Automation, AI, and the Intergenerational Transmission of Knowledge

Enrique Ide*

September 22, 2025

Abstract

Recent advances in Artificial Intelligence (AI) have sparked expectations of unprecedented economic growth. Yet, by enabling senior workers to accomplish more tasks independently, AI may inadvertently reduce entry-level opportunities, raising concerns about how future generations will acquire essential expertise. This paper develops a model to examine how advanced automation affects the intergenerational transmission of tacit knowledge—practical insights that resist codification and are critical for workplace success. The analysis shows that the competitive equilibrium features socially excessive automation of early-career tasks and reveals a critical trade-off: while such automation delivers immediate productivity gains, it can undermine long-term growth by hindering younger workers' acquisition of tacit skills. Back-of-the-envelope calculations suggest AI-driven entry-level automation could lower the long-run annual growth rate of U.S. per capita output by 0.05 to 0.35 percentage points, depending on its scale. The analysis further shows that AI co-pilots—systems providing access to tacit-like expertise once obtained only through direct experience—can partially offset these losses by assisting individuals who fail to develop adequate skills early in their careers. However, co-pilots are not always beneficial, as they may also weaken junior workers' incentives to engage in hands-on learning. These findings challenge the view that AI will automatically lead to higher economic growth, highlighting the need to safeguard—or deliberately create—entry-level opportunities to fully realize AI's potential.

*Department of Economics, IESE Business School, Carrer de Arnús i de Garí 3-7, 08034 Barcelona, Spain (eide@iese.edu).
I am grateful for the valuable comments and suggestions of David Autor, Heski Bar-Isaac, Matthew Beane, David De la Croix, Andrea Galeotti, Luis Garicano, Chad Jones, Anton Korinek, Ben Moll, Cristóbal Otero, Sebastián Otero, Esteban Rossi-Hansberg, Raffaella Sadun, and Eduard Talamàs. I also acknowledge the financial support of IESE through the High Impact Initiative-course 2024/2025. I declare I have no relevant or material financial interests that relate to the research described in this paper.

1 Introduction

Recent advances in Artificial Intelligence (AI) have fueled optimistic forecasts of dramatic increases in productivity and economic growth. Unlike earlier technological innovations—such as robots or conventional computers—that can only execute predefined instructions, AI systems learn and adapt from examples, enabling them to perform work once considered inherently human (Autor, 2024; Brynjolfsson et al., 2025b; Ide and Talamàs, 2025a). As a result, many technologists anticipate an economic transformation that could rival or even surpass that of the Industrial Revolution (Amodei, 2024; Levy, 2025; Altman, 2025).

However, the widespread adoption of AI may carry unintended long-term costs. By enabling senior workers to accomplish more tasks independently, these technologies may reduce entry-level opportunities and threaten the implicit contract in which younger workers exchange their labor for training and the prospect of future promotion (Roose, *The New York Times*, 2025; The Journal Podcast, *The Wall Street Journal*, 2025; BloombergTV, *Bloomberg.com*, 2025). Such arrangements have historically been crucial for transmitting tacit knowledge—practical insights that resist codification and are acquired primarily through repeated observation of practical successes and failures—raising concerns about how future generations will acquire essential expertise (Beane, 2024a,b; Garicano, 2025).

To analyze these concerns, this paper develops a model to explore how advanced automation can disrupt the intergenerational transmission of tacit knowledge and influence long-run economic growth. Existing research has primarily studied the impact of automation and AI on short-term productivity, direct labor-market outcomes, and income and wealth distribution, but has largely overlooked how these technologies reshape the interpersonal interactions essential for transferring workplace skills. My framework addresses this gap by embedding automation and AI within a growth model of knowledge diffusion and learning (as in, e.g., Lucas, 2009; Lucas and Moll, 2014; De la Croix et al., 2018). This enables me to analyze both automation's immediate impact and its long-run consequences for the skills of future generations.

The baseline model comprises two overlapping cohorts—novices and experts—each living two periods. Novices are born identical and without wealth, whereas experts vary in skill. Production requires expert knowledge and the execution of routine tasks, the latter of which can be performed either by novices or machines. By undertaking these routine tasks, novices interact with randomly matched experts, whose skills they observe and assimilate. This gives rise to an apprenticeship-like arrangement wherein novices accept lower wages in exchange for learning. Novices also independently attempt to innovate by generating random ideas. Thus, when novices transition to expert status, their skill is determined either by the highest skill among the experts they encountered or by the quality of their innovative idea—whichever is greater. Long-run growth arises from both the diffusion of best practices as novices learn from experts and the independent innovations of novices.

The model remains agnostic about whether routine tasks primarily involve codifiable or non-codifiable knowledge. Rather, the critical characteristic of these tasks is that they constitute entry-level activities typically performed by novices with minimal prior experience. Consequently, the

model's insights broadly apply to various automation technologies—including robots—but are especially relevant to AI, given its distinctive capability to automate a wide range of tasks. In contrast, experts' skills embody deeply experiential, tacit knowledge acquired only through extensive practice and direct interaction with more experienced individuals.

The analysis shows that the competitive equilibrium features socially excessive automation of early-career tasks and converges to one of three long-run regimes, depending on the initial stock of tacit knowledge and the costs of machines. In the **Learning Breakdown (LB)**, extensive automation severely restricts novice-expert interactions, gradually eroding tacit knowledge and ultimately resulting in economic stagnation. Under the **Constrained Learning (CL)** regime, moderate levels of automation allow some interactions between novices and experts to persist, sustaining limited yet positive long-term growth. Finally, in the **Full Learning (FL)** regime, high automation costs lead to frequent novice-expert interactions, ensuring robust knowledge transmission and maximal long-run economic growth.[1]

Building on this characterization, I then examine the impact of a one-time (unexpected) improvement in automation, modeled—following Autor et al. (2003)—as a reduction in machine costs. The analysis reveals a critical trade-off: while improvements in automation reduce production costs—boosting the productivity and income of the current generation of experts—these immediate gains may come at the expense of future generations. By limiting novice-expert interactions, increased automation may disrupt the transmission of tacit knowledge, permanently lowering the economy's growth potential. This finding aligns with evidence from sectors such as finance and medicine, where the practical skills of junior professionals have deteriorated following the widespread adoption of automated financial tools and robotic surgical systems (Beane, 2019, 2024a,b).

This trade-off, however, does not extend to all technological advances. In particular, improvements in routine labor-augmenting technologies—defined, following Acemoglu et al. (2024), as those that raise novice productivity without displacing novice labor—can deliver immediate efficiency gains while reinforcing knowledge transmission, thereby supporting long-run growth. Such advances include, for example, better algorithms or machines that help novices avoid mistakes.

Unfortunately, early evidence indicates that firms currently view generative AI primarily as a tool for automation rather than augmentation, especially at the entry level. Brynjolfsson et al. (2025a), for instance, find that employment among early-career workers in the most AI-exposed occupations has fallen by about 13% after the technology's introduction, while employment for older workers in the same roles has held steady or even increased—a pattern also documented by Lichtinger and Maasoum (2025).[2] Similarly, Dario Amodei, CEO of Anthropic, has warned that generative AI could

[1]The conclusion that economic growth is maximized when automation is sufficiently costly should be interpreted with caution. It relies on the assumption that all entry-level tasks provide meaningful learning opportunities. If some tasks are so menial or repetitive that they offer no such value, automating them can raise efficiency without impairing long-term growth. However, unless all automatable tasks lack learning value—a highly implausible case—the competitive equilibrium will still over-automate from a welfare standpoint, reducing economic growth relative to the first-best.

[2]This is also consistent with earlier evidence from Berger et al. (2024) showing that job-postings for entry-level positions in AI-exposed occupations fell sharply after the release of ChatGPT in 2023.

soon eliminate up to half of all entry-level white-collar jobs (VandeHei and Allen, 2025). Reinforcing this outlook, a recent survey by The World Economic Forum (2025) reports that firms expect the share of tasks performed solely by humans to fall from 47% to 33% by 2030, while fully automated tasks are projected to rise from 22% to 34%. Meanwhile, tasks performed collaboratively by humans and machines are expected to increase only modestly—from 30% to 33%.

Motivated by concerns that AI-driven entry-level automation could disrupt knowledge transmission and growth, I then perform simple back-of-the-envelope calculations to gauge the potential magnitude of these effects. This exercise is neither a formal calibration nor a forecast. It deliberately abstracts from compensating factors, such as AI's potential to stimulate innovation. Instead, its goal is to illustrate how recent advances in AI might undermine long-run economic growth by weakening the intergenerational transmission of tacit knowledge.

The theoretical model yields a simple expression that captures how automating entry-level tasks reduces the economy's steady-state growth. Specifically, the reduction in the long-run annual growth rate of per capita output, denoted by Δg^Y and expressed in percentage points per year, is given by:

$$\Delta g^Y = 100 \times (1 + g^Y)\big[1 - (1 - ax^{1/\theta})^{\theta/T}\big],$$

where g^Y is the economy's anual steady-state growth rate of per capita output before the automation shock, a represents the fraction of entry-level tasks automated due to the shock, x denotes the share of steady-state growth driven by the diffussion of tacit knowledge ($1 - x$ is the share of growth due to new ideas), θ captures the dispersion of tacit knowledge across experts, and T corresponds to the time (in years) individuals spend at each career stage—first as novices, then as experts.

I set the baseline growth rate g^Y to 2%, reflecting the average per-capita growth rate in the U.S. over the past 150 years (Jones, 2016). I choose $T = 20$, implying that individuals spend 20 years as novices followed by another 20 years as experts. Following Lucas (2009), Lucas and Moll (2014), and De la Croix et al. (2018), I adopt $\theta = 0.5$ as my baseline specification. Additionally, guided by empirical estimates from the international trade and knowledge diffusion literature (Eaton and Kortum, 1999; Santacreu, 2015), I assume that about two-thirds ($x = 65\%$) of steady-state growth is driven by the transmission of tacit knowledge.

Finally, informed by recent forecasts assessing the potential one-time productivity gains from AI adoption, I consider two illustrative scenarios for the fraction a of entry-level tasks that are automated. The first is a conservative scenario ($a = 5\%$) inspired by Acemoglu's (2024) cautious projection of modest productivity improvements of approximately 0.71% over ten years. The second is an aggressive scenario ($a = 30\%$), reflecting Aghion and Bunel's (2024) more optimistic forecasts of more extensive AI use and substantial productivity gains of around 7% over the same period. Although both studies remain agnostic about whether their productivity projections primarily reflect automation or augmentation, here I explicitly assume that AI adoption exclusively takes the form of automation, consistent with the early evidence presented above.

The results indicate that automating 5% of entry-level jobs lowers the long-run annual growth rate by about 0.05 percentage points. When combined with Acemoglu's (2024) projection of a one-time

0.71% productivity boost from adoption, this modest gain vanishes within 15 years, leaving output 2.72% below the no-AI baseline by 2100. In a more aggressive scenario, where 30% of entry-level jobs are automated, the growth rate falls by about 0.35 percentage points. Paired with Aghion and Bunel's (2024) projection of a 7% one-time gain, the initial boost is exhausted within 20 years, and output ultimately ends up 14.15% below baseline by the end of the century.[3]

Thus far, the analysis has focused on the negative implications of entry-level automation. However, automation represents only one dimension of AI's broader capabilities. AI-powered decision-support systems—commonly known as AI co-pilots—offer scalable access to specialized expertise previously attainable only through extensive hands-on experience (Autor, 2024; Ide and Talamàs, 2025a). These systems thus have the potential to offset some of the adverse effects of entry-level automation by democratizing access to expert-level, tacit-like knowledge.

As a final step, I analyze the impact of AI co-pilots on the intergenerational transmission of tacit knowledge, building on the framework of Ide and Talamàs (2025a). I model co-pilots as AI systems that experts can consult to augment their skills. The analysis shows that these systems can indeed partially offset the negative effects of entry-level automation by assisting experts who, as novices, failed to acquire adequate skills. Yet co-pilots are not unambiguously beneficial: their availability may also weaken novices' incentives to engage in hands-on learning. Thus, whereas entry-level automation reduces the supply of apprenticeship-like contracts offered by experts, AI co-pilots diminish novices' demand for these arrangements. Both mechanisms undermine interpersonal knowledge transfer, potentially constraining long-term economic growth.

In sum, this paper challenges the optimistic view that AI will automatically lead to higher economic growth, emphasizing instead that fully realizing AI's economic promise depends crucially on preserving or creating roles that facilitate novice-expert interactions.[4] Policymakers, therefore, face a delicate balancing act: promoting AI adoption while safeguarding entry-level opportunities. Effective policy responses might include targeted subsidies for mentorship and training programs, taxes on entry-level automation, or explicit incentives for the development of AI systems designed to complement rather than replace human labor. Without thoughtful intervention, the long-term economic benefits of AI risk being substantially diminished by a silent erosion of tacit-like expertise.

Related Literature

This paper bridges two strands of literature. First, it extends the literature on automation and AI by analyzing how these technologies influence the intergenerational transmission of tacit knowledge,

[3] I also consider a case where entry-level displacement is temporary (lasting only twenty years), as firms eventually introduce new junior roles. In this scenario, the long-run growth rate is preserved. However, the economy is left with a permanently smaller stock of tacit knowledge, which erases the one-time productivity gains from AI adoption in both the $a = 5\%$ and $a = 30\%$ scenarios.

[4] A more radical possibility is that AI becomes sufficiently transformative to render existing forms of tacit knowledge obsolete, fundamentally reshaping the skills required by future generations. In Section 7, I further elaborate on what this "new tacit knowledge" might be, and why it remains unlikely that current tacit knowledge will become entirely irrelevant.

with significant implications for long-term economic growth. Second, it contributes to the literature on idea flows and economic growth by explicitly modeling how automation and AI affect knowledge diffusion across generations.

The literature on automation has largely adopted a task-based framework, pioneered by Zeira (1998), Autor et al. (2003), and Acemoglu and Autor (2011). In this approach, automation technologies do not replace entire occupations but rather specific tasks within jobs, particularly those that are routine and easily codifiable. Building on this framework, Acemoglu and Restrepo (2018, 2019, 2022), and Acemoglu et al. (2024) examine how technological innovations—such as automation and labor-augmenting technologies—displace or complement human labor in existing tasks, generate new ones, and reshape employment patterns, wages, inequality, and productivity. Complementary research by Acemoglu (2024) and Aghion and Bunel (2024) applies this task-based approach to forecast potential one-time productivity increases from AI adoption. Finally, Autor and Thompson (2025) extend the framework by treating occupations as bundles of tasks, uncovering implications that traditional task-based models missed.[5]

More recently, Ide and Talamàs (2024, 2025a,b) build on the literature on knowledge hierarchies developed by Garicano (2000) and Garicano and Rossi-Hansberg (2004, 2006) to study AI's distinctive implications. As highlighted earlier, AI differs from earlier automation technologies in its ability to acquire and deploy non-codifiable knowledge. The knowledge-hierarchies framework is a specialized version of the task-based approach, placing non-codifiable knowledge at its core and where the links between tasks emerge endogenously from organizational considerations (Garicano and Rossi-Hansberg, 2015; Ide and Talamàs, 2025a). Leveraging this perspective, Ide and Talamàs (2024, 2025a,b) analyze AI's implications for occupational choices, organizational structures, labor income, and international trade patterns. Their findings reveal mechanisms absent from conventional task-based analyses, emphasizing the critical roles of AI autonomy, communication, and endogenous organizational responses.

I contribute to this literature by examining how automation and AI influence the intergenerational transmission of knowledge, thereby shaping the distribution of human skills across generations. I do so by embedding automation into an overlapping-generations framework, where the skill levels of successive cohorts critically depend on the extent of interactions between novices and experts. My baseline model—which focuses on entry-level automation—builds upon the traditional task-based approach, remaining agnostic as to whether routine tasks involve codifiable or tacit knowledge. Consequently, the model's insights apply broadly to various forms of automation but hold particular relevance for AI, given its unique capability to automate an increasingly wide range of tasks. In contrast, my extension analyzing AI co-pilots draws inspiration from Ide and Talamàs (2025a), particularly their treatment of these systems.

Kosmyna et al. (2025) also investigate how AI can affect human skills. In a recent pre-print, they present neuroscientific evidence that individuals who consistently rely on AI for essay writing ex-

[5]Other important contributions using the task-based approach include Aghion et al. (2017), Moll et al. (2022), Korinek and Suh (2024), and Acemoglu and Loebbing (2024).

hibit reduced neural connectivity and cognitive engagement compared with those completing the task unaided. While complementary, their findings differ fundamentally from the mechanism highlighted in this paper. My analysis emphasizes how automation indirectly reshapes skill distribution by disrupting intergenerational knowledge transmission through diminished novice-expert interactions. In contrast, Kosmyna et al. (2025) highlight a direct channel, showing that frequent personal AI use impairs cognitive skill development at the individual level.

The literature on idea flows and economic growth emphasizes how knowledge diffuses through interactions among individuals and firms, shaping productivity and growth. Kortum (1997) and Eaton and Kortum (1999, 2002) introduce extreme value theory for modeling the diffusion of ideas across firms and economies. Building on this foundation, Lucas (2009) develops a model where growth arises from continuous exchanges of ideas among heterogeneous agents. Lucas and Moll (2014) extend this framework by explicitly modeling individuals' optimal allocation of time between producing output and learning from others. Caicedo et al. (2019) explore how knowledge transmission within hierarchical organizations shapes career trajectories as workers advance by learning from coworkers.[6] My paper contributes to this literature by examining how automation and AI affect these interpersonal interactions—a crucial channel overlooked by existing models.

In the context of idea flows and economic growth, the paper closest to mine is De la Croix et al. (2018). They also develop an overlapping-generations model in which novices learn from experts, emphasizing the critical role European apprenticeship institutions played in fostering innovation and economic growth leading up to the Industrial Revolution. However, my analysis differs from theirs along three key dimensions.

First, whereas De la Croix et al. (2018) emphasize how medieval institutions—such as clans and guilds—helped mitigate experts' incentive to shirk on training effort, I abstract from this conflict to focus on how automation and AI reshape these training arrangements. Second, their model reflects medieval economic conditions, explicitly featuring population growth and Malthusian constraints. By contrast, my framework aligns more closely with Lucas (2009) in abstracting from these features to better capture contemporary economies. Third, while De la Croix et al. (2018) allow novices to pay upfront for apprenticeships and assume parental altruism to secure intergenerational knowledge transmission, my model assumes wealth-constrained novices and no altruism. Taken together, these differences yield fundamentally different long-run outcomes and comparative statics compared to those in De la Croix et al. (2018).

Finally, my paper also contributes to the literature on the economics of apprenticeships. As argued by Mokyr (2019), apprenticeship-like arrangements suffer from contractual incompleteness due to the tacit nature of the knowledge involved. These contractual frictions might lead to inefficient knowledge transmission (Garicano and Rayo, 2017; De la Croix et al., 2018; Fudenberg and Rayo, 2019).

I contribute to this body of work by examining the impact of automation and AI—an aspect

[6]Other important contributions include Perla and Tonetti (2014), Buera and Oberfield (2020), Jarosch et al. (2021). Perla et al. (2021), and Benhabib et al. (2021). For a survey on the literature, see Buera and Lucas (2018).

that remains largely unexplored. The only exception is the complementary work by Garicano and Rayo (2025), who extend the framework developed by Garicano and Rayo (2017) to analyze how AI alters the viability of apprenticeships. Their analysis, however, focuses on the dynamics of a single expert–novice relationship. By contrast, my paper examines how widespread disruptions to apprenticeship-like arrangements affect the intergenerational transmission of knowledge, with lasting consequences for economic growth.

1.1 Roadmap

The rest of the paper is organized as follows. Section 2 outlines four motivating ideas underpinning the analysis. Section 3 introduces the baseline model, which focuses on entry-level automation. Section 4 characterizes the long-run equilibrium of the baseline model and examines the implications of a one-time improvement in automation. Section 5 provides illustrative back-of-the-envelope calculations assessing the degree to which AI-driven entry-level automation might constrain output and growth. Section 6 extends the analysis by examining the impact of AI co-pilots. Section 7 concludes.

2 Four Motivating Ideas

This section documents four key motivating ideas that inform the subsequent analysis. First, advanced automation technologies, especially AI, have fueled optimistic forecasts of dramatic increases in productivity. Second, the most valuable workplace skills remain largely tacit, typically transmitted from experts to novices through personal interactions. Third, automation threatens the expert-novice relationship, potentially weakening the transmission of tacit knowledge. Fourth, despite such disruption, AI also opens new pathways to democratize tacit-like expertise through scalable decision-support systems (commonly referred to as AI co-pilots).

Automation, AI, and the Promise of Growth.—As highlighted in the Introduction, recent breakthroughs in AI have dramatically expanded the scope of automation beyond repetitive, clearly defined tasks. Modern AI technologies can now effectively perform sophisticated work—such as complex problem-solving, original research, and creative tasks—that have traditionally been considered the exclusive domain of humans (Autor, 2024; Brynjolfsson et al., 2025b; Ide and Talamàs, 2025a). Such expanded capabilities have driven optimistic forecasts of substantial increases in productivity and economic growth (Amodei, 2024; Levy, 2025; Altman, 2025).

AI has the potential to enhance productivity through two distinct channels (Aghion and Bunel, 2024). First, it can directly do so by reducing production costs. Some estimates suggest that this channel could lead to one-off productivity gains realized over ten years, equivalent to between 0.68 and 3.4 percentage points per year throughout this period (Goldman Sachs, 2023; McKinsey & Company, 2023; Aghion and Bunel, 2024).[7] Second, AI's ability to automate the production of ideas could stim-

[7]By contrast, Acemoglu (2024) is more skeptical, forecasting that AI adoption will yield a modest one-time TFP gain of

ulate innovation, potentially generating a permanent increase in productivity growth (Jones, 2024). However, despite widespread speculation, reliable estimates quantifying this last channel remain unavailable.

Tacit Knowledge, Entry-Level Jobs, and the Role of Apprenticeships.—Codifiable knowledge consists of explicit rules and procedures that can be readily communicated through manuals, books, and databases. By contrast, non-codifiable (or tacit) knowledge refers to intuitive skills and insights that individuals possess but find difficult to clearly articulate (Polanyi, 1966; Foray, 2004).

Although recent technological advances, such as the Information and Communication Technology (ICT) revolution, have enhanced the codifiability of knowledge (Cowan and Foray, 1997; Foray, 2004), the most valuable workplace skills remain predominantly tacit (MacKenzie and Spinardi, 1995; Beane, 2024b). Formal education, which focuses on codifiable knowledge, provides basic qualifications—what Beane (2024b, p. 4) calls "table stakes"—but true competence emerges only when workers integrate foundational knowledge with deeper insights gained through practice and hands-on collaboration with more experienced individuals.

This explains why apprenticeships have historically been crucial for skill transmission—and remain essential today. Though formal apprenticeships have declined, informal apprenticeships remain prevalent across modern economies. Medical training, particularly in surgery, exemplifies this through the principle "see one, do one, teach one" (Beane, 2024b, p. 3). Similar arrangements characterize professions such as law, consulting, and finance, where junior associates perform routine tasks under the supervision of senior colleagues, exchanging their labor for training, experience, and eventual promotion (Garicano and Rayo, 2017). These exchanges naturally arise because novices typically lack the financial resources to directly compensate experts for training, and the intangible, tacit nature of the knowledge they seek prevents its use as collateral (Becker, 1964).

Empirical evidence highlights the critical role of initial job placements in shaping workers' careers through their effects on early human capital accumulation. Oyer (2006) provides causal evidence that academic economists placed at higher-ranked institutions due to favorable labor market conditions exhibit persistently greater research productivity throughout their careers. Complementing this result, Kahn (2010) documents lasting wage penalties among college graduates entering the labor market during recessions, implicitly linking these persistent losses to diminished investments in early-career skills. Similarly, Arellano-Bover (2024) explicitly attributes the higher lifetime earnings of workers initially employed at larger firms to the superior training and developmental opportunities available at such firms.

Given their essential role in driving human capital accumulation, apprenticeships have historically been pivotal in fostering innovation and economic growth. Humphries (2006), for example, identifies England's apprenticeship system as a "neglected factor" underpinning the British Industrial Revolution. Similarly, De la Croix et al. (2018) argue that European apprenticeship institutions

just 0.71% over ten years—equivalent to approximately 0.07 percentage points per year. He therefore concludes that the technology's impact will fall substantially short of prevailing expectations.

significantly contributed to the continent's relative rise in technological creativity, population growth, and per-capita income in the centuries preceding the Industrial Revolution.

Advanced Automation May be Disrupting the Intergenerational Transmission of Knowledge.— Because non-codifiable knowledge spreads primarily through direct experience, its preservation depends crucially on the continual renewal of knowledgeable individuals across generations (Foray, 2004).[8] Advanced automation technologies may disrupt this renewal by enabling experts to operate independently, reducing juniors' opportunities for hands-on learning (Beane, 2024b).

Both anecdotal accounts and empirical studies document this disruption in investment banking and urological surgery (Beane, 2019, 2024a,b; Beane and Anthony, 2024). Junior analysts in investment banking are increasingly distanced from senior bankers, who now rely on AI for tasks that were previously delegated to juniors. Similarly, surgical residents are being sidelined from hands-on participation in complex procedures, as attending surgeons are turning instead to robotic systems. This lack of practical experience can significantly compromise juniors' productivity and performance once they advance to more senior roles. Beane (2024b, p. 9) vividly captures this concern through a conversation with a chief of surgery, who highlights that most new surgeons lack critical skills:

> "I mean these guys can't do it. They haven't had any experience doing it."

Recent advances in Generative AI may further exacerbate these disruptions by directly reducing entry-level job opportunities. Berger et al. (2024) exploit the unexpected release and rapid adoption of ChatGPT in 2023 as a natural experiment, showing that occupations more exposed to generative AI—measured using Eloundou et al.'s (2024) index—experienced an 18% decline in weekly job postings for entry-level white-collar positions. Consistent with this, Brynjolfsson et al. (2025a) and Lichtinger and Maasoum (2025) document significant employment losses among younger workers in the most AI-exposed occupations, even as employment for older workers in the same roles remained stable or increased.[9] Echoing these concerns, Dario Amodei, CEO of Anthropic, has warned that generative AI could ultimately eliminate up to half of all entry-level white-collar jobs (VandeHei and Allen, 2025).

AI Co-Pilots and Democratized Expertise.—While advanced automation may threaten intergenerational knowledge transmission by reducing direct interactions between experts and novices, AI simultaneously creates new opportunities to democratize expertise through decision-support systems (Autor, 2024). These systems—commonly known as AI co-pilots—could partially offset the erosion of tacit knowledge by enhancing the capabilities of less experienced individuals.

[8] MacKenzie and Spinardi (1995) vividly illustrates the fragility of tacit knowledge, documenting that—contrary to common belief—nuclear weapons can indeed be "disinvented." Historical accounts from British and Soviet experiences reveal that losing just one generation of skilled engineers can erase essential insights and procedures necessary for developing nuclear weapons, making subsequent redevelopment closer to reinvention than simple replication.

[9] Some studies report more optimistic findings. Johnston and Makridis (2025) document employment and wage gains for younger workers following the widespread adoption of AI, though Brynjolfsson et al. (2025a) argue these results may reflect insufficiently granular data. de Souza (2025) finds mixed effects in Brazil: AI appears to automate low-wage administrative jobs while augmenting younger and less-experienced workers in production roles.

The potential of AI co-pilots to democratize expertise arises from their ability to acquire tacit-like knowledge and scale its application through computational resources. This enables a broad range of individuals with the necessary foundational skills to perform more sophisticated tasks without directly acquiring the underlying expertise (Autor, 2024). Yet modern AI systems—especially Large Language Models (LLMs)—are opaque (Dell'Acqua et al., 2023): their embedded knowledge is encoded in complex internal structures that users can apply but cannot readily interpret or fully understand. Consequently, this knowledge exhibits a hybrid nature, combining the scalability typical of codifiable knowledge with the opacity and experiential features of non-codifiable knowledge (Ide and Talamàs, 2025a).

Recent experimental evidence supports the democratizing potential of AI co-pilots (Noy and Zhang, 2023; Peng et al., 2023; Brynjolfsson et al., 2025b). For instance, Brynjolfsson et al. (2025b) find that introducing an AI-based conversational assistant increased productivity by 30% among lower-skilled, less-experienced customer support agents, doubling the 15% average gain. Similarly, Noy and Zhang (2023) show that AI assistance reduced the correlation in individuals' performance across successive writing tasks from 0.41 in the control group (without AI) to 0.14 in the treatment group (with AI), highlighting AI's potential to narrow performance gaps.

3 The Baseline Model: Automation of Entry-Level Work

Motivated by concerns that advanced automation may prevent future generations from acquiring essential expertise, this section develops a formal model to examine how automating entry-level tasks may disrupt the transmission of tacit knowledge and affect long-run economic growth. While the framework applies broadly to automation, it is particularly relevant to AI, given its distinctive capability to automate a wide range of tasks. Later, I extend the baseline model to incorporate AI co-pilots, capturing AI's unique potential to democratize access to tacit-like expertise.

The model developed here is deliberately simple, built on several simplifying assumptions that ensure tractability and make the core mechanisms transparent. In what follows, Section 3.1 presents the formal mathematical structure, while Section 3.2 provides its interpretation and discusses the rationale and implications of these assumptions.

3.1 Formal Description

The Environment.— Time is discrete and infinite, indexed by $t = 0, 1, 2, \ldots$. The economy is populated by overlapping generations. At each point in time, two cohorts coexist: young agents ("novices") and old agents ("experts"). Each individual lives for two periods: first as a novice, then as an expert, before retiring. Both cohorts have unit mass, so the total population remains constant at two.

Novices at time t are born identical and without wealth. By contrast, experts in the same period are heterogeneous in skill. Specifically, expert i at time t has a skill level $q_{i,t}$, drawn from a Fréchet distribution with scale parameter k_t, shape parameter $1/\theta > 1$, and location parameter $q^{\min} > 0$. The

cumulative distribution function (CDF) of $q_{i,t}$ is, therefore:

$$G_t(q_{i,t}) = \exp\left\{-\left(\frac{q_{i,t} - q^{\min}}{k_t}\right)^{-1/\theta}\right\}$$

so $\bar{q}_t \equiv \mathbb{E}_t[q_{i,t}] = q^{\min} + k_t\Gamma(1-\theta)$, where $\Gamma(s) \equiv \int_0^\infty x^{s-1}e^{-x}dx$ is the Euler's gamma function.

I refer to k_t as society's stock of *tacit knowledge* at time t. A higher k_t raises the average expert skill level, shifting the entire distribution of skills upward. This knowledge stock evolves endogenously as novices gain experience interacting with experts in the preceding period and attempt to innovate on their own (as explained below). The initial stock of knowledge $k_0 > 0$ is exogenously given.

Production, Technology, and Labor Demand.— There is a single final good, which serves as the numeraire. Its production is linear in expert skill but requires the completion of a continuum of routine tasks with total measure $N > 1$. Formally, the output produced by an expert with skill $q_{i,t}$ at time t is given by:

$$Y_{i,t} = q_{i,t} \cdot \mathbb{1}\{\text{a measure } N \text{ of routine tasks are executed}\} \tag{1}$$

Routine tasks can be performed either by novices or machines. For simplicity, experts do not supply routine labor themselves. Experts incur labor costs when hiring novices and rental costs when using machines. Additionally, delegating tasks to novices generates supervision costs, reflecting the difficulties experts face in identifying and correcting mistakes made by novices. These supervision costs increase as novices are assigned more tasks, reflecting the greater complexity involved in effectively overseeing novices who take on greater responsibilities.

Formally, the total cost faced by expert i at time t, when assigning a measure $L_{i,t} \in [0, N]$ of tasks to novices (and thus, the remaining measure $N - L_{i,t}$ to machines), is:

$$C(L_{i,t}) = w_t L_{i,t} + r_t(N - L_{i,t}) + \frac{cL_{i,t}^2}{2}$$

where w_t denotes the novice wage, r_t is the rental rate of machines, and $c > 0$ governs the magnitude of supervision costs. I assume that novices cannot identify or select experts based on skill ex ante, explaining why all experts pay the same wages.

Moreover, following Acemoglu et al. (2024), I assume that each machine is produced instantaneously using $\rho > 0$ units of the final good and fully depreciates after one period. As a result, the rental rate of machines remains constant at $r_t = \rho$ for all $t \geq 0$. This simplifying assumption allows me to focus on the dynamics of knowledge accumulation rather than physical capital accumulation.

Experts seek to maximize their income during the final period of their economic lives. Therefore, expert i's optimization problem at time t is:

$$\max_{L_{i,t} \in [0,N]} I_{i,t} = q_{i,t} - C(L_{i,t}).$$

I assume $q^{\min} \geq \rho N$, ensuring that production is profitable even for experts possessing the lowest possible skill level $q^{\min}$.

The solution to this problem depends solely on current prices and parameters—not on individual skill or time. Thus, I omit subscripts and denote the common solution as $L^d(w_t, \rho)$, given explicitly by:

$$L^d(w_t, \rho) = \begin{cases} 0, & \text{if } \rho \leq w_t, \\ \frac{\rho - w_t}{c}, & \text{if } \rho - cN < w_t < \rho, \\ N, & \text{if } w_t \leq \rho - cN. \end{cases} \tag{2}$$

Intuitively, when novice wages exceed the rental cost of machines, experts assign all routine tasks to machines. Conversely, when novice wages are sufficiently low, all experts delegate all tasks to novices. Between these extremes, experts split tasks between novices and machines to optimally balance supervision, labor, and machine costs.

Finally, since all experts find it profitable to produce and all of them—regardless of individual skill—use the same amount of novices and machines, $L^d(w_t, \rho)$ also constitutes the aggregate demand for novice labor.

Novices' Labor Supply and Skill Acquisition.— A novice chooses her labor supply l_t to maximize expected lifetime income:

$$\max_{l_t \in \mathbb{R}_+} \left\{ w_t l_t - \frac{\mu}{2} l_t^2 + \beta \, \mathbb{E}_t[I^*_{i,t+1} \mid l_t] \right\}$$

where $\mu > 0$ captures the disutility of labor, $\beta \in (0, 1)$ is the discount factor, and $\mathbb{E}_t[I^*_{i,t+1} \mid l_t]$ denotes the novice's expected income as an expert in the next period. This expectation depends on the novice's current labor effort, as each task allows her to interact with a randomly matched expert, whose skill she observes and assimilates (i.e., supplying l_t units of labor exposes a novice to the skills of l_t different experts). The random matching reflects the assumption, introduced earlier, that novices cannot distinguish experts by skill ex ante.

Moreover, after this learning period—but before becoming an expert—the novice independently attempts innovation by generating a new idea of random quality. The quality of this new idea is independently distributed Fréchet:

$$\iota_{i,t} \sim \text{Fréchet}(\nu^\theta k_t, 1/\theta, q^{\min})$$

where $\nu \geq 0$ captures society's capacity for innovation.[10] The average quality of this idea depends on society's current stock of tacit knowledge k_t, reflecting that new ideas typically build upon existing insights and accumulated best practices (De la Croix et al., 2018; Buera and Oberfield, 2020).

The novice's skill upon becoming an expert in the next period is thus determined by the highest skill among experts she meets and her own generated idea:

$$q_{i,t+1} = \max\{q_{1,t}, \ldots, q_{l_t,t}, \iota_{i,t}\}$$

Since each expert requires the same amount of novice labor, novice-expert interactions are unbiased, with each encounter representing an independent draw from the current distribution of expert skills.

[10] All theoretical results remain valid with $\nu = 0$. I introduce this parameter only because it is relevant for the quantitative illustrations of Section 5.

Hence, $q_{1,t}, \ldots, q_{l_t,t} \sim_{\text{i.i.d.}} G_t$. Additionally, since the quality of the novice's own idea is also independently distributed, her future skill level as an expert—conditional on her labor supply l_t—is distributed according to:

$$\mathbb{P}(q_{i,t+1} \leq q \mid l_t) = [G_t(q)]^{l_t} \times F_t(q) \Longrightarrow q_{i,t+1} \mid l_t \sim \text{Fréchet}(k_{t+1} = (l_t + \nu)^{\theta} k_t, 1/\theta, q^{\min})$$

where F_t denote the CDF of idea quality at time t.

To accommodate a continuous labor supply $l_t \in \mathbb{R}_+$, I interpret $q_{i,t+1} = \max\{q_{1,t}, \ldots, q_{l_t,t}, \iota_{i,t}\}$ using fractional-order statistics, which generalizes traditional order statistics to non-integer sample sizes (Stigler, 1977). This is a standard assumption in the literature of idea flows and economic growth (e.g., Lucas, 2009; Lucas and Moll, 2014; De la Croix et al., 2018).

Given that future income is linear in skill, i.e., $I^*_{i,t+1} = q_{i,t+1} - C(\cdot)$, and $\mathbb{E}_t[q_{i,t+1} \mid l_t] = q^{\min} + (l_t + \nu)^{\theta} k_t \Gamma(1-\theta)$, the problem faced by a novice at time t can be written as:

$$\max_{l_t \in \mathbb{R}_+} \left\{ w_t l_t - \frac{\mu}{2} l_t^2 + \beta \Big[q^{\min} + (l_t + \nu)^{\theta} k_t \Gamma(1-\theta) - C(\cdot) \Big] \right\} \tag{3}$$

where $C(\cdot)$ does not depend on l_t, since experts use the same amount of novices and machines—and pay the same wages—regardless of their skill.

As evident from (3), novices may accept lower wages if compensated through valuable learning opportunities: For a given wage w_t, the marginal benefit of supplying an additional unit of labor increases with the current stock of knowledge k_t. In a sense, novices "purchase" knowledge through their labor, giving rise to an apprenticeship-like arrangement. However, as novices are born without wealth, wages cannot fall below zero; in other words, novices cannot pay upfront for training.

Given that the objective in (3) is strictly concave in l_t, it has a unique solution which, moreover, depends only on the current wage w_t and knowledge stock k_t. Furthermore, since novices are identical and each cohort has unit mass, the aggregate labor supply at time t, denoted by $L^s(w_t, k_t)$, coincides with the optimal labor supply of an individual novice. That is, $L^s(w_t, k_t)$ is also given by the solution to (3).

Competitive Equilibrium.— Due to novices' incentives to learn from experts, the labor market may fail to clear at non-negative wages due to an excess supply of novice labor. In such cases, I assume proportional rationing: each novice in equilibrium supplies a fixed fraction of their desired labor, determined by the aggregate imbalance between supply and demand. This approach ensures that all novices remain exposed to a representative cross-section of experts, preserving the Fréchet distribution of expert skills over time.

I can now define a competitive equilibrium. For this definition, let $\omega(k_t)$ be the (possibly negative) wage that clears the labor market if the stock of knowledge is k_t; that is, $L^d(\omega(k_t), \rho) = L^s(\omega(k_t), k_t)$. In Appendix A, I show that for any $k_t > 0$, there exists a unique $\omega(k_t)$.

Definition (Competitive Equilibrium). A competitive equilibrium is a sequence of wages and knowledge levels, $\{w_t^*, k_t^*\}_{t \geq 0}$, along with a sequence of novice labor, $\{L_t^*\}_{t \geq 0}$, such that for all $t \geq 0$:

- Given (w_t^*, k_t^*), experts and novices behave optimally at time $t \geq 0$, i.e.,

$$L^d(w_t^*, \rho) \text{ is given by (2) and } L^s(w_t^*, k_t^*) \text{ by the solution to (3).}$$

- Markets clear (possibly via proportional rationing):
  - If $\omega(k_t^*) \geq 0$, then $w_t^* = \omega(k_t^*)$ and $L_t^* = L^d(w_t^*, \rho) = L^s(w_t^*, k_t^*)$.
  - If $\omega(k_t^*) < 0$, then $w_t^* = 0$ and $L_t^* = L^d(0, \rho) = \lambda_t^* L^s(0, k_t^*)$, where $\lambda_t^* = L^d(0, \rho)/L^s(0, k_t^*)$.
- Knowledge evolves according to $k_{t+1}^* = (L_t^* + \nu)^\theta k_t^*$, with $k_0^* = k_0 > 0$.

As shown in Appendix A, for any initial stock of knowledge $k_0 > 0$, a competitive equilibrium always exists and is unique. Note also that the law of motion of knowledge implies that:

$$k_{t+1}^* > k_t^* \text{ iff } L_t^* > 1 - \nu, \quad k_{t+1}^* = k_t^* \text{ iff } L_t^* = 1 - \nu, \quad \text{and} \quad k_{t+1}^* < k_t^* \text{ iff } L_t^* < 1 - \nu$$

Some Additional Definitions.— In this context, equilibrium aggregate output at time t equals the average skill of experts:

$$Y_t^* = \mathbb{E}_t[Y_{i,t}^*] = \bar{q}_t = q^{\min} + k_t^* \Gamma(1 - \theta)$$

Experts' aggregate income at time t, in turn, is defined as aggregate output minus production costs (novice wages, machine rental, and supervision costs):

$$I_t^* = \mathbb{E}_t[I_{i,t}^*] = q^{\min} + k_t^* \Gamma(1 - \theta) - \left[w_t^* L_t^* + \rho(N - L_t^*) + \frac{c}{2}(L_t^*)^2 \right]$$

Society's per-period welfare at time t is then defined as total income earned by experts and novices, net of machine costs, supervision costs, and the disutility associated with novice labor effort. Formally, $\mathcal{W}_t^* = \mathcal{W}(L_t^*, k_t^*; \rho)$, where:

$$\mathcal{W}(L_t, k_t; \rho) \equiv q^{\min} + k_t \Gamma(1 - \theta) - \rho(N - L_t) - \left(\frac{c + \mu}{2} \right) L_t^2$$

Finally, the equilibrium growth rate of the stock of knowledge at time t is given by $g_t^k = (k_{t+1}^* - k_t^*)/k_t^* = (L_t^* + \nu)^\theta - 1$, while the growth rate of output is:

$$g_t^Y \equiv \frac{Y_{t+1}^* - Y_t^*}{Y_t^*} = \frac{k_{t+1}^* - k_t^*}{q^{\min}/\Gamma(1 - \theta) + k_t^*} = \frac{(L_t^* + \nu)^\theta - 1}{1 + \frac{q^{\min}}{k_t^* \Gamma(1 - \theta)}}.$$

This implies that when knowledge growth stalls ($g_t^k = 0$), output growth also ceases ($g_t^Y = 0$). Conversely, as knowledge accumulates ($k_t^* \to \infty$), the influence of the lower bound $q^{\min}$ becomes negligible, and output growth converges precisely to knowledge growth:

$$\lim_{k_t^* \to \infty} g_t^Y = g_t^k = (L_t^* + \nu)^\theta - 1.$$

3.2 Discussion and Interpretation

Before proceeding, I discuss the interpretation of the model as well as the motivation behind several simplifying assumptions.

On the Nature of Tasks and Roles.— The model remains agnostic on whether routine tasks involve codifiable or non-codifiable knowledge, making it broadly applicable to various forms of automation. The critical feature of these tasks is that they represent entry-level work traditionally assigned to novices, requiring minimal prior hands-on experience.

For an example involving codifiable knowledge, consider robotic technologies such as the da Vinci Surgical System. These systems automate repetitive, well-defined surgical tasks—such as controlled incisions, suturing, and tissue manipulation—that were traditionally performed by resident surgeons. For an example involving non-codifiable knowledge, consider entry-level tasks in law. AI systems are increasingly automating activities such as document review and initial contract drafting, which, despite requiring nuanced judgment, remain sufficiently bounded in complexity and were previously handled by junior associates.

In contrast, expert skills embody deeply experiential, non-codifiable knowledge. For instance, attending surgeons must manage unforeseen complications and anatomical variations, while senior lawyers handle complex client relationships, settlement negotiations, and sophisticated legal analyses. Such tasks inherently require extensive practical experience.

Finally, note that in the baseline model, machines can only perform routine, entry-level tasks and cannot replicate experts' deeply experiential, tacit knowledge. In Section 6, I extend the model by introducing AI co-pilots, which are systems that experts can consult to augment their skills. Moreover, in Section 7, I also discuss the more radical possibility that AI might render existing forms of tacit knowledge obsolete, fundamentally reshaping the skills required by future generations.

On Supervision Costs.— The baseline model assumes that experts incur supervision costs only when assigning tasks to novices, not when employing machines. This assumption is not essential: the main insights remain unchanged if supervising machines also entails positive and increasing costs. Introducing such costs, however, complicates the analysis without providing additional insight.

The decision to impose supervision costs on novices rather than machines is deliberate. Doing so allows me to compare automation improvements—modeled as reductions in machine costs (ρ)—with improvements in routine labor-augmenting technologies. The latter naturally correspond to reductions in novice supervision costs c, capturing scenarios in which algorithms or machines assist novices in avoiding mistakes, thereby enhancing their productivity without substituting their labor.

In practice, whether novices or machines require more supervision is context-dependent. In some settings, machines clearly outperform novices. For example, senior investment bankers appear to prefer platforms like Factset and CapIQ—which automate data gathering, interpretation, and financial calculations—over junior analysts, whose output is typically less reliable (Beane and Anthony, 2024). Similarly, surgical residents tend to operate more slowly and with less precision than robotic

surgical systems, potentially prolonging procedures and increasing patient risk (Beane, 2024b). In other settings, however, machines themselves may require substantial oversight. Generative AI, for example, can produce outputs that appear plausible but are factually incorrect (Autor, 2024; Becker et al., 2025). Thus, supervision costs are specific to the technology and tasks involved.

On Fréchet and Fractional-Order Statistics.— Primarily for tractability, I assume that both experts' skills and idea quality are distributed Fréchet and employ fractional-order statistics. These assumptions ensure that the distribution of expert skills remains Fréchet over time and are standard in the literature on idea flows and economic growth (e.g., Lucas, 2009; Lucas and Moll, 2014; De la Croix et al., 2018; Buera and Lucas, 2018).

A side-effect of adopting the Fréchet distribution is that it implies unbounded support, suggesting—somewhat unrealistically—that all productive knowledge already exists at the outset. De la Croix et al. (2018) show, however, that unbounded support can be naturally interpreted as an approximation of their more realistic "Da Vinci model." In this model, highly advanced ideas—such as Leonardo Da Vinci's visionary inventions—initially exist only theoretically and become feasible only as society's knowledge accumulates over time.

On the Production Function.— The production function (1) is a special case of the more general formulation (inspired by Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2018, 2019):

$$Y_{i,t} \equiv q_{i,t} \exp\left(N \int_0^1 \ln\big(a_{i,t}(s)\gamma_\ell(s) + [1 - a_{i,t}(s)]\gamma_m(s)\big)\, ds \right), \tag{4}$$

where $\gamma_\ell(s)$ and $\gamma_m(s)$ denote the productivity of novices and machines in task s, respectively, and $a_{i,t}(s) \in \{0, 1\}$ indicates whether expert i assigns task s to a novice ($a_{i,t}(s) = 1$) or to a machine ($a_{i,t}(s) = 0$).

The baseline specification corresponds to the special case $\gamma_\ell(s) = \gamma_m(s)$ for all $s \in [0, 1]$ (both normalized to 1). This restriction ensures that all experts employ the same number of novices, regardless of their skill. As a result, the distribution of expert skills retains its Fréchet structure across periods, allowing for a tractable and transparent equilibrium analysis.[11]

This analytical convenience, however, is not unique to the baseline model. Equal novice demand across experts can also emerge under alternative parameterizations. For instance, suppose novices complete all tasks at productivity $\gamma_\ell(s) = h > 0$, while machines perform only a subset $A < N$ of tasks at much higher productivity ($\gamma_m(s) = m > 0$ for $s \leq A/N$, and $\gamma_m(s) = 0$ otherwise, with m sufficiently large). In this case, all experts optimally assign the A feasible tasks to machines and the

[11] In general, when $\gamma_\ell(s)$ and $\gamma_m(s)$ vary across tasks, experts' hiring becomes skill-dependent and the equilibrium must be solved numerically. Two margins then matter. First, whether higher-skilled experts hire more novices—accelerating diffusion—or fewer—slowing it down. Second, whether high- or low-skilled experts adjust more strongly to automation. If high-skill experts reduce novice hiring more aggressively, the disruption to tacit knowledge transmission is amplified; if low-skill experts do so, the effect is partly mitigated. In all cases, however, the baseline insight remains intact: equilibrium novice–expert interactions are inefficiently scarce, and better automation further reduces them, lowering long-run growth.

remaining $N - A$ to novices, yielding:

$$Y_{i,t}^{*} = q_{i,t} m^{A} h^{N-A}.$$

In this case, although equilibrium novice labor differs, the law of motion for tacit knowledge and output is identical to the baseline (see the Online Appendix for details). This equivalence, in fact, holds for any parametrization such that (i) only a subset of tasks is automatable, and (ii) all experts optimally allocate all such tasks to machines. No other assumptions on $\gamma_{\ell}(s)$ and $\gamma_{m}(s)$ are required.

This "all-automatable-tasks-get-automated" specification is widely used in the automation literature (see, e.g., Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2019; Autor and Thompson, 2025). For most of this paper, however, I work with the baseline production function (1), because it allows me to characterize the equilibrium in closed form across all parameter values, not just in cases where machines are sufficiently more productive in the tasks they can perform. That said, the "all-automatable-tasks-get-automated" specification is still useful, and I return to it in Section 5, as it provides a convenient framework to incorporate external estimates of one-off productivity gains from AI adoption.

Finally, for simplicity, the model assumes that all entry-level tasks have identical learning value. Although real-world tasks vary in complexity and educational content, introducing such heterogeneity would significantly complicate the analysis without yielding additional insights.[12] Moreover, given the lack of systematic evidence on precisely which entry-level tasks—high- or low-learning-value—are most likely to be automated, I abstract from these differences, implicitly assuming that automation affects roughly equal proportions of both types of tasks.

On Novice-Expert Matching, the Zero-Lower Bound on Wages, and Contract Incompleteness.— As emphasized in the Related Literature, apprenticeship-like contracts are inherently incomplete because the knowledge involved is tacit. In the model, this contractual incompleteness manifests in two forms.

First, expert skills are unobservable ex ante, so novice—expert matching is random. Experts are thus compensated only by the cohort's average level of skill, rather than at their own. This assumption reflects the difficulty of assessing tacit expertise in advance (De la Croix et al., 2018). In practice, some partial observability exists—for instance, prestigious firms use reputation to attract juniors at lower pay—but abstracting from it preserves tractability and maintains the Fréchet structure of skills across cohorts. Introducing partial observability would require numerical methods to characterize equilibrium—reducing transparency—and is unlikely to overturn the model's core insights.[13]

Second, novices are wealth-constrained, so wages cannot fall below zero. This reflects both their

[12] If automation primarily targets entry-level tasks with low learning value, its negative impact on knowledge transmission would likely be minimal. Conversely, automating tasks that provide substantial learning opportunities would further disrupt the intergenerational transmission of knowledge.

[13] Partial observability would tend to accelerate knowledge transmission and reduce equilibrium inefficiency relative to the baseline, as higher-skilled experts would optimally hire more novices. Yet the equilibrium would remain inefficient, and lower automation costs would still disrupt knowledge transmission. Whether this disruption is amplified or mitigated relative to the baseline is ambiguous, hinging on which group of experts adjusts more strongly.

lack of financial resources to pay experts directly for training and the fact that the tacit, intangible nature of the knowledge sought cannot serve as collateral (see Section 2). Moreover, trainees must earn at least subsistence wages while in training. Together, these constraints render negative wages practically infeasible.[14]

As discussed in Section 4.1, these frictions generate socially insufficient learning in equilibrium. The Online Appendix compares the specific form of contractual incompleteness assumed here with the frictions emphasized in previous work (e.g., Garicano and Rayo, 2017; De la Croix et al., 2018).

On the Transmission of Tacit Knowledge (Breadth vs. Depth).— Finally, in the model, eventual skill depends on the best expert encountered, so novice labor is valuable for learning because it increases the *breadth* of expert interactions. Alternatively, the relationship between novice labor and learning could stem from the *depth* of interactions: novices may work under one or a few mentors, and more labor allows them to better internalize and improve upon that mentor's knowledge.

Whether learning depends on breadth or depth is unlikely to change the qualitative insights, but focusing on breadth is analytically convenient since it allows me to exploit the tools developed in the idea flows and growth literature (Kortum, 1997; Lucas, 2009; De la Croix et al., 2018).

4 The Impact of Automating Entry-Level Work

In this section, I characterize the economy's long-run equilibrium outcomes. I then analyze how improvements in the automation of entry-level work—modeled as a reduction in machine costs—can disrupt intergenerational knowledge transmission, ultimately reducing long-term growth. These results set the stage for Section 5, which provides numerical illustrations of the potential magnitude of these disruptions in the context of recent advances in AI.

4.1 Long-Run Equilibrium Outcomes

In the long run, the economy converges to one of three stable regimes—Learning Breakdown (LB), Constrained Learning (CL), and Full Learning (FL)—depending on the cost of machines and the initial stock of knowledge. Each regime differs in how routine tasks are allocated between novices and machines, influencing the effectiveness with which tacit knowledge is transmitted across generations.

Proposition 1. *Fix an initial knowledge stock $k_0 > 0$ and let:*

$$\underline{\rho} \equiv c(1-\nu), \quad \bar{\rho} \equiv cN, \quad \textit{and} \quad k^{\dagger} \equiv \frac{(c+\mu)(1-\nu)-\rho}{\beta\,\theta\,\Gamma(1-\theta)}$$

Except for the knife-edge parameter values described in Remarks 1 and 2 below, the competitive equilibrium converges to one of the following stable long-run regimes:

[14]The analysis does not fundamentally depend on a zero wage floor. Indeed, the model easily accommodates a general lower bound on wages, $w_t \geq \underline{w}$, where this threshold $\underline{w}$ may be positive or negative, depending on the institutional context.

1. **Learning Breakdown (LB).** *If $\rho \leq \underline{\rho}$ or $k_0 < k^\dagger$, the economy converges asymptotically to the LB regime, characterized by:*

$$L_t^* \to \frac{\rho}{c+\mu} < 1, \quad w_t^* \to \frac{\mu\rho}{c+\mu} > 0, \quad k_t^* \to 0, \quad g_t^k \to 0,\ g_t^Y \to 0, \quad \textit{and} \quad Y_t^* \to q^{\min}$$

During the transition, k_t^*, L_t^*, and w_t^* are strictly decreasing, weakly decreasing, and weakly increasing in time, respectively.

2. **Constrained Learning (CL).** *If $\underline{\rho} < \rho < \bar{\rho}$ and $k_0 > k^\dagger$, the economy reaches in finite time the CL regime, characterized by:*

$$L_t^* = \frac{\rho}{c} \in (1, N), \quad w_t^* = 0, \quad g_t^k = \left(\frac{\rho}{c} + \nu\right)^\theta - 1 > 0, \quad \textit{and} \quad \lim_{t\to\infty} g_t^Y = \left(\frac{\rho}{c} + \nu\right)^\theta - 1$$

During the transition, k_t^* and L_t^* strictly increase in time, while w_t^* strictly decreases in time.

3. **Full Learning (FL).** *If $\rho \geq \bar{\rho}$ and $k_0 > k^\dagger$, the economy reaches in finite time the FL regime, characterized by:*

$$L_t^* = N, \quad w_t^* = 0, \quad g_t^k = (N+\nu)^\theta - 1 > 0, \quad \textit{and} \quad \lim_{t\to\infty} g_t^Y = (N+\nu)^\theta - 1$$

During the transition, k_t^* and L_t^* strictly increase in time, while w_t^* strictly decreases in time.

Remark 1 (Knife-Edge Stable Equilibrium). *When $\rho = \underline{\rho}$ and $k_0 \geq k^\dagger$, the economy immediately (at $t = 0$) reaches a stable steady state characterized by $L_t^* = 1 - \nu$, $w_t^* = 0$, and $k_t^* = k_0$ for all $t \geq 0$.*

Remark 2 (Unstable Equilibrium). *When $\underline{\rho} < \rho$ and $k_0 = k^\dagger$, the economy immediately reaches an unstable steady state characterized by $L_t^* = 1 - \nu$, $w_t^* = \rho - c(1-\nu)$, and $k_t^* = k^\dagger$ for all $t \geq 0$.*

Proof. See Appendix A. □

Proposition 1 is illustrated in Figure 1. The figure depicts the parameter combinations of initial knowledge stock k_0 and machine costs ρ that lead to each of the three long-run regimes.

As the figure shows, when machines are sufficiently cheap ($\rho < \underline{\rho}$), the economy converges to the Learning Breakdown (LB) regime. Even at zero wages, experts do not allocate enough tasks to novices to sustain knowledge accumulation ($L^d(0, \rho) = \rho/c < 1-\nu$).[15] Consequently, tacit knowledge erodes over time, novice labor declines, and wages rise, shrinking output. This occurs because each generation is less knowledgeable than the previous one, so novices see fewer learning opportunities and therefore demand higher wages over time. These rising wages then make substituting novice labor with machines increasingly attractive, further accelerating the erosion of knowledge. In the long run, knowledge and output converge to the minimal skill level $q^{\min}$.

The outcome changes when machine prices rise to intermediate levels ($\underline{\rho} < \rho < \bar{\rho}$). Experts are then willing to allocate a larger share of tasks to novices—enough to sustain knowledge accumulation

[15] Note that convergence to the LB regime can occur only if $\nu < 1$. Intuitively, if $\nu \geq 1$, society can sustain—and potentially expand—its knowledge stock even without direct intergenerational transmission, because innovation provides an apprenticeship-independent source of knowledge.

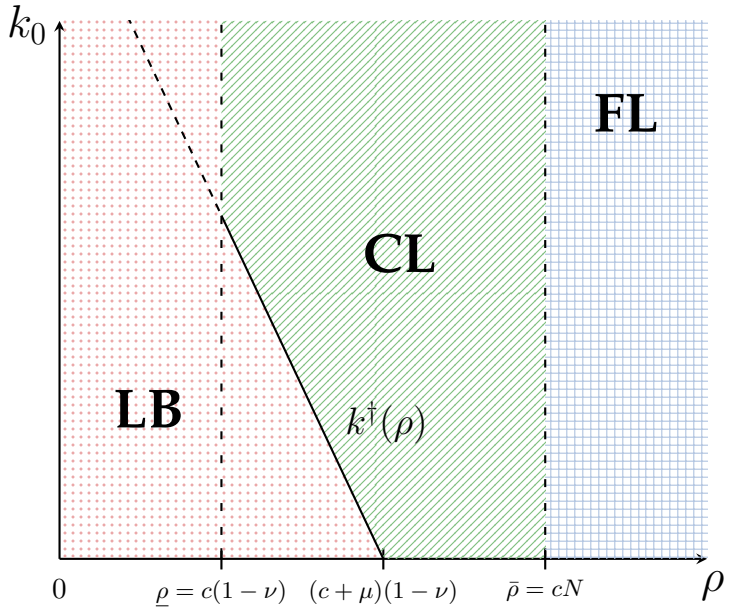


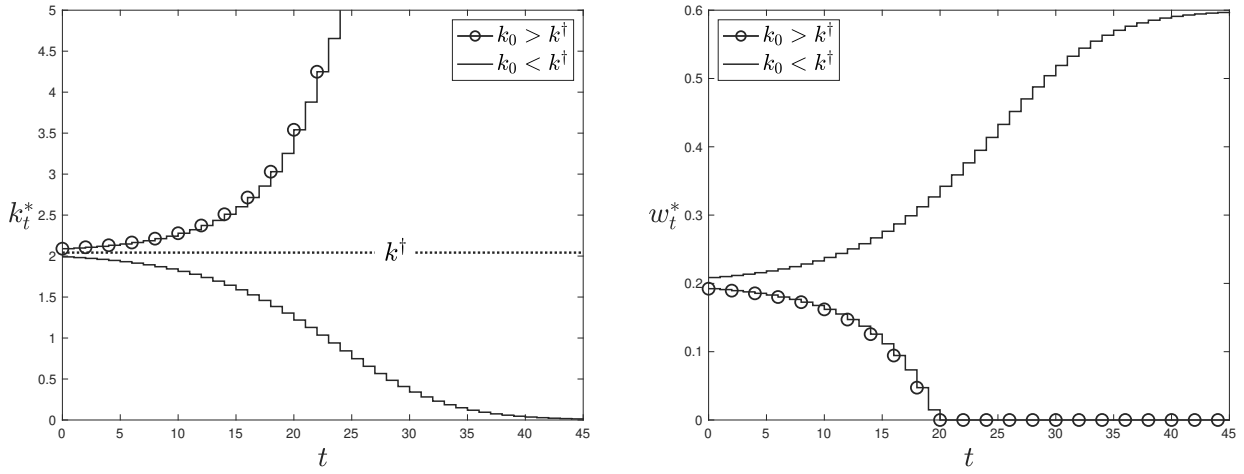
Figure 1: Long-Run Equilibrium Outcomes as a Function of (ρ, k_0)

Notes. Parameter Values: $N = 3$, $\theta = 0.45$, $c = 1$, $\mu = 1$, $\nu = 0$, and $\beta = 0.96^{20}$ (plus any $q^{\min} \geq \rho N$). The region with the **dotted pattern** corresponds to combinations of (ρ, k_0) for which the economy converges to the Learning Breakdown (LB) steady state. The **northeast diagonal lines** indicate convergence to the Constrained Learning (CL) steady state, and the **grid pattern** to the Full Learning (FL) steady state.

over time—as long as wages remain sufficiently low. In this scenario, the long-run outcome depends on the initial knowledge stock. If $k_0 > k^\dagger$, the economy converges to the Constrained Learning (CL) regime, where knowledge and output grow, though at a relatively moderate rate. Otherwise, the economy again descends into the LB regime.

The intuition is as follows. When $k_0 > k^\dagger$, the first generation of novices is willing to work for very low wages in exchange for learning from relatively skilled experts. This prompts experts to hire more novices, who in turn become highly skilled. This dynamic reinforces itself over time, with each cohort willing to work for less to learn from increasingly knowledgeable predecessors, leading to steady knowledge accumulation. Eventually, however, wages hit the zero lower bound. At that point, novices would pay to displace machines and learn more—but lack the means to do so. Consequently, machines remain in use, limiting the growth of knowledge and output. The solid lines with circle markers in Figure 2, panels (a) and (b), illustrate the resulting wage and knowledge dynamics in this case.

By contrast, when $k_0 < k^\dagger$, the initial generation of novices is less eager to learn and demands higher wages. Experts respond by using more machines, thereby reducing available learning opportunities. Each new generation of experts thus becomes less knowledgeable than the previous one, prompting subsequent cohorts of novices to demand even higher wages. This negative feedback loop results in declining knowledge and ultimately a collapse in growth, as illustrated by the solid lines without markers in Figure 2, panels (a) and (b).



(a) Tacit Knowledge (b) Wages

Figure 2: Equilibrium Dynamics when $\underline{\rho} < \rho < \bar{\rho}$

Notes. Parameter Values: $N = 3$, $\theta = 0.5$, $c = 1$, $\mu = 1$, $\beta = 0.96^{20}$, $\rho = 1.2$, $\nu = 0$, and $q^{\min} = 3.6$. These parameter values imply that $k^\dagger = 2.04$. The solid line with the marker corresponds to $k_0 = 2.09 > k^\dagger$, so the economy converges to the CL regime. The solid line without a marker corresponds to $k_0 = 1.99 < k^\dagger$, so the economy converges to the LB regime.

Finally, as Figure 1 shows, when machine prices are sufficiently high ($\rho \geq \bar{\rho}$), the economy converges to the Full Learning (FL) regime. In the particular case shown in the figure, convergence occurs regardless of the initial knowledge stock, as the condition $\rho \geq \bar{\rho}$ alone ensures $k_0 > k^\dagger$. More generally, however, both high machine costs and a sufficiently large initial knowledge stock are required—mirroring the logic of the CL regime. In the FL regime, experts assign all routine tasks to novices, maximizing novice learning. As a result, each generation reaches its full potential, and output grows at the highest feasible rate.[16]

As formally shown in the Online Appendix, the competitive equilibrium is inefficient. Relative to the first-best allocation—defined as the one that maximizes the discounted sum of welfare—the equilibrium exhibits excessive entry-level automation and insufficient learning. This inefficiency stems from the two contractual frictions discussed in Section 3.2: the unobservability of tacit knowledge (which prevents experts from capturing the full social value of their expertise), and novices' wealth constraints (which prevent novices from compensating experts to displace machines when learning is valuable). Together, these frictions generate underinvestment in knowledge acquisition.

The unobservability of tacit knowledge also implies that even if negative wages were feasible, inefficiencies would persist and the model's key insights would remain unchanged (see the Online Appendix).[17] The non-negative wage assumption is nevertheless useful: beyond its practical justification (see Section 3.2), it ensures the existence of the Constrained Learning (CL) regime—a balanced

[16] As noted in footnote 1, the result that growth is maximized when automation is sufficiently costly rests on the assumption that all routine tasks provide meaningful learning opportunities. If, instead, some entry-level tasks offer no learning value, automating them could increase efficiency without necessarily reducing long-term growth.

[17] Put differently, socially excessive automation does not hinge on both contractual frictions operating simultaneously; either one alone is sufficient.

growth path in which both machines and novices perform a fraction of routine tasks in steady state. Without it, the economy would inevitably converge either to the Learning Breakdown (LB) or Full Learning (FL) regime, ruling out this empirically plausible intermediate scenario.

In sum, Proposition 1 characterizes the long-run equilibrium outcomes of this economy. Building on this characterization, the next section shows that improvements in automation can boost short-term productivity but ultimately disrupt intergenerational knowledge transfer, significantly undermining long-term growth and welfare.

4.2 The Impact of Improvements in Automation

Consider an economy that at time $\tau - 1$ is either in the FL regime, the CL regime, or approximately at the LB regime. The next proposition characterizes the effect of a one-time (unexpected) automation improvement, modeled—following Autor et al. (2003)—as a permanent reduction in the cost of machines from ρ to a lower level $\rho' < \rho$. This shock takes place at the beginning of period τ. I restrict attention to cases where $\rho' < \bar{\rho}$; otherwise, the improvement would be inconsequential, as machines would remain unused both before and after the shock.

Proposition 2. *Let $\mathcal{W}_t^*(\rho) \equiv \mathcal{W}(L_t^*(\rho), k_t^*(\rho); \rho)$ be society's equilibrium per period welfare at time t given machine cost ρ.*

- *Suppose the economy is initially in the CL or FL regime at time $\tau - 1$. Then, a permanent reduction in machine costs at the beginning of time τ immediately increases welfare but eventually reduces it in the long run relative to the no-shock scenario. Formally,*

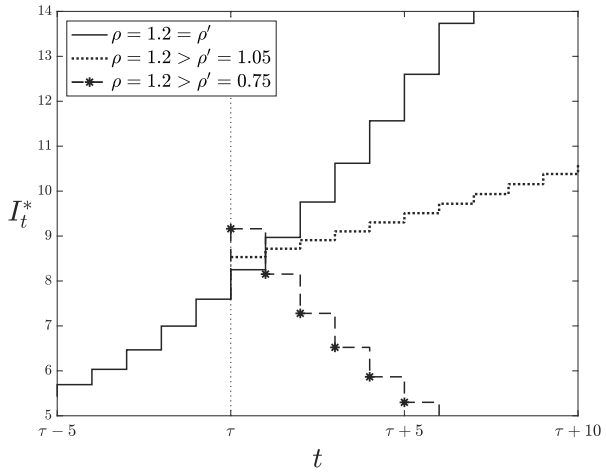
$$\mathcal{W}_\tau^*(\rho') > \mathcal{W}_\tau^*(\rho) \quad \textit{and} \quad \exists\, T \in (\tau, \infty) \ \textit{s.t.}\ \mathcal{W}_t^*(\rho') < \mathcal{W}_t^*(\rho),\ \forall\, t \geq T.$$

- *Suppose instead the economy is approximately in the LB regime at time $\tau - 1$. Then, a permanent reduction in machine costs at the beginning of time τ permanently improves welfare. Formally,*

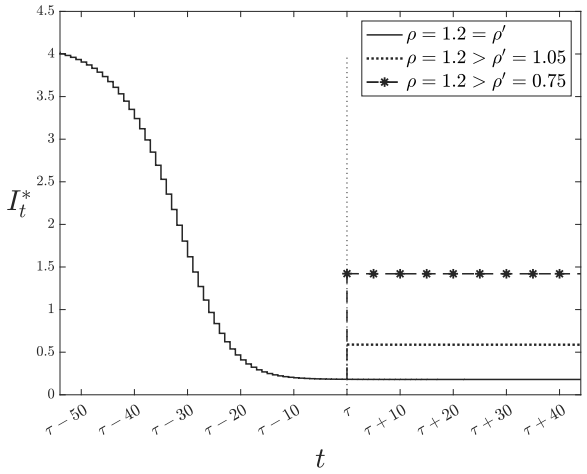
$$\mathcal{W}_t^*(\rho') > \mathcal{W}_t^*(\rho), \quad \forall\, t \geq \tau.$$

Proof. See Appendix A. □

The intuition behind the first part of the proposition is as follows. When the economy begins in the CL or FL regime, reducing machine costs immediately lowers production expenses, increasing the income and productivity of current experts, who now produce the same output at a lower cost. This results in higher per-period welfare. However, these short-term gains come at the expense of future generations. By substituting novice labor with machines, automation reduces the expert-novice interactions crucial for transmitting tacit skills. Consequently, subsequent generations—including current novices—become progressively less skilled and productive. Thus, the initial welfare gains enjoyed by today's experts are gradually offset, and eventually reversed, by the cumulative erosion of human capital in future cohorts.



(a) Initial Steady State: CL Regime



(b) Initial Steady State: LB Regime (Approximately)

Figure 3: The Impact of Improvements in Automation

Notes. Parameter values: Both panels use $N = 3$, $\theta = 0.5$, $c = 1$, $\mu = 1$, $\beta = 0.96^{20}$, $\rho = 1.2$, $\nu = 0$, and $q^{\min} = 3.6$, implying a critical threshold of $k^\dagger = 2.04$. Panel (a) sets the initial capital stock at $k_0 = 2.09 > k^\dagger$, so the economy is initially in the CL regime prior to the automation shock. In contrast, panel (b) sets $k_0 = 1.99 < k^\dagger$, placing the economy initially in the LB regime (convergence here is asymptotic). Each panel compares experts' income under the baseline scenario—in which machine costs remain constant at $\rho = \rho' = 1.2$—with two alternative scenarios involving automation improvements at time τ, reducing machine costs to either $\rho' = 1.05$ or $\rho' = 0.75$.

Figure 3(a) illustrates this phenomenon by examining how an automation shock affects experts' income when the economy initially lies in the CL regime. The figure compares a baseline scenario, in which machine costs remain constant at $\rho = \rho' = 1.2$, with two scenarios featuring automation improvements at time τ that reduce machine costs to either $\rho' = 1.05$ or $\rho' = 0.75$. As shown, the automation shock initially raises experts' income but subsequently slows—and may even reverse—its growth. Specifically, under a moderate cost reduction ($\rho' = 1.05$), experts' income continues growing, albeit more slowly, as the economy still converges to the CL regime. In contrast, a larger reduction in machine costs ($\rho' = 0.75$) accelerates skill erosion, ultimately driving the economy toward the LB regime and causing the income of future cohorts of experts to decline over time.

The intuition for the second part of the proposition follows from the observation that once the economy reaches the LB regime, its stock of tacit knowledge has already been largely depleted, and intergenerational learning has effectively ceased. Under these conditions, the trade-off between immediate income gains and future skill erosion disappears. Automation improvements become unequivocally beneficial since there is little remaining human capital to be lost. Figure 3(b) illustrates this scenario, showing that the same automation shocks examined in panel (a) result in permanent increases in experts' income when the economy is initially in the LB regime.

So far, I have focused exclusively on the effects of improvements in automation, defined—following Acemoglu et al. (2024)— as technologies that displace novices from tasks they currently perform. It is instructive, however, to contrast these task-displacing technologies with improvements in routine labor-augmenting technologies, defined as those that that raise novice productivity without displac-

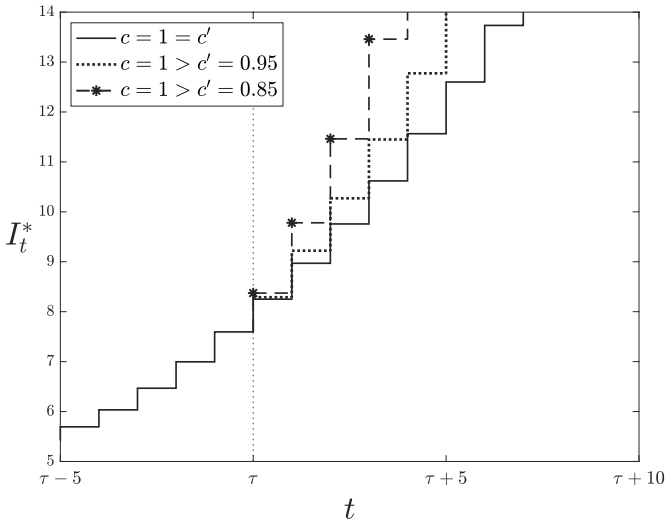


Figure 4: The Impact of Improvements in (Routine) Labor Augmenting Technologies

Notes. Parameter values: $N = 3$, $\theta = 0.5$, $\rho = 1.2$, $\mu = 1$, $\beta = 0.96^{20}$, $\nu = 0$, $k_0 = 2.09$, and $q^{\min} = 3.6$. These parameter choices imply a threshold $k^\dagger = 2.04 < 2.09 = k_0$, situating the economy in the CL regime prior to the technological improvement. The figure compares experts' income under the baseline scenario—in which supervision costs remain constant at $c = c' = 1$—with two alternative scenarios involving technological improvements (such as enhanced algorithms or better machinery) that reduce supervision costs to either $c' = 0.95$ or $c' = 0.85$ at time τ.

ing novice labor. Such labor-augmenting improvements can naturally be modeled as reductions in the cost of supervising novices (c), capturing scenarios in which algorithms or machines help novices avoid mistakes.

Figure 4 examines this scenario starting from an economy initially situated in the CL regime. The figure compares a baseline scenario—in which supervision costs remain constant at $c = c' = 1$—with two alternative cases in which supervision costs decline moderately (to $c' = 0.95$) or substantially (to $c' = 0.85$) at time τ. As illustrated, lower supervision costs boost the productivity and thus the income of the current generation of experts—mirroring the short-term effects of automation improvements. In sharp contrast to automation, however, these labor-augmenting improvements strengthen novice-expert interactions by increasing experts' incentives to hire novices. Thus, labor-augmenting technologies deliver immediate income gains while simultaneously promoting long-term growth.

As discussed in Section 2, recent forecasts—such as those by Goldman Sachs (2023), McKinsey & Company (2023), and Aghion and Bunel (2024)—project substantial productivity gains from the widespread adoption of AI. These anticipated improvements primarily represent one-off efficiency gains realized upon initial implementation. However, these studies do not explicitly differentiate whether productivity enhancements primarily stem from automation—where AI substitutes for human labor—or from augmentation—where AI boosts productivity without displacing workers.

The results presented in this section underscore the importance of this distinction. Advances in the automation of entry-level tasks generate immediate efficiency benefits but simultaneously disrupt novice-expert interactions. Such disruptions pose substantial risks to future productivity. By contrast, improvements in augmentation technologies yield immediate benefits without compromising—and

potentially even strengthening—the novice-expert interactions vital for sustained long-run productivity growth.

Unfortunately, early evidence suggests that firms currently view generative AI primarily as a tool for automation rather than augmentation, particularly at the entry-level. As discussed in Section 2, this evidence includes empirical findings by Berger et al. (2024), Brynjolfsson et al. (2025a), and Lichtinger and Maasoum (2025) and expert commentary, such as Dario Amodei's interview with VandeHei and Allen (2025).

This perspective is further reinforced by The World Economic Forum's (2025) *Future of Jobs Survey*, which gathers insights from over 1,000 employers representing more than 14 million workers across 22 industry clusters and 55 economies. The survey results indicate that, currently, 47% of workplace tasks are performed exclusively by humans, 30% are completed collaboratively by humans and machines, and 22% are fully automated by machines or algorithms. By 2030, however, respondents anticipate a notable shift: tasks performed solely by humans will decrease to 33%, tasks involving human-machine collaboration will increase modestly to 33%, and tasks handled entirely by machines will rise significantly to 34%.

Collectively, this early evidence raises concerns about the underlying nature of projected AI-driven productivity gains. If these gains primarily result from entry-level automation, the immediate productivity improvements associated with AI adoption may mask significant and lasting damage to human capital formation, ultimately jeopardizing sustained economic growth.

5 Quantitative Illustration

Motivated by concerns that AI-driven entry-level automation could disrupt knowledge transmission and growth, I now perform simple back-of-the-envelope calculations to gauge the potential magnitude of these effects. This exercise is neither a formal calibration nor a forecast. It does not attempt to match empirical data and deliberately abstracts from possible compensating factors, such as AI's potential to stimulate innovation. Instead, its goal is to illustrate how advances in AI might undermine long-run economic growth by weakening the intergenerational transmission of tacit knowledge.

5.1 Steady State Losses

The quantitative illustration of long-term growth losses is based on a minimal set of assumptions that are fully consistent with—but do not require—the complete structure of the baseline model developed earlier. This ensures that the results are not overly sensitive to specific modeling choices.

Specifically, the exercise rests on three central assumptions:

(a) **OLG Structure**: Each generation participates in the economy for exactly two "OLG periods"—initially as novices and subsequently as experts.
(b) **Steady-State Growth**: In steady-state, the growth rate of aggregate per capita output equals the

growth rate of the stock of tacit knowledge.

(c) **Evolution of Tacit Knowledge**: Society's stock of tacit knowledge, denoted by k_t, corresponds to the scale parameter of a Fréchet distribution and evolves according to:

$$k_{t+1} = (L_t^* + \nu)^\theta k_t$$

where L_t^* denotes the measure of entry-level tasks performed by novices, $1/\theta$ is the shape parameter of the Fréchet distribution, and $\nu > 0$ is a parameter capturing society's capacity for innovation.

An example beyond the baseline that is consistent with these three assumptions is the alternative "all-automatable-tasks-get-automated" specification introduced in Section 3.2, where machines can only execute a subset $A < N$ of tasks but do so at significantly higher productivity than novices.

Consider now a one-time improvement in automation resulting in a fraction $a \in [0, 1]$ of tasks previously performed by novices becoming automated. Under the above assumptions, the economy's annual steady-state growth rate of per capita output prior to the automation shock is:

$$g^Y = (L^* + \nu)^{\theta/T} - 1,$$

while the growth rate after the shock becomes:

$$g^{Y\prime} = [(1-a)L^* + \nu]^{\theta/T} - 1.$$

where T denotes the number of calendar years corresponding to one OLG period (i.e., the number of years individuals spend in each career status), and L^* represents the steady-state measure of tasks carried out by novices before the automation improvement.

The resulting reduction in the steady-state growth rate from increased automation (expressed in percentage points per year) is then given by:

$$\Delta g^Y \equiv 100 \times (g^Y - g^{Y\prime}) = 100 \times (1 + g^Y)\big[1 - (1 - ax^{1/\theta})^{\theta/T}\big], \text{ where } x \equiv \left(\frac{L^*}{L^* + \nu}\right)^\theta \tag{5}$$

Here, x denotes the fraction of steady-state growth attributable to the diffusion of tacit knowledge, while $1 - x$ represents the fraction attributable to the creation of new ideas.

Thus, to explore the long-term growth consequences of the automation shock, I only need five pieces of information:

(i) The conversion of OLG periods into calendar years, T.
(ii) The baseline annual growth rate of output prior to the automation shock, g^Y.
(iii) A plausible value for the parameter θ.
(iv) The fraction of steady-state growth driven by learning/diffusion, x.
(v) The proportion a of entry-level tasks that become automated due to AI.

I set $T = 20$, implying individuals spend roughly ages 25 to 45 as novices and ages 45 to 65 as experts. Regarding the baseline annual growth rate, I follow Jones (2016) and set it to $g^Y = 2\%$, reflecting the average per-capita growth rate of output observed in the U.S. over the past 150 years.

I also adopt $\theta = 0.5$ as my preferred specification, in line with Lucas (2009), Lucas and Moll (2014), and De la Croix et al. (2018). Lucas (2009) initially calibrated this parameter to match observed earnings dispersion in contemporary U.S. data. However, for robustness, I also consider an alternative scenario with $\theta = 0.28$, as used by Caicedo et al. (2019).

I set $x = 65\%$ as my baseline, with robustness checks at $x = 50\%$ and $x = 80\%$. This baseline reflects evidence from Eaton and Kortum (1999), who estimate that international technology diffusion alone accounts for roughly 40% of productivity growth in the U.S., and Santacreu (2015), who provides a lower estimate of about 25% for developed economies. Taking the midpoint (32.5%) and doubling it—assuming domestic diffusion contributes equally to international diffusion—results in the baseline choice of $x = 65\%$. The robustness checks capture a more conservative scenario ($x = 50\%$), assuming a smaller domestic contribution, and a more generous scenario ($x = 80\%$), reflecting potentially larger domestic diffusion.

Finally, regarding the proportion a of entry-level tasks automated by AI, I consider three scenarios:

Scenario 1 - $a = 5\%$: This scenario is inspired by Acemoglu (2024). He uses estimates from Eloundou et al. (2024) to conclude that approximately 19.9% of (wage-bill-weighted) tasks are currently exposed to AI, and then multiplies this figure by the fraction of AI-exposed tasks currently profitable to automate (23%), as estimated by Svanberg et al. (2024).

Scenario 2 - $a = 15\%$: This scenario is based on the *Future of Jobs Survey* from The World Economic Forum (2025) discussed in Section 4.2. Survey respondents estimate that machines and algorithms currently perform approximately 22% of all workplace tasks without human involvement and project this share to rise to 34% by 2030. Thus, they estimate that approximately 15% of tasks currently carried out by humans—either independently or jointly with machines—will become fully automated by the end of the decade.

Scenario 3 - $a = 30\%$: This scenario is inspired by Aghion and Bunel (2024). Their calculation follows a similar logic to Acemoglu (2024), but relies on a broader measure of AI exposure by Pizzinelli et al. (2023), who find that around 60% of (employment-weighted) tasks are exposed to AI. Aghion and Bunel (2024) also assume it will be profitable to use AI in approximately half of these tasks, driven by rapid anticipated reductions in computing costs (about 22% annually, as projected by Besiroglu and Hobbhahn, 2022). These estimates are also roughly in line with recent industry reports, such as Goldman Sachs (2023), which suggests that around a quarter of employment-weighted tasks could potentially be automated using AI.

It is important to note that both Acemoglu (2024) and Aghion and Bunel (2024) remain agnostic as to whether the productivity gains they forecast primarily reflect automation or augmentation. In contrast, I assume throughout that AI adoption takes the form of automation, consistent with the evidence in Section 4.2 suggesting that firms currently view recent advances in generative AI mainly as opportunities for automation rather than augmentation.

In addition, I assume that automation is restricted to entry-level tasks, in line with evidence from

	$\boldsymbol{\theta = 0.5}$			$\theta = 0.28$		
	$x = 50\%$	$\boldsymbol{x = 65\%}$	$x = 80\%$	$x = 50\%$	$x = 65\%$	$x = 80\%$
$a = 5\,\%$	0.0321	**0.0544**	0.0829	0.0060	0.0154	0.0325
$a = 15\,\%$	0.0974	**0.1668**	0.2570	0.0181	0.0467	0.0999
$a = 30\,\%$	0.1986	**0.3450**	0.5421	0.0365	0.0950	0.2072

Table 1: Estimated Reduction in Steady State Growth (pp/year)

Notes. Remaining parameter values: $g^Y = 2\%$ and $T = 20$. The baseline scenario is indicated in bold.

Berger et al. (2024), Brynjolfsson et al. (2025a), and Lichtinger and Maasoum (2025). Moreover, I abstract from task re-bundling or the creation of new junior-level roles, either of which could mitigate the reduction in novice learning opportunities.[18] In Section 5.3, I revisit this assumption by considering a case in which entry-level automation is temporary: it affects the economy for a finite period of time before new junior positions emerge.

The results of this numerical exercise are summarized in Table 1. In the baseline scenario with $\theta = 0.5$ and $x = 65\%$, annual growth losses range from 0.05 percentage points when $a = 5\%$ to 0.35 percentage points when $a = 30\%$. These losses become larger when $x = 80\%$, as a higher proportion of baseline growth then stems from diffusion—the very channel disrupted by automation. Conversely, losses diminish when $\theta = 0.28$, since a lower θ shifts a greater share of growth toward innovation, thus mitigating automation's negative impact.

5.2 Short-Run Gains versus Long-Run Losses

Section 5.1 showed that entry-level automation lowers the steady-state growth rate of the economy. Yet these long-run losses are only part of the picture. Recent discussions of AI adoption emphasize potential short-run productivity gains, raising the question of how these immediate benefits compare with the longer-term costs. To address this question, I draw on external forecasts of one-time adoption gains—specifically those by Acemoglu (2024) and Aghion and Bunel (2024).

Acemoglu (2024) projects cumulative TFP growth of about 0.71% over ten years—about 0.07 percentage points annually—under the assumption that AI adoption affects around 5% of tasks. By contrast, Aghion and Bunel (2024) project cumulative TFP growth of about 7% over the same horizon (approximately 0.68 percentage points annually), under the assumption that AI affects closer to 30% of tasks. Together, these forecasts capture the range of short-run adoption gains emphasized in current debates and serve as natural benchmarks for the quantitative exercise that follows.

To make use of these forecasts, however, the estimates must be translated into the model's framework. The complication is that Acemoglu (2024) and Aghion and Bunel (2024) report productivity

[18] Another reason these calculations may overstate growth losses is that they assume knowledge diffusion occurs exclusively through novice–expert (vertical) interactions, abstracting from potential peer-to-peer (horizontal) transmission.

gains directly in terms of TFP increases. In contrast, the baseline model in Sections 3 and 4 represents improvements in automation as reductions in machine costs. This mismatch prevents a direct comparison between their projections and the model's steady-state growth implications.

To address this issue, I turn to the alternative specification introduced in Section 3.2. In this setting, machines can perform only a subset $A < N$ of tasks before the automation shock, but do so at much higher productivity than novices. As a result, all experts optimally allocate A tasks to machines and the remaining $L^* = N - A$ tasks to novices, so the equilibrium output pre-shock at time t is:

$$Y_t = \bar{q}_t m^A h^{N-A},$$

where $\bar{q}_t$ is the average expert skill at time t, m the productivity of machines, and h the productivity of novices.

As discussed in Section 5.1, this specification delivers the same long-run losses as the baseline model, but has the added advantage of allowing improvements in automation to be represented as direct increases in TFP. This alternative, therefore, provides a natural bridge between the external forecasts and the model's growth implications. In particular, if AI enables the (profitable) automation of an additional $aL^* = a(N - A)$ tasks at productivity $m_{\text{AI}} > h$, the resulting one-off productivity gain from adoption at time t is:

$$\Delta\text{TFP} \equiv \frac{Y_t^{\text{AI}}}{Y_t^{\text{no-AI}}} - 1 = \left(\frac{m_{\text{AI}}}{h}\right)^{aL^*} - 1$$

This expression captures exactly the one-time productivity gains estimated by Acemoglu (2024) ($\Delta\text{TFP} = 0.71\%$ when $a = 5\%$) and Aghion and Bunel (2024) ($\Delta\text{TFP} = 7\%$ when $a = 30\%$).

Following these papers, I assume that the AI automation shock occurs in 2035, ten years from now. I then trace its consequences using my OLG framework, calculating the output ratio $Y_t^{\text{AI}}/Y_t^{\text{no-AI}}$ at $t = 2035, 2055, 2075, \ldots$ (since $T = 20$). This implies that $Y_{2025}^{\text{AI}}/Y_{2025}^{\text{no-AI}} = 1$, and that:

$$\frac{Y_t^{\text{AI}}}{Y_t^{\text{no-AI}}} = \frac{\bar{q}_t^{\text{AI}}}{\bar{q}_t^{\text{no-AI}}} \cdot (1 + \Delta\text{TFP}) = \left(1 - ax^{1/\theta}\right)^{\frac{\theta(t-2035)}{T}} \cdot (1 + \Delta\text{TFP}), \text{ for } t = 2035, 2055, \ldots$$

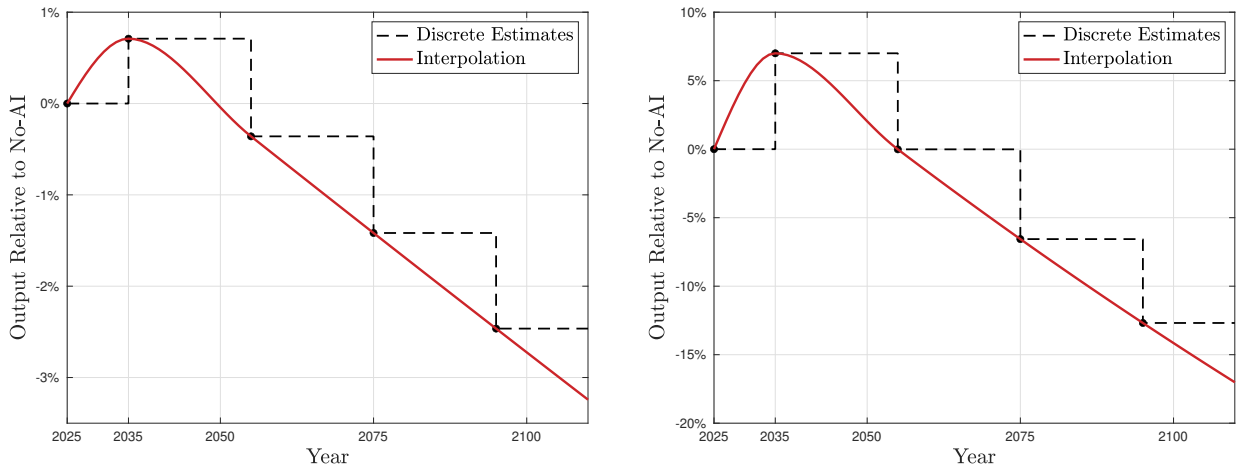
where the last equality follows because, for $t \geq 2035$:

$$\bar{q}_t^{\text{AI}} = \bar{q}_{2035} \cdot [(1-a)L^* + \nu]^{\frac{\theta(t-2035)}{T}} \text{ and } \bar{q}_t^{\text{no-AI}} = \bar{q}_{2035} \cdot (L^* + \nu)^{\frac{\theta(t-2035)}{T}}$$

I then render a smooth time path by interpolating the discrete OLG points with MATLAB's `interp1` function using the shape-preserving `pchip` method.

Figure 5 reports the results. It plots the percentage deviation of output under AI adoption from the no-AI baseline as a function of calendar time. Panel (a) corresponds to the conservative scenario of Acemoglu (2024), where the modest 0.71% gain from adoption is exhausted by 2050 and output falls 1.42% and 2.72% below the no-AI baseline by 2075 and 2100, respectively. Panel (b) depicts the more aggressive scenario of Aghion and Bunel (2024), in which the initial 7% gain disappears by 2055, with output falling 6.56% below the no-AI baseline by 2075 and 14.15% by the end of the century.[19]

[19]These calculations assume a one-to-one mapping between TFP and output, reflecting the baseline setup in which capital



(a) Scenario Inspired by Acemoglu (2024) ($a = 5\%$, $\Delta\text{TFP} = 0.71\%$)

(b) Scenario Inspired by Aghion and Bunel (2024) ($a = 30\%$, $\Delta\text{TFP} = 7\%$)

Figure 5: Quantitative Illustration - Short-run Gains vs. Long-Run Losses

Notes. Parameter values: $g_Y = 2\%$, $T = 20$, $\theta = 0.5$, and $x = 65\%$. The figure plots the percentage deviation of output under AI adoption from the no-AI baseline when the reduction in entry-level positions is permanent. Panel (a) shows the conservative scenario inspired by Acemoglu (2024) ($a = 5\%$, $\Delta\text{TFP} = 0.71\%$), where the initial gain of 0.71% in 2035 is exhausted by 2050 and output declines to 2.72% below baseline by 2100. Panel (b) shows the aggressive scenario inspired by Aghion and Bunel (2024) ($a = 30\%$, $\Delta\text{TFP} = 7\%$), where the initial 7% gain in 2035 vanishes by 2055 and output declines to 14.15% below baseline by 2100.

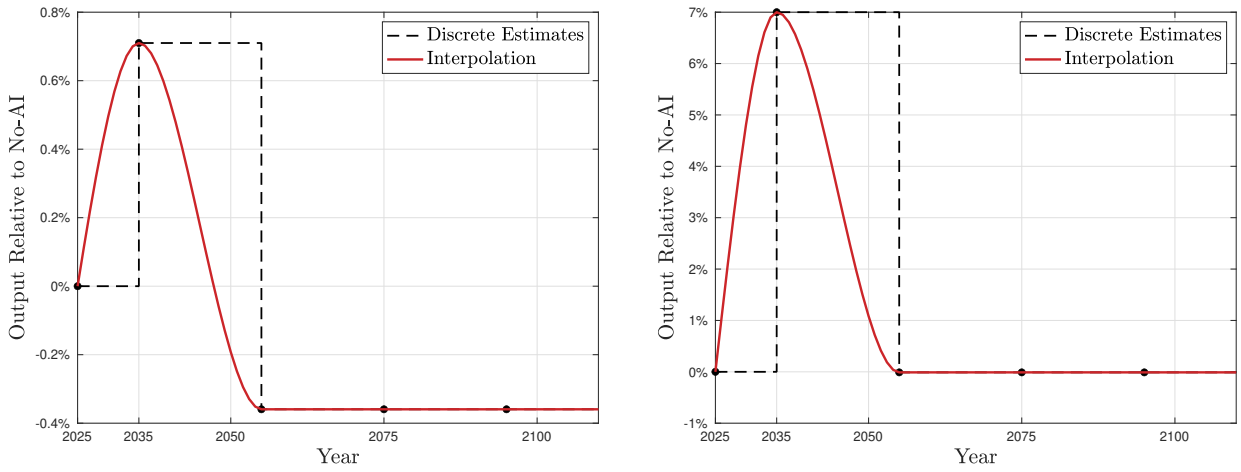
5.3 Temporary Disruption Only

The preceding analysis assumed that AI adoption permanently reduces entry-level positions. I now consider an alternative scenario in which the disruption is temporary: AI initially displaces novices, but new junior roles eventually emerge, restoring knowledge transmission. In this case, the long-run growth rate remains unchanged, but the period of reduced novice–expert interactions leaves the economy with a permanently smaller stock of tacit knowledge. The result is a permanent level effect on output that may offset the one-time productivity gains from AI adoption.

To illustrate this case, I adapt the exercise from the previous subsection, except that novices are displaced only during the first OLG period following the automation shock. In other words, entry-level positions contract between 2035 and 2055 but then return to their pre-AI level.

Figure 6 reports the quantitative results. In both scenarios, this temporary reduction in entry-level positions is sufficient to erase the initial productivity gains from AI adoption, leaving the economy

fully depreciates each period. By contrast, Acemoglu (2024) and Aghion and Bunel (2024) allow for capital deepening, so their TFP gains translate into more-than-proportional increases in output. For consistency, I adopt their TFP estimates. Moreover, as Acemoglu (2024) notes, once capital accumulation is introduced, consumer welfare depends on TFP rather than output, since the additional output reflects higher investment at the expense of consumption—strengthening the case for focusing on TFP in the baseline. For robustness, however, I repeat the exercise with capital deepening in the Online Appendix. The results change little: long-run output losses are only modestly smaller, and the break-even dates are only slightly later.



(a) Scenario Inspired by Acemoglu (2024) ($a = 5\%$, $\Delta\text{TFP} = 0.71\%$)

(b) Scenario Inspired by Aghion and Bunel (2024) ($a = 30\%$, $\Delta\text{TFP} = 7\%$)

Figure 6: Quantitative Illustration - Temporary Disruptions Only

Notes. Parameter values: $g^Y = 2\%$, $T = 20$, $\theta = 0.5$ and $x = 65\%$. The figure plots the percentage deviation of output under AI adoption from the no-AI baseline when the reduction in entry-level positions lasts only for one OLG period (2035–2055). Panel (a) shows the conservative scenario inspired by Acemoglu (2024) ($a = 5\%$, $\Delta\text{TFP} = 0.71\%$), where output stabilizes at -0.36% after 2055. Panel (b) shows the more aggressive scenario inspired by Aghion and Bunel (2024) ($a = 30\%$, $\Delta\text{TFP} = 7\%$), where the economy converges to -0.01% below baseline after 2055.

permanently below the no-AI baseline. In the conservative case inspired by Acemoglu (2024), output stabilizes at 0.36% below baseline after 2055. In the more aggressive case inspired by Aghion and Bunel (2024), the shortfall is negligible but still negative, with output converging to 0.01% below baseline after 2055.

5.4 Moving Beyond Automation: Decision-Support AI and Democratized Expertise

The results thus far indicate significant disruptions resulting from the extensive automation of entry-level tasks. Yet focusing solely on this dimension overlooks the broader transformative potential of AI. The next section expands the analysis to examine how the hybrid nature of AI-embedded knowledge not only democratizes tacit-like expertise but also reshapes the transmission of tacit knowledge.

6 AI Co-Pilots

As discussed in Section 2, AI co-pilots provide scalable access to sophisticated expertise previously attainable only through extensive hands-on experience. Hence, these systems have the potential to partially offset the loss of tacit knowledge caused by entry-level automation by assisting experts who, as novices, failed to acquire adequate skills.

However, as this section demonstrates, AI co-pilots are not always beneficial. By offering experts alternative access to tacit expertise, these systems may diminish novices' incentives for hands-

on learning, as novices anticipate relying on AI assistance in the future. This effect undermines human-driven knowledge transmission and may constrain growth through a mechanism distinct from the one discussed earlier: whereas the automation of entry-level tasks reduces the supply of apprenticeship-like arrangements offered by experts, AI co-pilots diminish novices' demand for them.

6.1 The Setting

Formal Description.— The model is identical to that of Section 3, except that experts now have the option to augment their skills using an AI co-pilot. However, because the reasoning behind AI recommendations remains opaque, novices cannot indirectly acquire this knowledge by interacting with experts who use these systems. Consequently, subsequent generations must independently seek the co-pilot's assistance to benefit from AI-generated expertise.

Formally, an expert with skill $q_{i,t}$ can access an AI co-pilot with skill level z_{AI} at a cost $\zeta \geq 0$, increasing her effective skill to $\max\{q_{i,t}, z_{\text{AI}}\}$. For simplicity, I assume that z_{AI} remains constant over time, though I later speculate on the case in which it grows exogenously. In either case, I abstract from modeling the underlying training process of AI systems, as this would require taking a stand on several dimensions—such as the resources, technology, and data required—that lie beyond the scope of this paper.[20]

As noted above, novices learn only from the expert's original skill $q_{i,t}$, not from the AI-augmented skill $\max\{q_{i,t}, z_{\text{AI}}\}$. An equivalent interpretation of this assumption is that novices interacting with AI-assisted experts merely learn that they "should consult the AI." I later consider the case in which AI's tacit-like knowledge is not opaque, allowing novices to acquire experts' AI-augmented skills.

Given that expert income is linear in skill, experts decide whether to employ the co-pilot following a threshold rule: experts with skills $q_{i,t} \leq \phi_{\text{AI}} \equiv z_{\text{AI}} - \zeta$ use it, while those with higher skills do not. I assume $\phi_{\text{AI}} > q^{\min}$ to ensure that at least some experts use the system. Thus, an expert's optimal income is given by $I^*_{i,t} = \max\{q_{i,t}, \phi_{\text{AI}}\} - C(\cdot)$, where $C(\cdot)$—which is independent of $q_{i,t}$—is the costs of hiring novices or renting machines to perform the routine tasks required for production (see Section 3).

This modification alters novices' expected returns from apprenticeship-based learning. Specifically, a novice's expected income at time $t+1$ conditional on supplying l_t units of labor at time t is now given by:

$$\mathbb{E}_t[I^*_{i,t+1} \mid l_t] = q^{\min} + e^{-u(k_t(l_t+\nu)^\theta)}(\phi_{\text{AI}} - q^{\min}) + k_t(l_t+\nu)^\theta \gamma(1-\theta, u(k_t(l_t+\nu)^\theta)) - C(\cdot)$$

$$\text{where } u(k) \equiv \left(\frac{\phi_{\text{AI}} - q^{\min}}{k}\right)^{-\frac{1}{\theta}} \text{ and } \gamma(s,u) \equiv \int_0^u x^{s-1}e^{-x}dx$$

[20] One potential approach is to model AI as an infinitely lived agent that gradually learns from its interactions with humans. Exploring this extension is an interesting avenue for future research.

Accordingly, the novice optimization problem at time t becomes:

$$\max_{l_t \in \mathbb{R}_+} \Big\{ w_t l_t - \frac{\mu}{2} l_t^2 + \beta \Big[q^{\min} + e^{-u(k_t (l_t+\nu)^\theta)} (\phi_{\text{AI}} - q^{\min}) + k_t (l_t+\nu)^\theta \gamma(1-\theta, u(k_t(l_t+\nu)^\theta)) - C(\cdot) \Big] \Big\} \tag{6}$$

The competitive equilibrium is defined exactly as in Section 3, except that novices' labor supply now follows equation (6) rather than (3). Since the tacit-like expertise provided by AI cannot be transmitted among humans, the dynamics of tacit knowledge accumulation remain unchanged, with $k_{t+1}^* = (L_t^* + \nu)^\theta k_t^*$. Nonetheless, AI co-pilots alter the effective distribution of expert skills, modifying aggregate output (gross of AI costs) at time t to:

$$Y_t^* = q^{\min} + e^{-u(k_t^*)} (z_{\text{AI}} - q^{\min}) + k_t^* \gamma(1-\theta, u(k_t^*))$$

A Brief Discussion.— This approach of modeling AI co-pilots is inspired by Ide and Talamàs (2025a). The access cost, ζ, represents the computational resources—or "inference compute"—required to use the system. Moreover, the assumption that experts who employ AI have their skills enhanced to $\max\{q_{i,t}, z_{\text{AI}}\}$ captures the intuitive idea that experts whose tacit knowledge already surpasses the AI's capability gain no additional insights from it, having internalized comparable expertise. This also aligns with the experimental evidence presented in Section 2, which highlights that AI co-pilots predominantly benefit less-skilled individuals.

The related assumption that only experts—but not novices—can use AI co-pilots is consistent with Autor's (2024) argument that individuals need "necessary foundational training" to meaningfully interpret and apply AI-generated recommendations. Although the prospect of fully autonomous "AI agents" substituting for human experts is intriguing, current technology cannot yet replicate the complete role of an expert—hiring, coordinating, and supervising novices or other machines to produce independently. Whether such capabilities will materialize in the coming years remains uncertain, and their analysis is left for future research.

Finally, the assumption that experts can readily use AI-generated insights while their reasoning remains opaque aligns with the discussion of scalability and interpretability in Section 2. As noted there, AI-embedded expertise can be easily duplicated, transferred, and scaled using computational resources, facilitating a broad democratization of sophisticated skills. However, the fact that AI encodes tacit-like knowledge into its complex internal structure—referred to as "model weights"—renders this knowledge largely opaque and inaccessible to human interpretation (Dell'Acqua et al., 2023; Ide and Talamàs, 2025a).

Limitations.— The model aims to examine how AI-embedded knowledge democratizes tacit expertise and reshapes its intergenerational transmission, while abstracting from additional channels. Although potentially important, these channels are unlikely to overturn its main qualitative insight: novices' incentives for hands-on learning weaken once they know they can rely on AI in the future.

First, the analysis does not consider AI as a *learning-enhancing tool*. In practice, AI co-pilots may help experts articulate tacit knowledge more clearly and enable novices to interpret expert behavior

more accurately. One way to represent this is to assume that, with such systems, each routine task exposes novices to multiple experts, thereby increasing the return to apprenticeship-based learning. The trade-off is that while AI could facilitate knowledge diffusion, it may also weaken incentives for hands-on learning if novices expect to rely on AI once they become experts.

Second, the model does not consider AI as a *creativity-enhancing tool*. Beyond transmitting existing expertise, AI can support the generation, refinement, and iteration of new ideas, thereby expanding innovative capacity. In the framework, this would correspond to raising the parameter ν, which governs the quality of new ideas. Incorporating this channel would likely amplify the contribution of innovation to growth and partly offset the losses from reduced tacit knowledge transmission.

6.2 The Impact of AI Co-Pilots

Having established the formal model, I now turn to the implications of AI co-pilots for long-run economic outcomes. As the following proposition shows, AI co-pilots raise the baseline skill level, thereby partially offsetting the negative effects of the Learning Breakdown (LB) regime. However, these systems may also crowd out essential hands-on learning, potentially leading the economy into an equilibrium with zero long-run growth for a broader range of parameter values.

Proposition 3. *Fix an initial knowledge stock $k_0 > 0$. Recall that $\underline{\rho} \equiv c(1-\nu)$ and $\bar{\rho} \equiv cN$, and define $k^\dagger_{\mathrm{AI}}$ as the unique solution to:*

$$\rho - (c+\mu)(1-\nu) + \beta\theta k\gamma(1-\theta, u(k)) = 0$$

so it satisfies $k^\dagger_{\mathrm{AI}} > k^\dagger$ whenever $k^\dagger > 0$. Except for the knife-edge parameter values described in Remarks 1 and 2 below, the competitive equilibrium converges to one of the following stable long-run regimes:

1. **Mitigated Learning Breakdown (MLB).** *If $\rho \leq \underline{\rho}$ or $k_0 < k^\dagger_{\mathrm{AI}}$, the economy converges asymptotically to the MLB regime. This regime is identical to the LB regime of Proposition 1 except that long-run output converges to $Y_t^* \to z_{\mathrm{AI}}$, rather than $Y_t^* \to q^{\min}$.*

2. **Constrained Learning (CL).** *If $\underline{\rho} < \rho < \bar{\rho}$ and $k_0 > k^\dagger_{\mathrm{AI}}$, the economy reaches in finite time the CL regime described in Proposition 1.*

3. **Full Learning (FL).** *If $\rho \geq \bar{\rho}$ and $k_0 > k^\dagger_{\mathrm{AI}}$, the economy reaches in finite time the FL regime described in Proposition 1.*

Remark 1 (Knife-Edge Stable Equilibrium). *When $\rho = \underline{\rho}$ and $k_0 \geq k^\dagger_{\mathrm{AI}}$, the economy immediately (at $t = 0$) reaches a stable steady state characterized by $L_t^* = 1-\nu$, $w_t^* = 0$, and $k_t^* = k_0$ for all $t \geq 0$.*

Remark 2 (Unstable Equilibrium). *When $\underline{\rho} < \rho$ and $k_0 = k^\dagger_{\mathrm{AI}}$, the economy immediately reaches an unstable steady state characterized by $L_t^* = 1-\nu$, $w_t^* = \rho - c(1-\nu)$, and $k_t^* = k^\dagger_{\mathrm{AI}}$ for all $t \geq 0$.*

Proof. See Appendix B. □

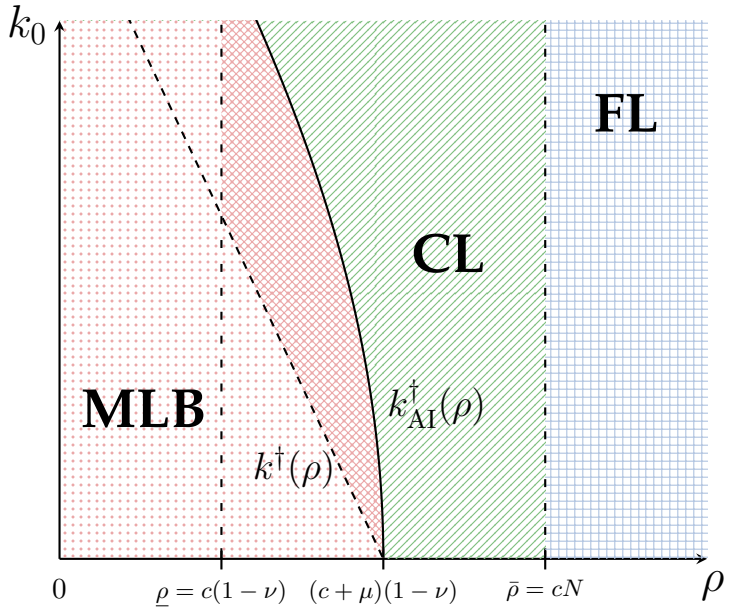


Figure 7: Long-Run Equilibrium Outcomes with AI Co-pilots as a Function of (ρ, k_0)

Notes. Parameter Values: $N = 3$, $\theta = 0.5$, $c = 1$, $\mu = 1$, $\nu = 0$, and $\beta = 0.96^{20}$, $q^{\min} = 3.6$, $\phi_{\text{AI}} = 12$. The region with the **dotted pattern** corresponds to combinations of (ρ, k_0) for which the economy converges to the Mitigated Learning Breakdown (MLB) steady state; the region with **northeast diagonal lines** indicates convergence to the Constrained Learning (CL) steady state; and the region with the **grid pattern** represents convergence to the Full Learning (FL) steady state. Finally, the combined **dotted and cross-hatch pattern** highlights parameter combinations that converge to the CL regime without AI but shift to the MLB regime when AI is introduced.

Proposition 3 is illustrated in Figure 7, which shows the economy's long-run equilibrium outcomes as a function of the cost of machines that perform routine tasks (ρ) and the initial stock of tacit knowledge (k_0). Comparing Figure 7 with Figure 1—which depicts the analogous long-run outcomes without AI co-pilots—reveals two key differences.

First, the Mitigated Learning Breakdown (MLB) regime replaces the Learning Breakdown (LB) regime. These two regimes differ only in their long-run equilibrium output: whereas output converges to the minimal skill level $q^{\min}$ in the LB regime, it converges instead to the higher, AI-provided skill level z_{AI} in the MLB regime. In this sense, AI co-pilots can partially offset the disruptions caused by entry-level automation by democratizing access to expertise.

Second, Figure 7 identifies a distinct region—indicated by a combined dotted and cross-hatched pattern—where adopting AI co-pilots shifts the economy from the Constrained Learning (CL) regime, characterized by sustained knowledge growth, into the MLB regime, characterized by persistent stagnation. This finding highlights that introducing AI co-pilots need not be unambiguously beneficial. Crucially, this region emerges through a mechanism that is fundamentally different from the one discussed in Section 4.

Indeed, as discussed previously, entry-level automation reduces novice-expert interactions by decreasing the supply of apprenticeship opportunities, as experts increasingly rely on machines rather

than novices for routine tasks. In contrast, AI co-pilots reduce novices' demand for apprenticeships, as novices anticipate future reliance on AI to complement their skills. Though distinct, both mechanisms weaken interpersonal knowledge transfer, progressively eroding society's accumulated stock of tacit human knowledge.

The mechanism underlying this demand-side effect warrants clarification. The primary reason AI co-pilots hinder intergenerational knowledge transmission is the opacity of their recommendations. If these systems were fully interpretable—so that novices could learn AI-generated insights simply by interacting with human experts who use them—the outcome would differ substantially. In that case, only the first generation of experts would ever consult AI. Subsequent cohorts would either surpass AI in knowledge—by learning from experts more skilled than AI—or directly acquire the AI's capabilities through exposure to AI-assisted experts. As a result, AI co-pilots would no longer weaken novices' incentives for hands-on learning.

In a sense, introducing interpretable AI co-pilots would resemble shifting the location parameter of the expert skill distribution from $q^{\min}$ to z_{AI} in the baseline model of Section 3.[21] Consequently, a direct policy implication of this analysis is that policymakers should actively promote the interpretability of AI systems—defined as their ability to articulate reasoning processes in ways understandable to humans (Doshi-Velez and Kim, 2017).[22]

Finally, up to this point, I have assumed that z_{AI} is constant over time. Under this assumption, opaque co-pilots expand the range of parameter values for which the economy converges to a Mitigated Learning Breakdown regime. One might therefore conclude that these systems are innocuous for economies with a sufficiently large stock of tacit knowledge, which could continue transmitting expertise across generations and sustaining growth as before. This conclusion, however, overlooks the possibility that z_{AI} could itself grow over time.

In that case, one can envision a scenario where even an economy with a high initial stock of knowledge eventually abandons hands-on learning once AI becomes sufficiently capable. From then on, output growth would depend solely on progress in z_{AI}, rather than on the transmission of human expertise.

If the AI frontier consistently grew faster than human transmission, AI would simply become the superior engine of growth. However, if it ultimately grew more slowly than tacit knowledge would have under continued transmission, the economy would be worse off in the long run than one without co-pilots. This risk would be particularly acute if z_{AI} rose rapidly at first—catching up with the stock of tacit knowledge and crowding out human learning—but later slowed or plateaued, leaving the economy locked into a weaker trajectory.

[21] The effects would resemble an increase in the location parameter, though they would not be mathematically equivalent. Consulting with AI transforms the initial Fréchet distribution of expert skills into a non-standard distribution with a point mass at z_{AI}, a feature that breaks the Fréchet structure for all subsequent periods.

[22] Note, however, that AI providers have incentives to maintain opacity. As explained above, enhanced interpretability would reduce long-term dependence on the system, potentially lowering recurring revenues for providers.

7 Final Remarks

In this paper, I develop a model to examine how advanced automation influences the intergenerational transmission of tacit knowledge. The analysis shows that the competitive equilibrium features socially excessive entry-level automation and reveals a critical trade-off: while improvements in automation deliver immediate productivity gains, they can weaken long-term economic growth by undermining the skills of future generations. Back-of-the-envelope calculations suggest AI-driven entry-level automation could lower the long-run annual growth rate of U.S. per capita output by 0.05 to 0.35 percentage points, depending on its scale. While AI co-pilots, which offer scalable access to specialized expertise, may partially mitigate these negative effects, they also risk discouraging younger workers from engaging in hands-on learning.

Taken together, these findings challenge the view that AI will automatically sustain productivity growth, suggesting instead that fully harnessing AI's potential requires safeguarding or creating new entry-level opportunities. Without such entry-level roles, AI-driven innovation would have to accelerate substantially—likely beyond current expectations—for long-term growth to match the revolutionary predictions of many.

A more radical possibility is that AI becomes sufficiently transformative to render existing forms of tacit knowledge obsolete, fundamentally reshaping the skills required by future generations. In this scenario, productivity growth would increasingly depend on developing new forms of tacit expertise—such as effectively prompting and interacting with AI systems. Yet, it is unlikely that current tacit knowledge would become entirely irrelevant. Human judgment, rooted in personal experience, would remain essential for assessing when AI performance is "good enough," especially given the opacity and occasional inaccuracies ("hallucinations") characteristic of today's AI systems.

Given the high stakes and considerable uncertainty surrounding AI's impact, policymakers face a delicate balancing act: promoting AI adoption while ensuring the preservation of essential entry-level opportunities. Effective policy responses might include targeted subsidies for mentorship and training programs, taxes designed to discourage excessive entry-level automation, explicit incentives for developing interpretable AI systems, and support for AI tools that complement—rather than replace—human labor. Without deliberate interventions, the subtle yet significant erosion of tacit knowledge could substantially diminish AI's long-term economic potential.

APPENDIX

A The Baseline Model: The Competitive Equilibrium

A.1 Existence and Uniqueness of a Competitive Equilibrium

Recall that $\omega(k_t)$ denotes the (potentially negative) wage that solves $L^d(\omega(k_t), \rho) = L^s(\omega(k_t), k_t)$. To establish the existence and uniqueness of a competitive equilibrium, it suffices to show that, for any

given $k_t > 0$, such a wage $\omega(k_t)$ exists and is unique.

Existence of the competitive equilibrium then follows directly from the existence of $\omega(k_t)$, as this ensures that the three defining conditions of a competitive equilibrium are well-defined and satisfied. Uniqueness also follows directly: for any given knowledge stock k_t^*, there is precisely one wage-labor pair (w_t^*, L_t^*) consistent with equilibrium. Since knowledge evolves deterministically, it follows that, from any initial knowledge stock $k_0 > 0$, the economy traces out a unique deterministic trajectory $\{w_t^*, k_t^*, L_t^*\}_{t\geq 0}$.

To formally establish the existence and uniqueness of $\omega(k_t)$, fix any $k_t > 0$. Observe that labor demand $L^d(w_t, \rho)$ is continuous, weakly decreasing in w_t, bounded between 0 and N, and satisfies $\lim_{w_t \to -\infty} L^d(w_t, \rho) = N$. Meanwhile, labor supply $L^s(w_t, k_t)$ is strictly increasing in w_t, and satisfies $\lim_{w_t \to -\infty} L^s(w_t, k_t) = 0$ and $\lim_{w_t \to \infty} L^s(w_t, k_t) = \infty$. Thus, for each $k_t > 0$, there exists a unique real number $\omega(k_t)$ that satisfies $L^d(\omega(k_t), \rho) = L^s(\omega(k_t), k_t)$.

I conclude this section by showing two additional properties of $\omega(k_t)$: first, that $\omega(k_t)$ is strictly decreasing in k_t, and second, that $\lim_{k_t \to \infty} \omega(k_t) = -\infty$. Because $L^d(w_t, \rho)$ does not depend on k_t and $L^s(w_t, k_t)$ is strictly increasing in k_t, it follows directly that $\omega(k_t)$ is strictly decreasing in k_t. Moreover, because $L^d(w_t, \rho)$ is bounded and $\lim_{k_t \to \infty} L^s(w_t, k_t) = \infty$, if $\omega(k_t)$ were bounded below, say $\omega(k_t) \geq \underline{w}$, then for k_t large enough we would have $L^s(\underline{w}, k_t) > N \geq L^d(\omega(k_t), \rho)$, contradicting market clearing. Hence, since $\omega(k_t)$ is strictly decreasing in k_t and no lower bound exists, it must be that $\lim_{k_t \to \infty} \omega(k_t) = -\infty$. □

A.2 Proof of Proposition 1

Recall that the law of motion of knowledge in equilibrium is $k_{t+1}^* = (L_t^* + \nu)^\theta k_t^*$, with $k_0^* = k_0 > 0$. Given that $\theta \in (0, 1)$, it follows immediately that:

$$k_{t+1}^* > k_t^* \text{ if } L_t^* > 1 - \nu, \quad k_{t+1}^* = k_t^* \text{ if } L_t^* = 1 - \nu, \quad \text{and} \quad k_{t+1}^* < k_t^* \text{ if } L_t^* < 1 - \nu$$

Recall also that $\omega(k_t)$ is strictly decreasing in k_t, and that:

$$\underline{\rho} \equiv c(1 - \nu), \quad \bar{\rho} \equiv cN, \quad \text{and} \quad k^\dagger \equiv \frac{(c + \mu)(1 - \nu) - \rho}{\beta\, \theta\, \Gamma(1 - \theta)}.$$

With these definitions in place, I now proceed to prove Proposition 1. First, I characterize the economy's long-run equilibrium outcomes through the following five lemmas. Afterward, I describe the transition dynamics toward the steady state.

Lemma A.2.1. *If $\rho < \underline{\rho}$, then the competitive equilibrium converges asymptotically to the LB regime. In this regime, $L_t^* \to \rho/(c + \mu) < 1 - \nu$, $w_t^* \to \mu\rho/(c + \mu)$, and $k_t^* \to 0$, so $g_t^k \to 0$ and $g_t^Y \to 0$.*

Proof. Note first that since $\rho > 0$, then for $\rho < \underline{\rho} = c(1 - \nu)$ it must be that $\nu < 1$. With this in mind, notice that $L^d(w_t, \rho) \leq L^d(0, \rho) = \rho/c < 1 - \nu$, as $\rho < \underline{\rho} = c(1 - \nu)$. Because $L_t^* = L^d(w_t^*, \rho)$, it immediately follows that $L_t^* < 1 - \nu$ for all $t \geq 0$. Given that $k_{t+1}^* = (L_t^* + \nu)^\theta k_t^*$, with $k_0^* = k_0 > 0$,

it must be that $k_t^* \leq (1-\varepsilon)^{\theta t} k_0$ for some $\varepsilon > 0$. Therefore, $\lim_{t\to\infty} k_t^* = 0$, implying that both g_t^k and g_t^Y converge to zero. Moreover, since $\lim_{t\to\infty} k_t^* = 0$, it also follows that $\lim_{t\to\infty} L^s(w_t^*, k_t^*) = L^s(w_t^*, 0) = w_\infty^*/\mu$. Thus, asymptotically, the labor market clears at a strictly positive wage, with $L^s(w_\infty^*, 0) = L^d(w_\infty^*, \rho) = L_\infty^* = \rho/(c+\mu) < 1-\nu$ and $w_\infty^* = \mu L_\infty^*$. □

Lemma A.2.2. *If $\underline{\rho} < \rho < \bar{\rho}$, then the competitive equilibrium converges in finite time to a CL regime if $k_0 > k^\dagger$, and converges asymptotically to the LB regime if $k_0 < k^\dagger$. In the CL regime, $L_t^* = \rho/c \in (1, N)$, $w_t^* = 0$, $g_t^k = \left(\frac{\rho}{c} + \nu\right)^\theta - 1$, and $\lim_{t\to\infty} g_t^Y = \left(\frac{\rho}{c} + \nu\right)^\theta - 1$. The outcome in the LB regime is the same as in Lemma A.2.1.*

Proof. Consider first the case $\nu > 1$. Since $L_t^* \geq 0$ and $k_{t+1}^* = (L_t^* + \nu)^\theta k_t^*$, this implies that $k_t^* \geq (1+\varepsilon)^{\theta t} k_0$ for some $\varepsilon > 0$, so $\lim_{t\to\infty} k_t^* = \infty$, and hence $\lim_{t\to\infty} \omega(k_t^*) = -\infty$. Thus, there must exist a finite time $T \geq 0$ at which $\omega(k_T^*) \leq 0$. From time T onward, the economy permanently settles into the CL regime, characterized by $w_t^* = 0$, $L_t^* = \rho/c \in (1-\nu, N)$ (due to proportional rationing), and $g_t^k = \left(\frac{\rho}{c} + \nu\right)^\theta - 1 > 0$. The fact that $\lim_{t\to\infty} g_t^Y = \left(\frac{\rho}{c} + \nu\right)^\theta - 1$, then follows because $\lim_{t\to\infty} k_t^* = \infty$.

Now consider $\nu \leq 1$. Note then that $k_{t+1}^* = k_t^*$ if and only if $L_t^* = 1-\nu$. Additionally, recall from the definition of a competitive equilibrium that $L_t^* = L^d(w_t^*, \rho)$. Hence, if $L_t^* = 1-\nu$, it must be the case that $w_t^* = \rho - c(1-\nu) > 0$. Because the labor market clears at a strictly positive wage when $L_t^* = 1-\nu$, it follows that $L^d(\rho - c, \rho) = L^s(\rho - c, k_t^*) = 1-\nu$. Using the first-order condition from problem (3), I find that:

$$\rho - (c+\mu)(1-\nu) + \beta\theta k_t^* \Gamma(1-\theta) = 0 \iff k_t^* = k^\dagger \equiv \frac{(c+\mu)(1-\nu) - \rho}{\beta\,\theta\,\Gamma(1-\theta)}.$$

Thus, when $\underline{\rho} < \rho$, the only knowledge level at time t consistent with $L_t^* = 1-\nu$ is exactly $k_t^* = k^\dagger$.

Next, recall from Appendix A.1 that $\omega(k_t)$ is strictly decreasing in k_t, with $\lim_{k_t\to\infty} \omega(k_t) = -\infty$. Therefore, when $k_t^* > k^\dagger$, I have $w_t^* = \max\{0, \omega(k_t^*)\} < \rho - c(1-\nu)$ and $L_t^* > 1-\nu$, and thus $k_{t+1}^* > k_t^*$. Consequently, $w_{t+1}^* < w_t^*$ and $L_{t+1}^* > 1-\nu$. Therefore, whenever $k_0 > k^\dagger$, it follows that $L_t^* > 1-\nu$ for all $t \geq 0$. This implies that $k_t^* \geq (1+\varepsilon)^{\theta t} k_0$ for some $\varepsilon > 0$, so $\lim_{t\to\infty} k_t^* = \infty$. Thus, by the same argument given above, there exists a finite time T such that the economy permanently settles into the CL regime.

Conversely, if $k_t^* < k^\dagger$, then $w_t^* = \omega(k_t^*) > \rho - c(1-\nu)$, implying $L_t^* < 1-\nu$, which in turn ensures $k_{t+1}^* < k_t^* < k^\dagger$. However, if so, then $w_{t+1}^* > w_t^*$ and $L_{t+1}^* < 1-\nu$. Thus, if $k_0 < k^\dagger$, it follows that $L_t^* < 1-\nu$ for all $t \geq 0$. Following an argument identical to the one used in Lemma A.2.1, I conclude that $\lim_{t\to\infty} k_t^* = 0$, implying that the economy asymptotically converges to the LB regime described in Lemma A.2.1. □

Lemma A.2.3. *If $\bar{\rho} \leq \rho$, then the competitive equilibrium converges in finite time to a FL regime if $k_0 > k^\dagger$, and converges asymptotically to the LB regime if $k_0 < k^\dagger$. In the FL regime, $L_t^* = N$, $w_t^* = 0$, $g_t^k = (N+\nu)^\theta - 1$, and $\lim_{t\to\infty} g_t^Y = (N+\nu)^\theta - 1$. The outcome in the LB regime is the same as in Lemma A.2.1.*

Proof. The proof closely follows that of Lemma A.2.2. The only difference arises when iterating forward from an initial knowledge stock $k_0 > k^\dagger$. In this case, once the equilibrium wage reaches zero

at some finite time $T \geq 0$, equilibrium novice labor remains constant at $L_t^* = N \leq \rho/c$ for all periods $t \geq T$. □

Lemma A.2.4. *If $\underline{\rho} < \rho$ and $k_0 = k^\dagger$, the economy immediately (at $t = 0$) reaches an unstable equilibrium with $L_t^* = 1 - \nu$, $w_t^* = \rho - c(1 - \nu)$, and $k_t^* = k^\dagger$ for all $t \geq 0$.*

Proof. In the proof of Lemma A.2.2, I have already shown that if $\underline{\rho} < \rho$ and $k_t^* = k^\dagger$, the unique competitive equilibrium must feature $L_t^* = 1 - \nu$ and $w_t^* = \rho - c(1 - \nu) > 0$, which implies $k_{t+1}^* = k_t^*$. Thus, if the initial knowledge stock is exactly $k_0 = k^\dagger$, the economy immediately settles into the stationary equilibrium described in the statement of this lemma. However, this equilibrium is unstable, since any small deviation from $k_0 = k^\dagger$ induces the economy to converge toward either the LB, CL, or FL regimes, as established in Lemmas A.2.2-A.2.3. □

Lemma A.2.5. *If $\rho = \underline{\rho}$ and $k_0 \geq k^\dagger$, the economy converges immediately (at $t = 0$) to a stable but non-generic equilibrium with $L_t^* = 1 - \nu$, $w_t^* = 0$, and $k_t^* = k_0$ for all $t \geq 0$. In contrast, if $\rho = \underline{\rho}$ and $k_0 < k^\dagger$, the economy converges to the LB regime described in Lemma A.2.1.*

Proof. By an argument analogous to the one I provided in the proof of Lemma A.2.2, when $\rho = \underline{\rho}$ and $k_t^* = k^\dagger$, the equilibrium novice labor satisfies $L_t^* = L^d(0, \rho) = L^s(0, k_t^*) = 1 - \nu$, with equilibrium wage $w_t^* = 0$. As a result, when $k_t^* > k^\dagger$, the equilibrium wage remains at $w_t^* = 0$, and the market experiences rationing, i.e., $L_t^* = L^d(0, \rho) = 1 - \nu < L^s(0, k_t^*)$. Thus, the knowledge stock remains constant at $k_{t+1}^* = k_t^*$. Consequently, if $\rho = \underline{\rho}$ and $k_0 \geq k^\dagger$, the economy immediately converges to the steady-state equilibrium with $L_t^* = 1 - \nu$, $w_t^* = 0$, and $k_t^* = k_0$ for all $t \geq 0$.

Conversely, if $\rho = \underline{\rho}$ and $k_t^* < k^\dagger$, I have $w_t^* = \omega(k_t^*) > 0$ and thus $L_t^* < 1 - \nu$. Hence, $k_{t+1}^* < k_t^*$, so $L_{t+1}^* < 1 - \nu$. Thus, whenever $\rho = \underline{\rho}$ and $k_0 < k^\dagger$, then $L_t^* < 1 - \nu$ for all $t \geq 0$. Following an identical reasoning as in the proof of Lemma A.2.1, I have that $\lim_{t\to\infty} k_t^* = 0$, so the economy asymptotically converges to the LB regime described in Lemma A.2.1. □

Finally, I characterize the transition dynamics to the steady state:

Lemma A.2.6.

- *In the transition to the LB regime, k_t^* strictly decreases in time, L_t^* weakly decreases in time, and w_t^* weakly increases in time.*
- *In the transition to the CL or FL regime, k_t^* and L_t^* strictly increase in time, while w_t^* strictly decreases in time.*

Proof. First, consider the transition to the LB regime. In this transition, $k_t^* > k_{t+1}^*$ for all $t \geq 0$ (by Lemmas A.2.1–A.2.3 and A.2.5). Since I previously established that the function $\omega(k_t)$ is strictly decreasing in k_t, it immediately follows that $\omega(k_t^*) < \omega(k_{t+1}^*)$. Given that $w_t^* = \max\{0, \omega(k_t^*)\}$, this implies that equilibrium wages w_t^* are weakly increasing over time. Furthermore, since $L_t^* = L^d(w_t^*, \rho)$ for all $t \geq 0$, and $L^d(w_t, \rho)$ is strictly decreasing in w_t, it follows that L_t^* is weakly decreasing in time.

Next, consider the transition to the CL or FL regime. By definition, during this transition, equilibrium wages are strictly positive ($w_t^* = \omega(k_t^*) > 0$) for all $t \geq 0$, as the economy has not yet reached a steady state. Furthermore, knowledge strictly increases over time ($k_t^* < k_{t+1}^*$ for all $t \geq 0$, by Lemmas A.2.1–A.2.3). Because $\omega(k_t)$ is strictly decreasing in k_t, it follows directly that equilibrium wages w_t^* are strictly decreasing over time. And since $L_t^* = L^d(w_t^*, \rho)$ for all $t \geq 0$, and $L^d(w_t, \rho)$ is strictly decreasing in w_t, it follows that L_t^* is strictly increasing in time. □

A.3 Proof of Proposition 2

To start, note that:

$$\begin{aligned}(7)\quad \Delta\mathcal{W}_t^* &\equiv \mathcal{W}(L_t^*(\rho'), k_t^*(\rho'); \rho') - \mathcal{W}(L_t^*(\rho), k_t^*(\rho); \rho) \\ &= \Gamma(1-\theta)[k_t^*(\rho') - k_t^*(\rho)] + \rho\left[N - L_t^*(\rho)\right] - \rho'\left[N - L_t^*(\rho')\right] + \left(\frac{c+\mu}{2}\right)\left(L_t^*(\rho)^2 - L_t^*(\rho')^2\right)\end{aligned}$$

where $L_t^*(\rho)$ and $k_t^*(\rho)$ denote the equilibrium labor and knowledge at time t given machine cost ρ. Also, for future reference, I denote by $w_t^*(\rho)$ the equilibrium wage at time t given machine cost ρ. Finally, recall that I restrict attention to cases where $\rho' < \bar{\rho}$, since otherwise, automation improvements have no effect—machines would remain unused.

With this established, I now proceed to prove Proposition 2 using the following two lemmas:

Lemma A.3.1. *Suppose the economy is initially in the CL or FL regime at time $\tau - 1$. Then, a permanent reduction in machine costs at the beginning of time τ immediately increases welfare but eventually reduces it in the long run:*

$$\mathcal{W}_\tau^*(\rho') > \mathcal{W}_\tau^*(\rho) \quad \textit{and} \quad \exists\, T \in (\tau, \infty) \ \textit{s.t.}\ \mathcal{W}_t^*(\rho') < \mathcal{W}_t^*(\rho),\ \forall\, t \geq T.$$

Proof. I first establish that $\Delta\mathcal{W}_\tau^* > 0$. Because the reduction in the cost of machines occurs at the beginning of period τ, I have that $k_\tau^*(\rho) = k_\tau^*(\rho')$. Hence,

$$\Delta\mathcal{W}_\tau^* = \rho\left[N - L_\tau^*(\rho)\right] - \rho'\left[N - L_\tau^*(\rho')\right] + \left(\frac{c+\mu}{2}\right)\left(L_\tau^*(\rho)^2 - L_\tau^*(\rho')^2\right)$$

Now, recall that $\omega(k_t, \rho)$ denotes the (potentially negative) wage that solves $L^d(\omega(k_t, \rho), \rho) = L^s(\omega(k_t, \rho), k_t)$, where I am making explicit that $\omega(k_t, \rho)$ depends on ρ. Because $L^d(w_t, \rho)$ is strictly increasing in ρ for any given w_t, it follows that $\omega(k_t, \rho)$ is also strictly increasing in ρ.

Note then that if the economy is initially in the CL or FL regime at time $\tau - 1$, wages have already reached zero, i.e., $w_{\tau-1}^*(\rho) = 0$, which implies $w_\tau^*(\rho) = 0$ as well. Since $\omega(k_t, \rho)$ is strictly increasing in ρ, it follows that after the reduction in ρ, wages remains at the zero lower bound at time τ, i.e., $w_\tau^*(\rho') = \max\{0, \omega(k_\tau^*, \rho')\} = 0$. This is because $\omega(k_\tau^*, \rho') < \omega(k_\tau^*, \rho) \leq 0$, where the final inequality holds due to the earlier observation that $w_{\tau-1}^*(\rho) = w_\tau^*(\rho) = 0$. The fact that $w_\tau^*(\rho') = 0$ then implies that $L_\tau^*(\rho') = \rho'/c$ for all $\rho' < \bar{\rho}$.

Next, consider each initial regime separately. If the economy is initially in the FL regime at $\tau - 1$, then $L^*_\tau(\rho) = N$. Evaluating $\Delta\mathcal{W}^*_\tau$ when $L^*_\tau(\rho) = N$ and $L^*_\tau(\rho') = \rho'/c$ yields:

$$\Delta\mathcal{W}^*_\tau = \frac{(\bar{\rho} - \rho')^2}{2c} + \frac{\mu(\bar{\rho} - \rho')(\bar{\rho} + \rho')}{2c^2} > 0$$

Similarly, if the economy is initially in the CL regime, then $L^*_\tau(\rho) = \rho/c$. Evaluating $\Delta\mathcal{W}^*_\tau$ when $L^*_\tau(\rho) = \rho/c$ and $L^*_\tau(\rho') = \rho'/c$ gives:

$$\Delta\mathcal{W}^*_\tau = \frac{(\rho - \rho')}{c^2}\left[c\left(\bar{\rho} - \frac{\rho}{2} - \frac{\rho'}{2}\right) + \frac{\mu(\rho + \rho')}{2}\right] > 0$$

where the last inequality follows because $\rho' < \rho < \bar{\rho}$ in this case. Thus, in both the CL and FL regimes at time $\tau - 1$, it holds that $\Delta\mathcal{W}^*_\tau > 0$.

I now show that the permanent reduction in machine costs at the beginning of time τ eventually decreases welfare in the long run. To begin, note that the one-off welfare gain from cheaper machines is bounded above:

$$\rho\left[N - L^*_t(\rho)\right] - \rho'\left[N - L^*_t(\rho')\right] + \left(\frac{c+\mu}{2}\right)\left(L^*_t(\rho)^2 - L^*_t(\rho')^2\right) \leq \rho N + \left(\frac{c+\mu}{2}\right)N^2 \equiv B < \infty$$

It follows immediately that $\Delta\mathcal{W}^*_t \leq \Gamma(1-\theta)[k^*_t(\rho') - k^*_t(\rho)] + B$.

Next, observe that equilibrium novice labor remains permanently lower from period τ onward. Specifically, the original path involves a constant amount of novice labor $L^*_t(\rho) = L^{\text{old}} \in \{\rho/c, N\} > 1$ for all $t \geq \tau$, while under the reduced machine cost scenario, the path involves a sequence $\{L^*_t(\rho')\}_{t=\tau}^{t=\infty}$, where $L^*_t(\rho') \leq \rho'/c < L^{\text{old}}$ for all $t \geq \tau$ (note that the economy need not immediately settle into a steady state following the reduction in ρ).

Since knowledge evolves according to $k^*_{t+1} = (L^*_t + \nu)^\theta k^*_t$, the knowledge stocks under the two paths evolve as follows:

$$k^*_t(\rho') \leq k_\tau\left(\frac{\rho'}{c} + \nu\right)^{\theta(t-\tau)} \quad \text{and} \quad k^*_t(\rho) = k_\tau\left(L^{\text{old}} + \nu\right)^{\theta(t-\tau)}$$

Because $L^{\text{old}} > 1 - \nu$ and $L^{\text{old}} > \rho'/c$, the difference $k^*_t(\rho') - k^*_t(\rho)$ decreases strictly over time and becomes arbitrarily negative as $t \to \infty$. Given that $\Gamma(1-\theta) > 0$ and $B < \infty$, this implies there must exist a finite date $T \in (\tau, \infty)$ such that $\Gamma(1-\theta)[k^*_t(\rho') - k^*_t(\rho)] + B < 0$ for all $t \geq T$. Since this expression provides an upper bound on $\Delta\mathcal{W}^*_t$, the desired result immediately follows. □

Lemma A.3.2. *Suppose the economy is approximately in the LB regime at time $\tau - 1$. Then, a permanent reduction in machine costs at the beginning of time τ permanently improves welfare:*

$$\mathcal{W}^*_t(\rho') > \mathcal{W}^*_t(\rho), \quad \forall\, t \geq \tau.$$

Proof. Given that the economy is approximately in the LB regime at time $\tau - 1$, this implies that $\rho \leq (c+\mu)(1-\nu) \leq c + \mu$ (otherwise, the economy would have converged either to the CL or FL regime). Moreover, $k^*_t(\rho) < \varepsilon$ and $k^*_t(\rho') < \varepsilon$ for all $t \geq \tau$, where $\varepsilon > 0$ but arbitrarily small. This

also implies that $L_t^*(\rho) < \rho/(c+\mu) + \varepsilon$ and $L_t^*(\rho') < \rho'/(c+\mu) + \varepsilon$. Using this expressions on (7) and taking $\varepsilon \to 0$, yields:

$$\lim_{\varepsilon \to 0} \Delta \mathcal{W}_t^* = \frac{(\rho - \rho')}{2(c+\mu)} \left[2N(c+\mu) - \rho - \rho'\right] > (N-1)(\rho - \rho') > 0, \text{ for all } t \geq \tau$$

where the second-to-last equality follows because $\rho' < \rho \leq c + \mu$, and the last inequality follows because $N > 1$ and $\rho > \rho'$. □

B AI Co-Pilots

B.1 Existence and Properties of $k_{\text{AI}}^\dagger$

Recall that Proposition 3 defines $k_{\text{AI}}^\dagger$ as the unique solution to:

$$\rho - (c+\mu)(1-\nu) + \beta\theta k \gamma(1-\theta, u(k)) = 0 \tag{8}$$

and states that $k_{\text{AI}}^\dagger > k^\dagger$ whenever $k^\dagger > 0$. In this appendix, I demonstrate both the uniqueness of the solution to (8) and verify that $k_{\text{AI}}^\dagger > k^\dagger$ whenever $k^\dagger > 0$.

To establish uniqueness, first note that the left-hand side of (8) approaches $-\infty$ as $k \to -\infty$ and $+\infty$ as $k \to +\infty$. Moreover, it is strictly increasing in k, as its derivative is positive: $\beta\theta\gamma(1-\theta, u(k)) + \beta u(k)^{1-\theta} e^{-u(k)} > 0$. Thus, the left-hand side of (8) crosses zero exactly once, from below, ensuring uniqueness.

Next, to verify that $k_{\text{AI}}^\dagger > k^\dagger$ whenever $k^\dagger > 0$, note that $k^\dagger$ is the unique solution to:

$$\rho - (c+\mu)(1-\nu) + \beta\theta k \Gamma(1-\theta) = 0 \tag{9}$$

This implies that $k^\dagger = 0$ if and only if $\rho = (c+\mu)(1-\nu)$, and $k^\dagger > 0$ whenever $\rho < (c+\mu)(1-\nu)$. Similarly, from (8), I have that $k_{\text{AI}}^\dagger = 0$ if and only if $\rho = (c+\mu)(1-\nu)$, and $k_{\text{AI}}^\dagger > 0$ whenever $\rho < (c+\mu)(1-\nu)$. Therefore, it remains only to show that $k_{\text{AI}}^\dagger > k^\dagger$ for all $\rho < (c+\mu)(1-\nu)$.

To show this, note that for all $k > 0$, the left-hand side of (9) is strictly pointwise higher than the left-hand side of (8). This inequality holds since $\phi_{\text{AI}} > z_{\text{AI}}$ implies $u(k) < \infty$, so:

$$\Gamma(1-\theta) = \int_0^\infty x^\theta e^{-x} dx > \int_0^{u(k)} x^\theta e^{-x} dx = \gamma(1-\theta, u(k))$$

Given that the left-hand side of (8) crosses zero exactly once and from below, this immediately implies $k_{\text{AI}}^\dagger > k^\dagger$ whenever $\rho < (c+\mu)(1-\nu)$, as desired.

B.2 Proof of Proposition 3

The proof closely parallels that of Proposition 1; therefore, I outline only the key steps.

First, if $\rho < \underline{\rho}$, then $L_t^* < 1 - \nu$ for all $t \geq 0$, and thus $\lim_{t\to\infty} k_t^* = 0$ (given that $k_{t+1}^* = (L_t^* + \nu)^\theta k_t^*$). Consequently, the economy converges asymptotically to the MLB regime. This regime is identical to

the LB regime from Proposition 1, except that long-run output converges to $Y_t^* \to z_{\text{AI}}$, rather than $Y_t^* \to q^{\min}$.

Next, suppose that $\rho \geq \underline{\rho}$. Observe that $k_{t+1}^* = k_t^*$ if and only if $L_t^* = 1 - \nu$. Moreover, whenever $L_t^* = 1 - \nu$, it must be that $w_t^* = \rho - c(1-\nu) \geq 0$. Thus, the first-order condition from problem (6) implies that the unique knowledge level at time t consistent with $L_t^* = 1 - \nu$ is precisely $k_t^* = k_{\text{AI}}^\dagger$, where $k_{\text{AI}}^\dagger$ solves:

$$\rho - (c + \mu)(1 - \nu) + \beta\theta k_{\text{AI}}^\dagger \gamma(1 - \theta, u(k_{\text{AI}}^\dagger)) = 0$$

It then follows—by exactly the same logic as in the proof of Proposition 1—that whenever $k_0 > k_{\text{AI}}^\dagger$, then $\lim_{t\to\infty} k_t^* = \infty$, so the economy reaches either the CL or the FL regime in finite time. Conversely, whenever $k_0 < k_{\text{AI}}^\dagger$, then $\lim_{t\to\infty} k_t^* = 0$, and the economy converges asymptotatically to the MLB regime. □

References

ACEMOGLU, D. (2024): "The Simple Macroeconomics of AI," *Economic Policy*, 40, 13–58.

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

ACEMOGLU, D., F. KONG, AND P. RESTREPO (2024): "Tasks At Work: Comparative Advantage, Technology and Labor Demand," NBER Working Paper No. 32872.

ACEMOGLU, D. AND J. LOEBBING (2024): "Automation and Polarization," 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, 1488–1542.

——— (2019): "8. Artificial Intelligence, Automation, and Work," in *The Economics of Artificial Intelligence: An Agenda*, ed. by A. Agrawal, J. Gans, and A. Goldfarb, University of Chicago Press, 197–236.

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

AGHION, P. AND S. BUNEL (2024): "AI and Growth: Where Do We Stand," LSE Working Paper.

AGHION, P., B. F. JONES, AND C. I. JONES (2017): "Artificial Intelligence and Economic Growth," NBER Working Paper No. 23928.

ALTMAN, S. (2025): "The Gentle Singularity," June, 2025. Available at https://blog.samaltman.com/the-gentle-singularity (accessed June 15, 2025).

AMODEI, D. (2024): "Machines of Loving Grace: How AI Could Transform the World for the Better," October, 2024. Available at https://darioamodei.com/machines-of-loving-grace (accessed June 15, 2025).

ARELLANO-BOVER, J. (2024): "Career Consequences of Firm Heterogeneity for Young Workers: First Job and Firm Size," *Journal of Labor Economics*, 42, 549–589.

AUTOR, D. (2024): "Applying AI to Rebuild Middle Class Jobs," Tech. rep., National Bureau of Economic Research.

AUTOR, D. AND N. THOMPSON (2025): "Expertise," Forthcoming in the *Journal of the European Economic Association*.

AUTOR, D. H., F. LEVY, AND R. J. MURNANE (2003): "The Skill Content of Recent Technological Change: An Empirical Exploration*," *The Quarterly Journal of Economics*, 118, 1279–1333.

BEANE, M. (2019): "Shadow Learning: Building Robotic Surgical Skill When Approved Means Fail," *Administrative Science Quarterly*, 64, 87–123.

——— (2024a): "How AI Could Keep Young Workers From Getting the Skills They Need," *The Wall Street Journal*, July 26, 2024. Available at https://www.wsj.com/lifestyle/careers/ai-training-young-employees-ed08cedc (accessed July 20, 2025).

——— (2024b): *The Skill Code: How to Save Human Ability in an Age of Intelligent Machines*, New York: Harper Business.

BEANE, M. AND C. ANTHONY (2024): "Inverted Apprenticeship: How Senior Occupational Members Develop Practical Expertise and Preserve Their Position When New Technologies Arrive," *Organization Science*, 35, 405–431.

BECKER, G. S. (1964): *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*, University of Chicago press.

BECKER, J., N. RUSH, E. BARNES, AND D. REIN (2025): "Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity," ArXiv preprint arXiv:2507.09089.

BENHABIB, J., J. PERLA, AND C. TONETTI (2021): "Reconciling Models of Diffusion and Innovation: A Theory of the Productivity Distribution and Technology Frontier," *Econometrica*, 89, 2261–2301.

BERGER, P. G., W. CAI, L. QIU, AND C. XINYI SHEN (2024): "Employer and Employee Responses to Generative AI: Early Evidence," Available at SSRN 4874061.

BESIROGLU, T. AND M. HOBBHAHN (2022): "Trends in GPU Price-Performance," Tech. rep., Technical Report, EpochAI 2022.

BLOOMBERGTV (2025): "AI Boom, Entry-Level Bust: Why College Grads Are Struggling to Land Jobs," August, 9, 2025. Available at https://www.bloomberg.com/news/videos/2025-08-09/ai-boom-entry-level-bust-grads-struggle-to-land-jobs-video (accessed August 11, 2025).

BRYNJOLFSSON, E., B. CHANDAR, AND R. CHEN (2025a): "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence," Working Paper.

BRYNJOLFSSON, E., D. LI, AND L. RAYMOND (2025b): "Generative AI at Work," *The Quarterly Journal of Economics*, 140, 889–942.

BUERA, F. J. AND R. E. LUCAS (2018): "Idea Flows and Economic Growth," *Annual Review of Economics*, 10, 315–345.

BUERA, F. J. AND E. OBERFIELD (2020): "The Global Diffusion of Ideas," *Econometrica*, 88, 83–114.

CAICEDO, S., J. LUCAS, ROBERT E., AND E. ROSSI-HANSBERG (2019): "Learning, Career Paths, and the Distribution of Wages," *American Economic Journal: Macroeconomics*, 11, 49–88.

COWAN, R. AND D. FORAY (1997): "The Economics of Codification and the Diffusion of Knowledge," *Industrial and Corporate Change*, 6, 595–622.

DE LA CROIX, D., M. DOEPKE, AND J. MOKYR (2018): "Clans, Guilds, and Markets: Apprenticeship Institutions and Growth in the Preindustrial Economy," *The Quarterly Journal of Economics*, 133, 1–70.

DE SOUZA, G. (2025): "Artificial Intelligence in the Office and the Factory: Evidence from Administrative Software Registry Data," FRB of Chicago Working Paper No. 2025-11.

DELL'ACQUA, F., E. MCFOWLAND III, E. R. MOLLICK, H. LIFSHITZ-ASSAF, K. KELLOGG, S. RAJENDRAN, L. KRAYER, F. CANDELON, AND K. R. LAKHANI (2023): "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality," Harvard Business School Technology & Operations Mgt. Unit Working Paper No. 24-013.

DOSHI-VELEZ, F. AND B. KIM (2017): "Towards A Rigorous Science of Interpretable Machine Learning," ArXiv preprint arXiv:1702.08608.

EATON, J. AND S. KORTUM (1999): "International Technology Diffusion: Theory and Measurement," *International Economic Review*, 40, 537–570.

——— (2002): "Technology, Geography, and Trade," *Econometrica*, 70, 1741–1779.

ELOUNDOU, T., S. MANNING, P. MISHKIN, AND D. ROCK (2024): "GPTs are GPTs: Labor market impact potential of LLMs," *Science*, 384, 1306–1308.

FORAY, D. (2004): *Economics of Knowledge*, MIT press.

FUDENBERG, D. AND L. RAYO (2019): "Training and Effort Dynamics in Apprenticeship," *American Economic Review*, 109, 3780–3812.

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

——— (2025): "The AI Becker Problem," July, 23, 2025. Available at https://www.siliconcontinent.com/p/the-ai-becker-problem (accessed July 23, 2025).

GARICANO, L. AND L. RAYO (2017): "Relational Knowledge Transfers," *American Economic Review*, 107, 2695–2730.

——— (2025): "Training in the Age of AI: A Theory of Apprenticeship Viability," Working Paper.

GARICANO, L. AND E. ROSSI-HANSBERG (2004): "Inequality and the Organization of Knowledge," *American Economic Review Papers and Proceedings*, 94, 197–202.

——— (2006): "Organization and Inequality in a Knowledge Economy," *The Quarterly Journal of Economics*, 121, 1383–1435.

——— (2015): "Knowledge-Based Hierarchies: Using Organizations to Understand the Economy," *Annual Review of Economics*, 7, 1–30.

GOLDMAN SACHS (2023): "Generative AI: Hype, or Truly Transformative?" July 5, 2023. Available at https://www.goldmansachs.com/pdfs/insights/pages/top-of-mind/generative-ai-hype-or-truly-transformative/report.pdf (accessed November 30, 2024).

HUMPHRIES, J. (2006): "English Apprenticeship: A Neglected Factor in the First Industrial Revolution," in *The Economic Future in Historical Perspective*, British Academy.

IDE, E. AND E. TALAMÀS (2024): "The Turing Valley: How AI Capabilities Shape Labor Income," ArXiv preprint arXiv:2408.16443.

——— (2025a): "Artificial Intelligence in the Knowledge Economy," Forthcoming in the *Journal of Political Economy*.

——— (2025b): "The Impact of AI on Global Knowledge Work," Working Paper. Prepared for the Carnegie-Rochester-NYU Conference on Public Policy.

JAROSCH, G., E. OBERFIELD, AND E. ROSSI-HANSBERG (2021): "Learning From Coworkers," *Econometrica*, 89, 647–676.

JOHNSTON, A. AND C. MAKRIDIS (2025): "The Labor Market Effects of Generative AI: A Difference-in-Differences Analysis of AI Exposure," Available at SSRN 5375017.

JONES, C. I. (2016): "The Facts of Economic Growth," in *Handbook of Macroeconomics*, Elsevier, vol. 2, 3–69.

——— (2024): "The AI Dilemma: Growth versus Existential Risk," *American Economic Review: Insights*, 6, 575–90.

KAHN, L. B. (2010): "The Long-Term Labor Market Consequences of Graduating from College in a Bad Economy," *Labour Economics*, 17, 303–316.

KORINEK, A. AND D. SUH (2024): "Scenarios for the Transition to AGI," NBER Working Paper No. 32255.

KORTUM, S. S. (1997): "Research, Patenting, and Technological change," *Econometrica*, 1389–1419.

KOSMYNA, N., E. HAUPTMANN, Y. T. YUAN, J. SITU, X.-H. LIAO, A. V. BERESNITZKY, I. BRAUNSTEIN, AND P. MAES (2025): "Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task," ArXiv preprint arXiv:2506.08872.

LEVY, S. (2025): "Demis Hassabis Embraces the Future of Work in the Age of AI," *Wired*, June 4, 2025. Available at (accessed June 15, 2025).

LICHTINGER, G. AND S. M. H. MAASOUM (2025): "Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Résumé and Job Posting Data," Working Paper.

LUCAS, R. E. (2009): "Ideas and Growth," *Economica*, 76, 1–19.

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

MACKENZIE, D. AND G. SPINARDI (1995): "Tacit Knowledge, Weapons Design, and the Uninvention of Nuclear Weapons," *American Journal of Sociology*, 101, 44–99.

MCKINSEY & COMPANY (2023): "The Economic Potential of Generative AI: The Next Productivity Frontier," June, 2023. Available at https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-AI-the-next-productivity-frontier/ (accessed June 15, 2024).

MOKYR, J. (2019): *The Economics of Apprenticeship*, Cambridge University Press, 20–43.

MOLL, B., L. RACHEL, AND P. RESTREPO (2022): "Uneven Growth: Automation's Impact on Income and Wealth Inequality," *Econometrica*, 90, 2645–2683.

NOY, S. AND W. ZHANG (2023): "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence," *Science*, 381, 187–192.

OYER, P. (2006): "Initial Labor Market Conditions and Long-Term Outcomes for Economists," *Journal of Economic Perspectives*, 20, 143–160.

PENG, S., E. KALLIAMVAKOU, P. CIHON, AND M. DEMIRER (2023): "The Impact of AI on Developer Productivity: Evidence from GitHub Copilot," ArXiv preprint arXiv:2302.06590.

PERLA, J. AND C. TONETTI (2014): "Equilibrium Imitation and Growth," *Journal of Political Economy*, 122, 52–76.

PERLA, J., C. TONETTI, AND M. E. WAUGH (2021): "Equilibrium Technology Diffusion, Trade, and Growth," *American Economic Review*, 111, 73–128.

PIZZINELLI, C., A. J. PANTON, M. M. M. TAVARES, M. CAZZANIGA, AND L. LI (2023): "Labor Market Exposure to AI: Cross-Country Differences and Distributional Implications," IMF Working Paper WP/23/216.

POLANYI, M. (1966): *The Tacit Dimension*, Doubleday & Co.

ROOSE, K. (2025): "For Some Recent Graduates, the A.I. Job Apocalypse May Already Be Here," *The New York Times*, May 30, 2025. Available at https://www.nytimes.com/2025/05/30/technology/ai-jobs-college-graduates.html (accessed August 11, 2025).

SANTACREU, A. M. (2015): "Innovation, Diffusion, and Trade: Theory and Measurement," *Journal of Monetary Economics*, 75, 1–20.

STIGLER, S. M. (1977): "Fractional Order Statistics, With Applications," *Journal of the American Statistical Association*, 72, 544–550.

SVANBERG, M., W. LI, M. FLEMING, B. GOEHRING, AND N. THOMPSON (2024): "Beyond AI Exposure: Which Tasks are Cost-Effective to Automate with Computer Vision?" Available at SSRN 4700751.

THE JOURNAL PODCAST (2025): "AI Is Coming for Entry-Level Jobs," *The Journal - WSJ Podcast*, Monday, July 7, 2025. Available at https://www.wsj.com/podcasts/the-journal/ai-is-coming-for-entry-level-jobs/e9f9eb31-14ad-498d-94f3-11ce91e9c464 (accessed July 19, 2025).

THE WORLD ECONOMIC FORUM (2025): "The Future of Jobs Report 2025," Tech. rep., The World Economic Forum, January 7, 2025. Available at https://www.weforum.org/publications/the-future-of-jobs-report-2025/ (accessed July 27, 2025).

VANDEHEI, J. AND M. ALLEN (2025): "Behind the Curtain: A White-Collar Bloodbath," *Axios*, May 28, 2025. Available at https://www.axios.com/2025/05/28/ai-jobs-white-collar-unemployment-anthropic (accessed June 27, 2025).

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