e_i) + f_o(e_i)$ to be constant (uniform on $[0, 1]$) gives the ODE

$$f_y'(e_i) + p \frac{f_y(e_i)}{1-e_i} + (1-p) f_y'(e_i) = 0,$$

whose solution with $f_y(0) = 1$ simplifies to $f_y(e_i) = (1 - e_i)^{\frac{p}{2-p}}$ (up to normalization), a strictly decreasing function for $p \in (0, 1)$.

Proposition 3. *A higher level of automation I_0 last period worsens the expertise distribution in a FOSD sense.*

Proof of Proposition 3: Consider a marginal increase $dI_0 > 0$. For $e_i > I_0$, equation (27) gives $\frac{\partial f_o(e_i;I_0)}{\partial I_0} = (q - p)\frac{f_y(I_0)}{1-I_0} < 0$: the marginal worker shifts from the p-integral to the q-integral, which strictly decreases the density of old workers at every expertise level above I_0. For $e_i < I_0$ the density is unchanged, and the lost mass is added to the atom of workers at their initial expertise. Since mass is only relocated downward, the distribution function increases at every point: $F(\cdot; I_0') \geq F(\cdot; I_0)$ for $I_0' > I_0$, a FOSD deterioration. For a discrete increase from I_0 to I_0', the same argument applies at every intermediate level, with the relocated mass added on the interval $(I_0, I_0']$. Combined with the argument in the proof of Proposition 1 that a first-order stochastically dominated expertise distribution lowers equilibrium aggregate output, this delivers the sign of $\frac{\partial \Omega}{\partial I_{t-1}}$ used throughout section 4. $\blacksquare$

References

Acemoglu, Daron. 1998. “Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality.” *Quarterly Journal of Economics* 113: 1055–89.

Acemoglu, Daron. 2002. “Directed Technical Change.” *Review of Economic Studies* 69 (4): 781–809.

Acemoglu, Daron, and David Autor. 2011. “Skills, Tasks and Technologies: Implications for Employment and Earnings.” In *Handbook of Labor Economics*, 4:1043–1171. Elsevier.

Acemoglu, Daron, Dingwen Kong, and Asuman Ozdaglar. 2026. “AI, Human Cognition and Knowledge Collapse.” Working Paper 34910, NBER.

Acemoglu, Daron, Fredric Kong, and Pascual Restrepo. 2024. “Tasks at Work: Comparative Advantage, Technology and Labor Demand.” 32872. NBER.

Acemoglu, Daron, and Todd Lensman. 2024. “Regulating Transformative Technologies.” *American Economic Review: Insights* 6 (3): 359–76.

Acemoglu, Daron, and Pascual Restrepo. 2018. “The Race Between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment.” *American Economic Review* 108 (6): 1488–1542.

Acemoglu, Daron, and Pascual Restrepo. 2019. “Artificial Intelligence, Automation, and Work.” In *The Economics of Artificial Intelligence: An Agenda*, 197–236. University of Chicago Press.

Agrawal, Ajay, Joshua S. Gans, and Avi Goldfarb. 2019. “Artificial Intelligence: The Ambiguous Labor Market Impact of Automating Prediction.” *Journal of Economic Perspectives* 33 (2): 31–50.

Arrow, Kenneth J. 1962. “The Economic Implications of Learning by Doing.” *Review of Economic Studies* 29 (3): 155–73.

Autor, David H., and Michael J. Handel. 2013. “Putting Tasks to the Test: Human Capital, Job Tasks, and Wages.” *Journal of Labor Economics* 31: S59–96.

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. “The Skill Content of Recent Technological Change: An Empirical Exploration.” *Quarterly Journal of Economics* 118: 1279–1333.

Autor, David, and Neil Thompson. 2025. “Expertise.” *Journal of the European Economic Association* 23 (4): 1203–71.

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson. 2024. “Artificial Intelligence, Firm Growth, and Product Innovation.” *Journal of Financial Economics* 151: 103745.

Baccara, Mariagiovanna, SangMok Lee, and Leeat Yariv. 2023. “Task Allocation and On-the-Job

Training." *Journal of Economic Theory* 207: 105587.

Bainbridge, Lisanne. 1983. "Ironies of Automation." *Automatica* 19 (6): 775–79.

Berger, Philip G., Wei Cai, Lin Qiu, and Cindy Xinyi Shen. 2024. "Employer and Employee Responses to Generative AI: Early Evidence." University of Chicago.

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen. 2025. "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence." Stanford University.

Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock. 2018. "What Can Machines Learn and What Does It Mean for Occupations and the Economy?" *AEA Papers and Proceedings* 108: 43–47.

Bulow, Jeremy. 1986. "An Economic Theory of Planned Obsolescence." *Quarterly Journal of Economics* 101 (4): 729–49.

Chen, Mark A., and Joanna Wang. 2024. "Displacement or Augmentation? The Effects of AI Innovation on Workforce Dynamics and Firm Value." Georgia State University.

Coase, Ronald H. 1972. "Durability and Monopoly." *Journal of Law and Economics* 15 (1): 143–49.

Eisfeldt, Andrea L., and Dimitris Papanikolaou. 2013. "Organization Capital and the Cross-section of Expected Returns." *Journal of Finance* 68 (4): 1365–1406.

Eisfeldt, Andrea L., Gregor Schubert, and Miao Ben Zhang. 2023. "Generative AI and Firm Values." 31222. NBER. Forthcoming, *Journal of Finance*.

Garicano, Luis, and Luis Rayo. 2025. "Training in the Age of AI: A Theory of Career Viability." Discussion Paper 20634, CEPR.

Gibbons, Robert, and Michael Waldman. 2004. "Task-Specific Human Capital." *American Economic Review* 94 (2): 203–7.

Grinschgl, Sandra, and Aljoscha C. Neubauer. 2022. "Supporting Cognition with Modern Technology: Distributed Cognition Today and in an AI-Enhanced Future." *Frontiers in Artificial Intelligence* 5: 908261.

Hampole, Menaka, Dimitris Papanikolaou, Lawrence D. W. Schmidt, and Bryan Seegmiller. 2025. "Artificial Intelligence and the Labor Market." 33509. NBER.

Haslbeck, Andreas, and Hans-Juergen Hoermann. 2016. "Flying the Needles: Flight Deck Automation Erodes Fine-Motor Flying Skills Among Airline Pilots." *Human Factors* 58 (4): 533–45.

Heckman, James, and José Scheinkman. 1987. "The Importance of Bundling in a Gorman–Lancaster Model of Earnings." *Review of Economic Studies* 54: 243–55.

Hosseini Maasoum, Seyed Mahdi, and Guy Lichtinger. 2025. "Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Résumé and Job Posting Data." Harvard University.

Ide, Enrique. 2025. “Automation, AI, and the Intergenerational Transmission of Knowledge.” IESE.

Ide, Enrique, and Eduard Talamàs. 2025. “Artificial Intelligence in the Knowledge Economy.” *Journal of Political Economy* 133 (12): 3762–3800.

Jones, Charles I. 2024. “The AI Dilemma: Growth Versus Existential Risk.” *American Economic Review: Insights* 6 (4): 575–90.

Jones, Charles I. 2026. “AI and Our Economic Future.” *Journal of Economic Perspectives* 40 (3): 3–22.

Kosmyna, Nataliya, Eugene Hauptmann, Ye Tong Yuan, Jessica Situ, Xian-Hao Liao, Ashly Vivian Beresnitzky, Iris Braunstein, and Pattie Maes. 2025. “Your Brain on ChatGPT: Accumulation of Cognitive Debt When Using an AI Assistant for Essay Writing Task.” arXiv Preprint arXiv:2506.08872.

Li, Wendy C. Y., and Bronwyn H. Hall. 2020. “Depreciation of Business R&D Capital.” *Review of Income and Wealth* 66 (1): 161–80.

Loewenstein, Mark A., and James R. Spletzer. 1994. “Informal Training: A Review of Existing Data and Some New Evidence.” Bureau of Labor Statistics.

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

Maug, Ernst. 2025. “Financing Human Capital: A Survey and Synthesis.” *Foundations and Trends in Finance* 15 (2): 213–376.

Nishesh, Naman, Paige Ouimet, and Elena Simintzi. 2024. “Labor and Corporate Finance.” In *Handbook of Corporate Finance*, edited by David J. Denis, 647–73. Edward Elgar.

Peters, Ryan H., and Lucian A. Taylor. 2017. “Intangible Capital and the Investment-q Relation.” *Journal of Financial Economics* 123 (2): 251–72.

Smith, Adam. 1776. *An Inquiry into the Nature and Causes of the Wealth of Nations.* London: W. Strahan; T. Cadell.

Zhai, Chunpeng, Santoso Wibowo, and Lily D. Li. 2024. “The Effects of Over-Reliance on AI Dialogue Systems on Students’ Cognitive Abilities: A Systematic Review.” *Smart Learning Environments* 11: 28.