Skill Mismatch over the Technology Life Cycle

Ashwini K. Agrawal   Daniel Kim   Prasanna Tambe*

March 16, 2026

Abstract

Despite the productivity gains associated with digital technologies, their diffusion across firms remains slow and uneven, and many firms continue to operate legacy systems with limited growth potential. A large literature attributes this pattern to skill mismatch between firms and workers, yet most existing studies treat mismatch as an exogenous barrier to adoption. This paper shows that skill mismatch instead evolves endogenously over the technology diffusion process. We develop a framework in which multidimensional skill mismatch—the misalignment between firms' skill requirements and workers' capabilities—changes as technologies move through their life cycle. To test these predictions, we construct granular measures of firm skill demand and worker skill supply using matched data on job postings and worker résumés. Our analysis yields three key findings. First, we document a systematic U-shaped pattern in mismatch over the technology life cycle: mismatch is high when new technologies are first adopted, declines as firms and workers adjust, and rises again as technologies mature into legacy systems. Second, we show that diffusion bottlenecks arise not only from gaps in technical expertise but also from shortages of complementary nontechnical capabilities, in particular managerial and coordination skills. Third, we demonstrate that skill mismatch helps explain the persistence of firms operating legacy systems with limited growth opportunities. Together, these findings highlight skill mismatch as a dynamic force shaping both the diffusion of new technologies and the persistence of older ones.

*Preliminary and Incomplete; do not circulate without authors' consent. Contact: Ashwini K. Agrawal, London School of Economics, e-mail: a.agrawal14@lse.ac.uk; Daniel Kim, University of Waterloo, e-mail: daniel.kim2@uwaterloo.ca; Prasanna Tambe, Wharton School, University of Pennsylvania, e-mail: tambe@wharton.upenn.edu. Kim gratefully acknowledges funding for this project provided by the CPA Ontario Centre for Sustainability Reporting and Performance Management (CSPM) at the University of Waterloo.

1 Introduction

Technological progress has accelerated in recent decades, yet the diffusion of new technologies across firms remains considerably slow and uneven (Kalyani et al., 2025). A growing literature points to the labor market as a central factor behind this diffusion gap. In particular, many studies argue that the adoption and productivity of frontier technologies is constrained by the limited supply of workers with complementary skills (Adão et al., 2024; Brunello and Wruuck, 2021; Acemoglu and Zilibotti, 2001). Consistent with this view, the business press is replete with accounts of firms' extraordinary—and often unsuccessful—efforts to recruit workers with technology-complementary capabilities.[1]

Yet despite the prominence of skill mismatch in discussions of technological change, we have a limited understanding of how mismatch itself arises and evolves as technologies diffuse. Much of the existing literature treats mismatch as exogenous, abstracting from firms' endogenous adjustments of skill requirements and workers' incentives to acquire new skills as technologies mature. As a consequence, we know little about how mismatch changes over the technology life cycle and how it may differ between firms adopting new technologies and those operating legacy systems. Moreover, prior empirical work typically relies on univariate measures of skill, obscuring the particular—and potentially multidimensional—skills that constrain diffusion. Identifying the specific skills demanded by firms and supplied by workers is inherently challenging, as many relevant capabilities are unobservable and difficult to infer from job descriptions. For example, a job posting requiring "Java programming," may reflect demand for a narrow technical task, broad problem-solving ability, and/or social skills needed for collaborative software development. Disentangling these dimensions is crucial for identifying which skills—and *combinations* of skills—impede technology diffusion.

In this paper, we study how multidimensional skill mismatch endogenously evolves and shapes the diffusion of information technologies. We develop a conceptual framework that

[1] ManpowerGroup's 2024 Global Talent Shortage survey reveals that 75% of employers worldwide report difficulty finding workers with appropriate skills (ManpowerGroup, 2024).

builds on the seminal work of Chari and Hopenhayn (1991), who provide a general equilibrium theory of technology diffusion and vintage human capital that has become a workhorse model in the literature. We incorporate insights from several different strands of the literature to extend their model to environments in which technologies require heterogeneous bundles of skills and workers accumulate those skills endogenously over the diffusion process.

The key departure is the introduction of multidimensional skill mismatch—–defined as the distance between the vector of skill weights required by firms and the vector of skill weights possessed by workers (Lazear, 2009). These skill vectors capture both technical and nontechnical dimensions of human capital that relate to technology adoption. Differences in skill requirements across firms stem from variation in the tasks embodied in the information technologies and complementary organizational structures that firms adopt (Acemoglu and Autor, 2011; Brynjolfsson and Hitt, 2000). We then use skill mismatch to microfound the productivity dynamics of workers that are otherwise assumed in the baseline model; we postulate that greater mismatch reduces worker productivity, consistent with the empirical evidence in Lise and Postel-Vinay (2020), Guvenen et al. (2020), and Coraggio et al. (2025). These model adaptations generate novel predictions about three key issues: (i) how multidimensional skill mismatch endogenously evolves over the technology life cycle, (ii) what skill attributes serve as technology diffusion bottlenecks, and (iii) how skill mismatch can help explain the prevalence of legacy technologies with limited growth opportunities.

To test these predictions, we construct a matched employer–employee dataset that links U.S. firms' online job postings with workers' online résumés. The job postings data contain detailed information on firm skill requirements while the résumé data allow us to observe the skills supplied by workers with comparable granularity. Linking the two data sources enables us to measure the skill attributes of new hires in close temporal proximity to firms' posted skill requirements. The job postings data also provide information that reflects firms' adoption of specific information technologies. We identify the year of inception for each of the technologies in our sample, and use these data to characterize the diffusion of information technologies across firms over time.

We harmonize the raw, unstructured skills text from job postings and résumés by applying large language models (LLMs) and standard natural language processing (NLP) techniques to construct multidimensional measures of skill demand and supply. Specifically, we use BERT (Bidirectional Encoder Representations from Transformers; Devlin et al. (2019)) in a zero-shot classification framework (Yin et al., 2019) to map each list of reported skills into a vector of weights across 25 broad skill categories, including managerial, financial, and technical competencies.[2] These deep learning tools (Dell, 2025) are well suited to our setting for several reasons. Most importantly, because LLMs are trained on a vast corpus of text, they are able to capture the semantic meaning of skill descriptions expressed by firms and workers more effectively than researcher-defined keyword dictionaries. For example, a skill requirement such as "Java programming" simultaneously conveys various skills that are difficult to encode using manually constructed taxonomies. LLM's, in contrast, are able to codify skill descriptions using high-dimensional embeddings that capture much (if not all) of the information that market participants associate with these terms in practice, thereby mitigating the limitations of ad-hoc skill measures for our analysis.

In our main empirical specifications, we measure skill mismatch as the Euclidean distance between firms' skill-demand vectors and the skill-supply vectors of newly hired workers.[3] This continuous measure provides an intuitive characterization of the extent to which newly hired workers possess skills that meet firms' requirements, with larger distances indicating greater mismatch. Because skill weights sum to one, this measure captures the degree to which workers' comparative advantages across skill dimensions align with firms' needs. Our vector-based approach allows us to compare skill mismatch consistently across firms and over time while incorporating a broad range of skill dimensions.

[2]The 25 skill categories are defined by the résumé data provider and are designed to parsimoniously and efficiently span the skill space observed in the data. As discussed later, our results are robust to alternative skill categorizations that have been used in the literature.

[3]We also present similar results using alternative distance metrics, such as Hellinger distances between skill vectors.

We document several empirical findings that support the predictions of the model. The first finding is that skill mismatch between firms and newly hired workers follows a U-shaped pattern over the technology life cycle: mismatch is high when new technologies are adopted, declines as firms and workers adjust, and rises again as technologies mature into legacy systems. Importantly, this process is slow, unfolding over several decades. Consistent with the model, mismatch is initially high because workers lack the specific combinations of skills required to work with new technologies. As adoption expands, workers acquire relevant skills and firms adjust their skill requirements, leading to closer alignment between skill supply and demand. Over time, however, as newer and more productive technologies emerge, workers have weaker incentives to join firms that continue to rely on older technologies, causing skill mismatch at these firms to increase. The empirical relationship between skill mismatch and technology diffusion that we document is new to the literature, and arises theoretically from relaxing the standard assumptions of exogenous, one-dimensional human capital. By modeling the joint, endogenous evolution of multidimensional skill mismatch and technology adoption, we show that mismatch itself can be explained by economic forces and that it is a quantitatively important constraint not only for new technologies but also for legacy systems.

Our second set of results identifies the specific skill attributes that serve as bottlenecks to information technology diffusion. We show that diffusion is constrained not by technical skills alone, but by shortages in combinations of technical and non-technical skills. We draw this conclusion from several complementary findings. First, we document a U-shaped pattern in skill mismatch across both technical occupations (such as engineers and data scientists) and non-technical occupations (such as managers and business analysts) that are central to the organizational changes that accompany IT adoption (Bresnahan et al., 2002; Bartel et al., 2007). Second, when examining mismatch along individual skill dimensions, we find that skill gaps are most pronounced among non-technical IT-complementary skills, particularly managerial and business capabilities. Moreover, these gaps appear alongside substantial shortfalls in technical skills such as advanced manufacturing and network security knowledge. By contrast, skills associated with tasks less central to IT diffusion (e.g., hospitality or public speaking) exhibit little mismatch and are relatively oversupplied on average. Taken together, these patterns indicate that skill mismatch reflects scarcities in technical and non-technical skills that underpin the effective use of information technologies, rather than narrow deficits in technical expertise alone.

Our third set of findings shows how skill mismatch contributes to the persistence of legacy technologies with limited growth opportunities. In the model, firms operating older technologies remain viable because they attract workers whose skill portfolios are poorly aligned with frontier systems. Skill mismatch at the technological frontier therefore allows firms deploying legacy systems to continue operating despite inferior performance. Consistent with this mechanism, we document that firms operating legacy systems exhibit lower market-to-book ratios than firms that adopt new technologies. Both types of firms are nevertheless prevalent, and continue to operate and hire workers despite the skill mismatch that they both face in the labor market. These results complement existing explanations for the persistence of legacy technologies, which otherwise emphasize adjustment costs, switching frictions, or technological lock-in. Our analysis shows that multidimensional skill mismatch provides a distinct economic mechanism that can explain the coexistence of frontier technologies and legacy systems despite vast differences in future growth prospects.

We conduct an extensive set of analyses to establish the robustness of our results. Our findings are robust to alternative skill taxonomies proposed in prior work, including those developed by Deming and Kahn (2018). We also show that our results are robust to alternative choices of job-posting and résumé sampling windows and consistent across different measures of technology vintage. We further show that the estimated patterns are even stronger when we replace BERT with alternative large language models, such as ModernBERT. Finally, we discuss several limitations of our approach. In particular, our methodology captures mismatch in workers' comparative advantages across skill dimensions rather than mismatch in absolute skill levels, which may represent an additional, empirically relevant dimension of skill mismatch.

The main contribution of this paper is to show how multidimensional skill mismatch endogenously evolves and shapes technology diffusion. We advance the literature on technological change and skill mismatch by relaxing the standard assumptions of exogenous mismatch and one-dimensional human capital. These innovations yield several new insights into labor market dynamics that are otherwise obscured in existing frameworks. First, skill mismatch evolves following a U-shaped pattern over the technology life cycle and is pronounced not only for firms adopting new technologies, but also for firms deploying legacy systems. Second, diffusion bottlenecks arise from shortages of specific combinations of nontechnical and technical skills rather than from purely technical skill gaps alone. Third, labor market segmentation induced by skill mismatch helps explain the persistence of legacy technologies despite their limited growth prospects.

A secondary contribution of our paper is methodological. We introduce a simple empirical approach for measuring multidimensional skill demand and supply using large language models applied to unstructured text from job postings and résumés. This approach allows us to construct granular measures of skill combinations that capture firms' job requirements and workers' capabilities at scale. Our methodology complements prior approaches based on univariate classifications or manually-curated taxonomies by characterizing human capital in settings where skill content is heterogeneous, difficult to observe, and continuously evolving.

The rest of the paper proceeds as follows. Section 2 discusses the related literature, Section 3 presents our theoretical framework, Section 4 details our empirical approach, Section 5 describes our data, Section 6 presents our findings, Section 7 presents robustness analysis, and Section 8 concludes.

2 Related Literature

A longstanding question in economics concerns the gap between the rapid pace of technological progress and the much slower diffusion of new technologies across firms and countries. A growing literature emphasizes labor market forces as central to technology adoption and productivity dynamics. A prominent contribution is Acemoglu and Zilibotti (2001), who show that productivity differences across countries reflect mismatches between the skill requirements of frontier technologies and the human capital of the local workforce. Related work by Adão et al. (2024) argues that the speed of technological adoption depends critically on the size of skill gaps between workers and new technologies. Together with practitioner surveys documenting skill shortages as a primary barrier to adoption (MIT Technology Review Insights, 2023), these studies provide compelling evidence of skill mismatch as a central determinant of technology diffusion and aggregate productivity.

A common feature of this literature, however, is that skill mismatch is taken as exogenous. Gaps between technological requirements and individual worker skills are treated as fixed constraints, rather than as endogenous outcomes that are continuously shaped by firms' and workers' responses to technological diffusion. This abstraction is well suited to studying the consequences of skill mismatch for adoption and productivity, but it leaves open a fundamental question: what determines the evolution of skill mismatch in the first place?

Understanding the forces that shape skill mismatch is crucial for assessing how mismatch dissipates or persists as technologies diffuse, thereby influencing long-run adoption patterns and the coexistence of new technologies with legacy systems. Absent a clear account of the mechanisms that shape mismatch, it is difficult to interpret observed skill shortages or to distinguish temporary adjustment frictions from persistent bottlenecks. Anecdotally, skill mismatch is not confined to firms adopting new technologies, as firms operating legacy technologies also frequently report difficulty hiring workers with desirable skills (Monaghan and Bass, 2020). The nature of skill mismatch facing these firms, however, may reflect fundamentally different economic forces. Finally, understanding the endogenous determination of skill mismatch is a prerequisite for evaluating policies aimed at workforce development, as effective policy design requires separating market-driven adjustment dynamics from exogenous skill constraints.[4]

Progress on these issues also requires understanding the multidimensional nature of skill gaps between firm requirements and worker capabilities. Empirical studies of technological change and skill mismatch have typically examined mismatch through a univariate lens, focusing on broad categorizations of skill or education (Deming, 2023). This abstraction is sufficient for existing work to explain aggregate patterns of technology diffusion and productivity, however, it is insufficient for identifying the specific skill attributes that explain technology differences across firms. For example, we have a limited understanding of the extent to which technology diffusion is constrained by technical skill gaps alone or by gaps in IT-complementary nontechnical skills such as managerial and organizational capabilities. Moreover, there is little established evidence on the *combinations* of these skills as diffusion bottlenecks, even though there are numerous accounts of hiring managers reporting difficulties in finding employees with the right "mix" of skills (OECD, 2024).

The relevance of multidimensional skill mismatch for technology diffusion is further underscored by a growing literature that studies multidimensional mismatch more broadly. Prominent theoretical contributions include Lindenlaub and Postel-Vinay (2023) and Moscarini (2001), who study worker sorting based on multidimensional skill attributes. On the empirical front, Baley et al. (2022) and Lise and Postel-Vinay (2020) examine multidimensional skill mismatch arising from job search frictions, while Guvenen et al. (2020) study mismatch driven by uncertainty about workers' own skills. These studies rely on occupation-level skill requirements, drawn from O*NET, which are well suited to analyzing the implications of mismatch for worker earnings, mobility, and productivity in frictional labor markets.

Our paper complements this literature by studying multidimensional skill mismatch in

[4]See Goldin and Katz (2008) for a comprehensive discussion of these issues in the context of technical change and the wage structure.

the context of technology diffusion and analyzing variation in mismatch across firms. Rather than taking mismatch as a consequence of search frictions or informational constraints, we model skill mismatch as an endogenous outcome of firms' technology choices and workers' skill acquisition decisions over the technology life cycle. Empirically, we depart from occupation-level abstractions by measuring skill demand and supply at the firm level using job postings and résumés, allowing us to capture substantial heterogeneity in skill requirements even within narrowly defined occupations. This approach enables us to study how multidimensional skill mismatch evolves as technologies diffuse and to identify the specific combinations of skills that act as bottlenecks to adoption—dimensions that are difficult to observe using standard occupational classifications.

To study multidimensional skill mismatch and technology diffusion, we develop a conceptual framework that builds on the classic theory of technology diffusion and vintage human capital developed by Chari and Hopenhayn (1991). Their general equilibrium framework provides a workhorse model of how human capital evolves endogenously as firms adopt successive waves of more productive technologies. In the original model, workers are differentiated along a single skill dimension—skilled versus unskilled—and newly hired workers enter firms with a common initial skill level. This parsimonious structure is well suited to explaining aggregate patterns of diffusion and human capital accumulation, but it abstracts from heterogeneity in the degree and composition of skill mismatch across firms and technologies.

We adapt this framework to allow for multidimensional skill mismatch between firms and newly hired workers, reflecting the rich, high-dimensional nature of human capital emphasized in contemporary discussions of technology adoption and skill gaps. In our model, firms differ in their skill requirements across multiple dimensions, and newly hired workers vary in the extent to which their skill portfolios align with these requirements. This extension allows mismatch to vary systematically across firms, technologies, and stages of the technology life cycle. The model thus generates new empirical predictions about the joint evolution of multidimensional skill mismatch and technology diffusion that go beyond the original scope of

Chari and Hopenhayn (1991) and speak directly to the central questions of this paper.

3 Theoretical Framework

We present our conceptual framework by first summarizing the relevant aspects of Chari and Hopenhayn (1991). The model is an overlapping-generations economy with an infinite horizon (time t = 0,1,2...), where workers of two types (young and old, each of mass 1) live for two periods. In each time period t, a new technology is introduced that is more productive than all previous vintages; a specific vintage is denoted by τ (which means that the technology was first introduced at time t-τ). Firms adopt a technology of a particular vintage and hire U unskilled workers (type u) and S skilled workers (type s), where skill is a one-dimensional measure of human capital that is specific to a given technology vintage. A firm's production function $Y = \gamma^{t-\tau} f(U, S)$ is constant returns to scale, strictly increasing and concave in the number of skilled workers, and the productivity parameter γ is greater than 1.

A key assumption in the model is that unskilled workers and skilled workers are complements: the marginal productivity of an unskilled worker increases in the number of skilled workers who are simultaneously working in the same firm (i.e. $\frac{\partial f_U}{\partial S_t} > 0$). Young workers, who begin life unskilled, choose to work at a firm with a particular technology vintage. After one period of experience, they become skilled workers by acquiring technology vintage-specific human capital. When old, these workers can either join a firm with the same technology or switch to a firm that uses a different technology, in which case the workers are now considered unskilled. Equilibrium requires that workers sort across firms to maximize lifetime utility, firms choose technologies and make hiring decisions to maximize profits, and markets clear. Chari and Hopenhayn (1991) prove the existence of a unique stationary equilibrium in which the distribution of skilled workers (s_τ) over technology vintages $\tau \in \{0,1,2...\}$ is single-peaked: human capital specialization in a given technology rises with vintage, peaks, and then declines to zero once the technology becomes obsolete and adoption ceases.

Our first adaptation is to introduce an intuitive and well-accepted definition of multidimensional skill mismatch to provide a richer characterization of skilled and unskilled workers in the baseline model. Specifically, we define skill mismatch between firms and workers as the discrepancy between firms' multidimensional skill requirements and workers' multidimensional skill attributes, each measured using skill weight vectors following Lazear (2009). This measure has precedence in the academic literature and the practitioner sphere (International Labour Organization, see https://www.ilo.org/skills-mismatches). In the context of our model, when a firm adopts information technology of a particular vintage, the firm designs a production system with complementary assets built around the technology (Brynjolfsson et al., 2002; Brynjolfsson and Hitt, 2000). Production entails the completion of various tasks; the skills required of workers to perform these tasks represent the firm's skill requirements. This connection between technology, tasks, and skill requirements is rooted in the task-based approach to human capital (Acemoglu and Autor, 2011).[5]

Formally, a firm that adopts information technology of vintage τ will have skill requirements for a worker of type $j \in (u,s)$ denoted by $F^{\tau,j} \equiv (F_1^{\tau,j}, F_2^{\tau,j} ... F_N^{\tau,j})$, where $0 \leq F_i^{\tau,j} \leq 1$ for N skills indexed by i = 1,2,...N and $\sum_{i=1}^{N} F_i^{\tau,j} = 1$. The magnitude of each skill requirement $F_i^{\tau,j}$ measures the relative weight of the skill among the total set of skill requirements of the firm. Workers of type $j \in (u,s)$ possess multidimensional skill attributes, denoted by $W^{\tau,j} \equiv (W_1^{\tau,j}, W_2^{\tau,j} ... W_N^{\tau,j})$ where $0 \leq W_i^{\tau,j} \leq 1$ for N skills indexed by i = 1,2,...N and $\sum_{i=1}^{N} W_i^{\tau,j} = 1$. The magnitude of each skill attribute $W_i^{\tau,j}$ measures the relative weight of the skill within a worker's skill portfolio. We assume that skills are imperfectly substitutable and include IT-complementary skills such as programming skills and managerial ability (we denote the set of such skills by $\mathcal{C}$), as well as IT-neutral skills such as public speaking and

[5] There is a vast literature that examines the channels through which information technologies shape firm production decisions and organizational scope, worker tasks, and skill requirements. See for example, Bartel et al. (2007), Bresnahan et al. (2002), Autor et al. (2003), Caroli and vanReenen (2001), Acemoglu and Zilibotti (2001), Acemoglu and Restrepo (2018), Beaudry et al. (2010), Goos and Manning (2007), Goos et al. (2014), Akerman et al. (2015), Michaels et al. (2014), Autor et al. (2024).

event planning (we denote the set of such skills by $\mathcal{N}$).

We denote mismatch for skill i between a firm using technology of vintage τ and a worker of type $j \in (u,s)$ by $M_i^{\tau,j} \equiv F_i^{\tau,j} - W_i^{\tau,j}$ for $i \in$ (1,2,...N). Mismatch across all skills between a firm using technology of vintage τ and a worker of type $j \in (u,s)$ is denoted by Euclidean distance $M^{\tau,j} \equiv \|F^{\tau,j} - W^{\tau,j}\|_2 = \sqrt{\sum_{i=1}^{N}(F_i^{\tau,j} - W_i^{\tau,j})^2}$. We use these definitions to recast unskilled and skilled workers in the original model to be defined in terms of multidimensional skill mismatch. Specifically, we redefine a skilled worker as one who possesses skill attributes $W^{\tau,s}$ that equal the skill requirements $F^{\tau,s}$ of a firm using technology of vintage τ (i.e., $W_i^{\tau,s}$=$F_i^{\tau,s}$ for all individual skill dimensions i=1,2,...N). The skill mismatch $M^{\tau,s}$ of a firm-worker pairing of this type will thus be *zero*; vintage-specific human capital in the original model is thus represented by the specific skill weights desired by the firm, consistent with the notion of specific human capital defined by Lazear (2009). Correspondingly, we redefine an unskilled worker as one who possesses skill attributes $W^{\tau,u}$, where $W_i^{\tau,s} \neq F_i^{\tau,s}$ for at least one $i \in \{1,2,...N\}$; the skill mismatch $M^{\tau,u}$ of this firm-worker pairing will be *non-zero*.

Our second adaptation is to assume that technological change is skill-biased and that the demand for IT-complementary skills rises faster than the speed by which workers acquire these skills. This assumption is grounded in a large literature that examines changes in the wage structure through the lens of a race between technological change and education (Goldin and Katz, 2008). A number of papers find that changes in information and communication technologies during the recent decades have been associated with relative wage and employment changes that reflect rapid growth in the demand for IT-complementary skills that outpaces their supply. Formally, we assume that new technology vintages have greater IT-complementary skill requirements than older technologies: $\frac{\partial F_i^{\tau,u}}{\partial \tau} < 0$ for IT-complementary skills $i \in \mathcal{C}$. Because skill weights must add up to 1, this assumption implies that $\frac{\partial F_i^{\tau,u}}{\partial \tau} \geq 0$ for IT-neutral skills $i \in \mathcal{N}$. We model the relatively slower speed of workers' IT-complementary skill acquisition by assuming $|\frac{\partial W_i^{\tau,u}}{\partial \tau}| < |\frac{\partial F_i^{\tau,u}}{\partial \tau}|$ for IT-complementary skills $i \in \mathcal{C}$.

Our third adaptation is to use our definition of skill mismatch to microfound the production complementarity between unskilled and skilled workers that is otherwise assumed in Chari and Hopenhayn (1991). Specifically, we decompose the original assumption of $\frac{\partial f_U}{\partial S_t} > 0$ into the product of two components: $(\frac{\partial M^{\tau,u}}{\partial S_t}) \times (\frac{\partial f_U}{\partial M^{\tau,u}}) > 0$. For the first term, we assume that as the firm employs more skilled workers s_τ, the firm faces lower adjustment costs when assigning tasks and facilitating knowledge transfer to unskilled workers so that the skill requirements facing unskilled workers $F^{\tau,u}$ align more closely with their skill endowments $W^{\tau,u}$, thereby reducing skill mismatch $M^{\tau,u}$. Formally, this means $\frac{\partial M^{\tau,u}}{\partial S_t} < 0$. In the Appendix C, we provide a microfoundation for this assumption by setting up the firm's problem of designing skill requirements to minimize skill mismatch subject to adjustment costs that are increasingly alleviated as the firm employs more skilled workers. To provide a concrete example: as unskilled workers gain experience and become skilled, they develop managerial expertise that lowers the need for incoming workers to possess managerial capabilities.

For the second term, we assume that the marginal productivity of an unskilled worker $f_U(U_t, S_t)$ is decreasing and concave in the degree of skill mismatch $M^{\tau,u}$ between the worker and the firm. Formally, this implies $\frac{\partial f_U}{\partial M^{\tau,u}} < 0$ and $\frac{\partial^2 f_U}{\partial (M^{\tau,u})^2} < 0$. This assumption is supported by prior work that finds a negative association between mismatch and worker productivity (Coraggio et al., 2025; Guvenen et al., 2020; Lise and Postel-Vinay, 2020). The decomposition that we introduce to the model implies that an increase in the number of skilled employees allows the firm to raise the marginal productivity of its unskilled workers by altering their skill requirements and reducing the mismatch they exhibit with the firm.

Taken together, the adaptations that we introduce microfound key primitives of Chari and Hopenhayn (1991) without altering its mechanics. As a result, we are able to generate the same equilibrium conditions and comparative statics as the original model. For example, the distribution of skilled workers s_τ remains single-peaked across technology vintages. The enriched foundations that we provide, however, yield novel testable implications that enable us to address issues that are beyond the scope of the original framework: namely, the relationship between multidimensional skill mismatch and information technology diffusion. We derive these implications below, and provide a graphic illustration of our model in Figure 1.

Prediction 1: U-shaped skill mismatch over the technology life cycle.

In equilibrium, the model predicts that the mismatch between firms and unskilled workers, $M^{\tau,u}$, follows a U-shaped pattern with respect to technology vintage τ. As discussed earlier, we postulate that skill mismatch for unskilled workers decreases as the firm employs more skilled workers; i.e. $\frac{\partial M^{\tau,u}}{\partial S_t} < 0$. Because the equilibrium distribution of skilled labor, s_τ, is single-peaked across technology vintages—a key result shown by Chari and Hopenhayn (1991)—it follows that $M^{\tau,u}$ declines as s_τ rises, reaches its minimum when s_τ attains its peak, and increases thereafter as s_τ falls.

Intuitively, skill mismatch is high for newly introduced technologies because few workers initially possess the capabilities required to utilize them effectively. As technologies mature, more unskilled workers gain experience working with them and develop relevant skills. As firms employ a larger pool of skilled workers, the adjustment costs of assigning tasks more efficiently to unskilled workers decline, reducing mismatch for new employees. However, once technologies become sufficiently old and less productive, workers have dampened incentives to join these firms and become skilled in these technologies. Firms then face higher adjustment costs when reorganizing tasks and skill requirements for new hires, causing skill mismatch to rise once again.

Prediction 2: IT-complementary skill requirements exceed worker capabilities

The model predicts that firms' demand for IT-complementary skills exceeds unskilled workers' supply of these capabilities. That is, the model predicts positive mismatch, $M_c^{\tau,u} > 0$, for unskilled workers across IT-complementary skills $i \in \mathcal{C}$. This prediction stems from the assumption that technological progress is skill-biased towards IT-complementary skills and that firms adopting newer technologies will adjust their skill requirements for IT-complementary skills faster than unskilled workers are able to acquire these skills on their own. In equilibrium, the steady-state distribution of technology adoption will thus be characterized by firm requirements exceeding unskilled worker capabilities for IT-complementary skills. A corollary of this prediction is that there will be an excess supply of IT-neutral skills because skill weights sum to one. The model therefore predicts that $M_c^{\tau,u} \leq 0$ for IT-neutral skills $i \in \mathcal{N}$. These predictions impose intuitive discipline on the model mechanism: if skill mismatch indeed constrains IT adoption, the constraint should arise along the skill dimensions most complementary to IT. If, on the other hand, workers' endowments meet or exceed firms' requirements for IT-complementary skills, then the evidence would diminish the relevance of the model as an explanation for IT diffusion patterns.

Prediction 3: U-shaped skill mismatch for IT-complementary skills

The model predicts that skill mismatch along IT-complementary skill dimensions should be U-shaped across technology vintages. The logic for this prediction stems from the arguments used for predictions 1 and 2. As s_τ increases, firms will prioritize adjustments in IT-complementary skill requirements, as these skill requirements exhibit the largest gaps between firm demands and worker capabilities. As technology vintage increases and s_τ rises, the magnitude of these adjustments declines until s_τ peaks. Thereafter, as s_τ falls, firms find it increasingly difficult to reallocate tasks and reduce mismatch, leading to a rise in misalignment. The model thus predicts a U-shaped relationship between IT-complementary skill mismatch and technology vintage that drives the patterns we observe for the full array of skills in the sample.

Prediction 4: Firm growth opportunities decrease with technology vintage

In the model, newer technologies are inherently more productive than older vintages, implying that firms operating legacy technologies face weaker growth opportunities due to their lower productive capacity. Nevertheless, skill mismatch allows these firms to continue attracting workers and remain active, as workers lack the capabilities required to use newer, more productive technologies. Thus, rather than observing only firms that operate with high growth opportunities, the model predicts that we should observe a wide distribution of firms exhibiting a negative relationship between technology vintage and firm market-to-book ratios, a standard empirical proxy for growth opportunities. Finding such a pattern would be consistent with the model's implication that skill mismatch sustains the persistence of legacy firms despite their limited growth prospects. In contrast, if the data indicate that investment opportunities do not decline with technology vintage, then the evidence would suggest that the model has limited explanatory power for observed patterns of technology diffusion.

4 Empirical Approach

4.1 Challenges

Empirically characterizing skill mismatch and its relationship with firm technologies is challenging. Ideally, we would like to be able to identify the specific skills that different firms seek and the specific skills that workers possess. We could then compute the distance between the skill attributes of firms and matched workers to estimate skill mismatch, and examine how skill mismatch varies over the technology life cycle.

There are numerous challenges with implementing this procedure in practice. First, there are limited data of sufficient granularity that adequately describe the skill attributes of firms and workers. Prior studies of skill mismatch either rely on survey data of managers and workers, or administrative datasets that contain information on worker demographics, education, and experience. These data lack the detail required to meaningfully characterize skill mismatch across firms that utilize different information technologies. Firms relying on different IT systems, for example, will often require broadly similar levels of education, yet very different levels of expertise with specific technologies.

Second, it is conceptually challenging to identify the specific skills that firms truly demand, as these skills are fundamentally unobservable and difficult to discern from observable data on firm skill requirements. For example, a job posting that asks for "Java programming" skills may signal demand for a specific coding task, general problem-solving ability, and/or social skills that facilitate teamwork with software engineers. Additionally, firms will have implicit demand for many skills that are not explicitly specified in job postings, such as basic language and communication skills. The same points apply when attempting to identify workers' true skill portfolios from observed skills listed on workers' resumes.

Prior studies that analyze skills data in job postings frequently devise keyword tables to classify skill requirements. For example, one might look for keywords in job postings and then use those keywords to label skills as "cognitive" or "routine". This bag-of-words approach is well-suited for many applications, however, this approach is inappropriate for our setting, because any mapping of observable skills-to-unobservable skills defined by the econometrician may fail to resemble the skill mapping used by market participants. This disparity can lead to severe measurement error that distorts estimates of skill mismatch.

Third, it is computationally challenging to process skills data from job postings and worker resumes. There are typically thousands of unique and inconsistently defined skill requirements that vary significantly across firms. Deming and Kahn (2018) study U.S. job postings and find that skill requirements vary greatly across narrowly defined occupations; they estimate that firm-specific factors account for nearly 30% of the variation in posted skill requirements. Similar problems apply to the lexecon of skills that workers use to describe their human capital on resumes. These aspects of skills data make it challenging to produce estimates of skill mismatch that are computationally feasible and comparable across firms.

4.2 Empirical Methodology

We devise an empirical methodology that helps us overcome the aforementioned challenges and empirically characterize skill mismatch over the technology life cycle. First, we construct a novel matched employer-employee panel dataset that combines information from U.S. firm job postings (provided by Lightcast, formerly known as Burning Glass) with information from workers' online professional profiles (provided by Revelio). The job postings data contain detailed information on the skill requirements of firms for various positions. Job postings also list the specific technologies that companies want workers to be familiar with. The online professional profile data provide detailed information on the skills that workers possess, along with the start and end dates at various companies where workers have previously been employed. These raw data represent the most detailed information currently available for describing the skill attributes of firms and workers across the U.S.

From these data, we construct new firm-worker matches each year by identifying workers who join a firm within one year of the posting of a job ad and aggregating all such workers and associated job ads for each firm-year. This method provides a first approximation to the actual firm-worker matches realized each year; there are many workers who do not have online professional profiles and there are many vacancies that are not advertised online. We evaluate these sampling issues and discuss the robustness of our findings to alternative sampling criteria and different matching specifications in Section 5. We merge our data with firm-level information from Compustat and CRSP to characterize firms' annual balance sheet characteristics, investment and operating activities, and market valuations. Estimates of TFP are provided by Imrohoroglu and Tuzel (2014) and estimates of investment in intangible capital are provided by Peters and Taylor (2017). We describe the construction of the dataset and relevant sources in more detail in Appendix A.

We then employ the zero-shot classification algorithm developed by Yin et al. (2019) and the BERT large language model (Devlin et al., 2019) to map the raw, free-text skill descriptions reported by firms and workers into a common set of 25 predefined skill groups. This classification step addresses the conceptual challenges noted above by producing comparable measures of skill attributes that can be used to estimate mismatch. The key advantage of using the BERT large language model is that it leverages contextual information in text to interpret what firms and workers mean when they report skills in unstructured form. In effect, we rely on the model's semantic representation of skill descriptions to approximate how market participants themselves interpret and categorize observable skills. We believe that this procedure provides a more accurate and relevant mapping for our context than could be achieved by an econometrician manually coding skill attributes using a bag-of-words approach.

To describe this classification step more precisely, we employ a predefined set of 25 skill categories developed by Revelio Labs. These categories are constructed by Revelio using clustering algorithms applied to the large corpus of skills that they extract from online professional profiles. Distinct skill expressions are first grouped into unique skills, which are then iteratively clustered into broader categories. Each cluster is assigned a representative label chosen to be both frequent within the cluster and central to its semantic meaning.[6] We adopt Revelio's skill taxonomy rather than define our own skill taxonomy, under the assumption that their large-scale clustering procedure spans the raw skill space more efficiently and systematically than any ad hoc grouping. We believe that this approach minimizes measurement error and reduces "researcher bias" that could arise when classifying raw skills into comparable skill groups used to generate estimates of mismatch.

The zero-shot classification algorithm maps the raw skill data for each firm or worker into a 25-dimensional vector, where each coordinate j=1,2,...,25 represents the model's estimated confidence that the raw skill profile entails Revelio skill category j. Formally, for each input skill profile k and predefined skill label j, the language model computes a logit score for the hypothesis "label j is true given input k." This entailment score, or degree of entailment, is defined relative to the alternative hypotheses that skill label j is contradicted or neutral given input k. Entailment logits are computed for each of the 25 skill categories and then passed through a softmax transformation, producing a normalized distribution that sums to one. Intuitively, a higher score for skill category j indicates greater model confidence that raw skill input k is consistent with category j relative to the other predefined skill groups. The resulting vector of scores captures the degree of alignment between the reported raw skills and each standardized category, providing a harmonized representation that can be directly compared across firms and workers.

We use these skill vector outputs from zero-shot classification to estimate skill mismatch.

[6] See Revelio's full methodology description at: https://www.data-dictionary.reveliolabs.com/methodology.html#clustering-skills.

Specifically, for each firm-year observation, we define skill mismatch between the firm and its workers as the Euclidean distance between firm skill vector $\mathbf{f} = (f_1, f_2, \ldots, f_{25})$ and worker skill vector $\mathbf{w} = (w_1, w_2, \ldots, w_{25})$:

$$m(\mathbf{f}, \mathbf{w}) = \sqrt{\sum_{j=1}^{25} (f_j - w_j)^2}.$$

for skill attributes j=1, 2,...25. Intuitively, our skill mismatch measure captures how far apart a firm's skill requirements and a worker's skill portfolio are in the 25-dimensional "skill space". If a firm and a worker place similar weight on the same skill categories, their vectors will be close together, resulting in a low level of skill mismatch. By contrast, if the firm emphasizes skills the worker does not possess—or the worker has skills that are not demanded by the firm—their vectors will differ more sharply, leading to a higher level of mismatch.

Finally, we follow the prior literature (eg. Tambe (2014)) to characterize firms' information technologies by identifying the vintage years of each of the technologies that are listed in firms' job postings, where the vintage year is the year in which a technology is invented. We compute the average (and median) vintage year of all the technologies that appear in the firms' job postings. This measure captures the stage of the technology cycle that corresponds to the information technologies that are being adopted by firms and listed in job postings to convey worker skill requirements.

Armed with measures of skill mismatch and firm IT capital vintages, we perform a number of empirical analyses to present new facts that characterize skill mismatch in the context of firms' technology decisions. We examine how skill mismatch varies across firms and within-firms, we examine how mismatch varies across specific skill categories, and we study how skill mismatch is related to measures of firm performance such as operating profitability and total factor productivity. We also perform a number of analyses to establish the robustness of our conclusions to alternative variable definitions and empirical specifications.

5 Data

We construct a matched employer-employee dataset that combines information from three separate databases. First, we obtain job postings for U.S. firms from Lightcast for the years 2010 to 2024. Second, we merge these data with firm-level information from Compustat. Third, we merge our firm-year data with Revelio job profile data for workers. Most of the job profile data in Revelio comes from Linkedin. The data thus allow us to classify the starting years for employees of firms in the sample. The Revelio and Lightcast data contain information on the multidimensional skill attributes of workers and firms, respectively, and are the key data that we use to constructing our measures of skill mismatch. The specific skill categories that Revelio uses to classify worker skills is presented in Table A2. We break these skills up into three groups, based on the literature on IT and productivity: IT-complementary technical skills, IT-complementary nontechnical skills, and Peripheral skills that are unrelated to information technology deployment.

Table 1 describes our sample of firms (Panel A); we present data from Compustat to illustrate the population of firms and describe how our sample compares with the population. Our sample firms are generally larger and more productive than the average Compustat firm, reflecting sample selection toward technology-intensive companies that are more likely to post jobs online and have employees with LinkedIn profiles. Table 1 also describes firm characteristics such as age, capital intensity, etc.

[Insert Table 1]

Figure 2 presents a histogram that depicts the ages of information technologies utilized by firms in our sample. The figure depicts the distribution of firm-year observations that correspond to different age buckets. The most commonly observed technology vintage in our sample of job postings reflects information technologies that have an age of approximately 28 to 31 years. A substantial portion of our firms utilize older technologies that have been in place for greater than 36 years.

To get a sense of the types of technologies that appear in job postings, Table A1 presents the top ten most frequently cited technologies that are listed in each technology bin in our sample. As the table indicates, we see that frontier technologies in our sample include cloud computing and modern programming languages, while technologies listed in the latest bins correspond to lower level programming languages and applications that are associated with mainframe systems and older IT systems.

[Insert Figure 2]

Table 2 presents summary statistics that describe our central measures of multidimensional skills, across each of the 25 skill groups that we use in our main analysis. The statistics for demand corresponds to the entailment scores for a given skill group across all job postings in the sample, while the mean statistic for supply corresponds to the entailment scores for workers in our sample. The "gap" measure corresponds to the absolute value of the difference between the demand and supply measures for every firm-year observation in the sample.

[Insert Table 2]

6 Empirical Findings

6.1 Skill mismatch over the technology life cycle

The model's first empirical prediction is that multidimensional skill mismatch follows a U-shaped pattern over the technology life cycle. We test this prediction using both graphical evidence and regression analysis. Figure 3 plots average skill mismatch against technology vintage for all firm–year observations in the sample. Technology vintage is grouped into ten equal-weighted bins spanning more than 40 years. The figure reveals a pronounced U-shaped relationship: skill mismatch is relatively high for firms implementing new technologies, declines steadily as technologies mature, and rises again for older vintages. The standard-error bars indicate that differences in mismatch across vintages are statistically significant for a range of bin comparisons. Skill mismatch is highest among firms using the newest technologies (2.1), declines by approximately 40 percent to its minimum after 33.8 years, and subsequently increases (to 1.2) for the oldest technologies in the sample—overlapping the levels observed at earlier stages of the life cycle.

Several additional features of Figure 3 shed light on the relationship between skill mismatch and the technology diffusion process. The long time-series of observed technology vintages in the sample reflects the slow pace of adoption in the data, consistent with studies of gradual diffusion relative to the rate of technological progress (Kalyani et al., 2025). The declining slope of mismatch among younger technologies aligns with the model's diminishing-adjustment mechanism: as the stock of skilled workers expands, firms first close their largest skill gaps, generating progressively smaller reductions in mismatch over time. Finally, mismatch reaches its minimum roughly 35 years after technology inception, suggesting that it can take several decades for information technologies to fully diffuse and for skill mismatch to be minimized.

Table 3 formalizes these patterns using regression estimates of skill mismatch as a function of technology vintage. Across specifications, the coefficient on the linear term in vintage is negative (between -2.2 and -2.3), while the quadratic term is positive (between 3.3 and 3.5), confirming the U-shaped relationship observed in the figure. These results are robust to the inclusion of year and industry fixed effects, indicating that the predicted relationship holds both across and within industries. We also observe statistically significant results consistent with the model using specifications with year and firm fixed effects, with the caveat that these results rely on significantly less variation than other specifications, as technology vintage varies little within most firms over time. The robustness of the results across specifications shows that the U-shaped pattern observed in Figure 3 is not an artifact of particular industry characteristics or firm attributes, but instead reflects a systematic relationship between skill mismatch and technology vintage that is consistent with the adjustment dynamics emphasized in the model.

Taken together, Figure 3 and Table 3 support the view that skill mismatch evolves endogenously over the technology life cycle. Mismatch declines as technologies mature and the supply of skilled workers expands, consistent with the model's prediction that firms initially close the largest skill gaps as complementary skills accumulate. Once technologies become sufficiently old, however, mismatch rises again as worker incentives to join firms using less productive technologies weaken.

These findings rule out several plausible alternative priors. First, skill mismatch is not confined to firms adopting new technologies: it remains an empirically relevant constraint for firms operating legacy systems. This pattern is consistent with practitioner surveys and business-press accounts documenting hiring difficulties among firms reliant on older technologies. Second, mismatch is not a persistent, technology-invariant feature of labor markets. Instead, the U-shaped pattern indicates that systematic forces govern how skill mismatch increases and decreases over the technology life cycle.

An important implication of the findings is that similar observed levels of skill mismatch can reflect fundamentally different economic forces. We observe comparable mismatch among firms deploying frontier technologies and among firms using legacy technologies. According to the model, mismatch in the former case reflects a scarcity of workers who possess the skills required to operate new technologies, whereas mismatch in the latter case partly reflects weak worker incentives to join firms that rely on technologies that are relatively less productive. Distinguishing between these mechanisms is essential for interpreting observed skill shortages and for evaluating policies aimed at workforce development and technology adoption.

6.2 Characterizing Skill Bottlenecks to Technology Diffusion

The model also makes predictions about the key skill attributes that serve as bottlenecks to information technology diffusion. Prediction 2 states that firm requirements should exceed worker capabilities for IT-complementary skills. These skills may include technical skills such as programming and troubleshooting, as well as non-technical skills such as management and organizational capabilities Brynjolfsson and Hitt (2000). To test this prediction, we examine skill mismatch across each of the individual skill categories in our data. We graphically depict average skill vector weights for firm requirements and worker attributes in Figure 4. We also aggregate the skills in our sample into three groups following prior work on IT and productivity (Bartel et al., 2007): IT-complementary technical skills, IT-complementary nontechnical skills, and IT-neutral skills, and present these aggregated skills in Figure 6.

Consistent with this model, Figures 4 and 6 show that mismatch for IT-complementary skills—and especially for IT-complementary nontechnical skills—is substantially larger than mismatch for IT-neutral skills across all technology vintages. These results indicate that firms' skill requirements exceed workers' capabilities precisely in the dimensions most relevant for IT adoption. It is also notable that the largest bottlenecks to diffusion arise not only in technical competencies but also in nontechnical, IT-complementary skills that the literature identifies as critical for organizational and process changes accompanying new technologies (Bartel et al., 2007; Bresnahan et al., 2002).

For example, the most pronounced mismatches emerge for nontechnical capabilities such as managerial ability and related organizational skills, as well as for technical capabilities tied to IT infrastructure—technology support, networking, and manufacturing-related software. These attributes are directly complementary to IT-intensive production systems. By contrast, attributes that are in relatively low demand and exhibit little sensitivity to technology vintage—such as legal writing, biotech, and risk management—display minimal mismatch and, in several cases, excess supply. Taken together, these findings demonstrate that overall mismatch does not simply reflect a shortfall of "STEM skills" alone, but rather a targeted scarcity in high-weight, IT-complementary attributes, including managerial and finance-related capabilities. This pattern aligns with the model's weighting mechanism: gaps in attributes receiving the highest weights become the most binding for aligning worker skills with firm requirements.

To test the model's third main prediction, we next examine Euclidean distance measures of mismatch in two additional ways. First, we examine skill mismatch separately for each of the three skill groups in Figure 4. Figure 7 shows that the U-shaped pattern in overall mismatch is driven primarily by IT-complementary nontechnical skills. This evidence supports the hypothesis that firms reduce mismatch by prioritizing adjustments in the attributes that are farthest from workers' capabilities—which, in our setting, are precisely IT-complementary skills. We reinforce this conclusions by presenting regression analysis in Tables 5, 6, and 7. The coefficients for technology vintage, estimated while controlling for year and industry fixed effects, echo the results in our main specification for IT-complementary skills and support the model prediction.

In a second analysis, we examine skill mismatch and technology vintage separately for each 2-digit occupational code in our sample. We estimate our main regression specification of skill mismatch on technology vintage (using both a quadratic and a piecewise linear specification) for each occupational code, and present the results in Table 4. The findings indicate that we observe significant skill mismatch that decreases over time for occupations such as management, computer and math intensive jobs, engineers, sales, and production worker roles (the individual skill vector weights for these workers and their employers are depicted in Figure 5). In contrast, Table 4 shows that occupations that are typically less related to information technologies do not show meaningful variation in mismatch across technology vintages. The results support the view that skill mismatch for occupations that are central to IT-enabled changes within the firm are critical for explaining IT-diffusion rates and exhibit patterns that are consistent with the model mechanism.

6.3 Multidimensional skill mismatch and the returns to investment across technology vintages

The model also predicts that skill mismatch can help explain the persistence of legacy systems with relatively poor growth opportunities. Firms operating older technologies remain viable because they attract workers whose skill portfolios are poorly aligned with frontier systems.

Skill mismatch at the technological frontier therefore allows firms deploying legacy systems to continue operating despite inferior performance.

To test this prediction, we estimate sample firms' market-to-book ratios—a commonly used proxy for firm growth prospects—and study how these ratios vary with technology vintage. We graph the relationship between market-to-book ratios and technology vintage in Figure 8. Consistent with the model mechanism, we document that firms operating legacy systems exhibit lower market-to-book ratios than firms that adopt new technologies. Both types of firms are nevertheless prevalent, and continue to operate and hire workers despite the skill mismatch that they both face in the labor market. The graphical results are further reinforced by the regression analysis presented in Table 8. Across specifications that control for year and industry fixed effects, the estimated coefficient on technology vintage is negative and statistically significant, ranging from -1.2 to -2.8.

These findings are consistent with the model's prediction that technologies with lower marginal returns to investment persist because of labor market mismatch rather than because they remain equally attractive investment opportunities. These results reject several alternative priors. First, they rule out the possibility that, in our sample, older technologies generate marginal returns to investment comparable to those of newer technologies. Second, they are inconsistent with explanations in which the coexistence of old and new technologies reflect persistent industry-level differences unrelated to technology vintage. Finally, the results suggest that the persistence of legacy technologies is not solely driven by technological or adjustment frictions on the capital side; instead, the findings point to skill mismatch and worker sorting as central forces shaping firms' investment opportunities across the technology life cycle. These results complement existing explanations for the persistence of legacy technologies, which otherwise emphasize adjustment costs, switching frictions, or technological lock-in. Our analysis shows that multidimensional skill mismatch provides a distinct economic mechanism that can explain the coexistence of frontier technologies and legacy systems despite vast differences in future growth prospects.

7 Robustness Analysis

Our empirical approach to measuring skill mismatch and characterizing its relationship with technology vintage takes advantage of recent computational advances in large language models. Our approach has many strengths, but it also has several potential limitations. In this section, we discuss various concerns about our empirical methodology and we present additional analyses to assess the robustness of our findings.

7.1 Choice of language model

In our analysis, we rely on the BERT large language model. BERT has become the workhorse model for text classification because it is bidirectional (relative to uni-directional large models that only read text in one direction), and therefore better suited for capturing contextual information than many other language models. This makes BERT well-suited for our skill classification task, because it can capture the meaning of explicit raw skill data and it can also capture the skills that are implied but not explicitly mentioned in observable data.

There have been recent advances to BERT, such as ModernBERT, that are reportedly even better at capturing contextual information than BERT. The problem with these more advanced models is that they are computationally intensive; running ModernBert on our sample requires several weeks of run-time. Neverthless, to explore the robustness of our results to the choice of large language model, we repeat our analysis for firm-worker matches in the year 2020 using ModernBert. We find results that are even stronger than what we find using Bert. For example, as shown in Table 9, column 4, the U-shape of skill mismatch over technology vintage is even more pronounced than what we estimate using BERT. Thus our findings do not appear to be driven solely by the particular choice of large language model that we use for our analysis.

[Insert Table 9]

7.2 Technology vintage measure

In our main specifications, we use the average vintage for all information technologies listed in a firm's job postings as an estimate of the IT capital vintage of the firm. The average IT vintage that we estimate might be skewed by software that is particularly old and so commonplace that it may not be particularly relevant for characterizing changes in production systems that are empirically most relevant in the context of skill mismatch. For example, around 25% of the information technologies listed in job postings belong to Microsoft Office suite.

We address this concern in two ways. First, we show that using alternative measures for IT capital vintage, such as median vintage year for IT capital, does not significantly alter our results. Second, we re-run our analysis, excluding Microsoft office software from our measures of IT capital vintage, and we find that our results remain the same. Our findings thus show that the patterns of skill mismatch that we present are not significantly related to the particular measures of IT capital vintage that we employ.

7.3 Unobserved skill proficiency

Our analysis does not account for the relative level of proficiency associated with skills in our sample. For example, when two workers list "Python programming" on their professional profiles, we do not distinguish between the two workers based on their relative levels of proficiency with Python. Similarly, if two firms list the same skill requirements in their job postings, we do not distinguish between minimum levels of proficiency in the skills required by the two firms. In reality, there is likely to be significant heterogeneity in unobserved ability that further adds to the degree of skill mismatch that we measure in our sample.

One concern for our analysis is that firm-worker matches that exhibit large mismatch using our measures may actually be close skill matches once unobservable proficiency is taken into account. This concern is important, but it is unlikely to significantly alter our main conclusions. First, our analysis implicitly assumes that firms and workers have a common (minimum) standard of proficiency in a particular skill when it is explicitly listed online. Thus, our measure of skill mismatch is still informative about relative differences in skill attributes between firms and workers, taking into account skill proficiency in the decision to list a particular skill online or not. Second, we note that our measure of skill mismatch is associated with lower firm productivity and lower wages, which suggests that our measures of skill mismatch are unlikely to be systematically biased by unobserved levels of skill proficiency.

7.4 Misidentifying firm-worker matches

We show that our results are robust to using alternative time windows for matches firms with workers. For example, we show that our results are similar when we match workers who join the firms within the same year of the firm's job postings. We also find that the results are the same when we identify workers who join firms two years after jobs are posted online.

7.5 Alternative skill labels

We show that our results are also robust to alternative skill labels, such as the categories of skills specified in Deming and Kahn (2018). In the literature, it is common for skills associated with technology to be classified as cognitive, routine, etc. When we use these skill classifications in place of the skill classifications defined by Revelio, we observe similar results for all of our tests. Thus our findings are not driven by specific skill labels per se; the results reflect skill mismatch that is consistent with the predictions of our model.

8 Conclusion

This paper presents evidence of multidimensional skill mismatch as a dynamic factor that shapes the diffusion of information technologies. We document that skill mismatch is U-shaped in technology vintage and is especially pronounced for IT-complementary skills—both technical and nontechnical. We also show that the patterns of capital investment across vintages are consistent with the view that skill mismatch shapes firms' investment in both intangible and tangible capital. In particular, our results help rationalize J-curve dynamics in which early investment in intangible assets is high, measured TFP appears low at the onset of new technology adoption, and productivity gains materialize only gradually as skills adjust. Together, these findings support models of technological change in which delays in skill adjustment play a key role in generating firm-level heterogeneity in productivity.

Our analysis highlights several promising directions for future research. On the theoretical side, understanding the incentives of workers to acquire skills that complement new technologies—and the incentives of firms to adopt new technologies given workers' human-capital investment behavior—remains an open and important question. More broadly, identifying which specific skills are most critical for facilitating technology diffusion, and how these skills can be acquired or developed most efficiently, is central to understanding the productivity impact of ongoing technological change. These issues have particular urgency today, as policymakers and business leaders increasingly point to skill gaps as a first-order barrier to innovation and growth. On the empirical side, our findings underscore the potential of large language models and other modern computational tools for measuring and analyzing the evolving nature of multidimensional skill demand and supply in the labor market. Such tools offer a powerful way to study the micro-foundations of technological diffusion and may help open new avenues for understanding how skills, technology, and productivity co-evolve.

References

Acemoglu, D. and Autor, D. (2011). *Skills, Tasks and Technologies: Implications for Employment and Earnings*, volume 4, Part B.

Acemoglu, D. and Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542.

Acemoglu, D. and Zilibotti, F. (2001). Productivity differences. *Quarterly Journal of Economics*, 116(2):563–606.

Adão, R., Beraja, M., and Pandalai-Nayar, N. (2024). Fast and slow technological transitions. *Journal of Political Economy Macroeconomics*, 2(2):183–227.

Akerman, A., Gaarder, I., and Mogstad, M. (2015). The skill complementarity of broadband internet. *Quarterly Journal of Economics*, 130(4):1781–1824.

Autor, D., Chin, C., Salomons, A., and Seegmiller, B. (2024). New frontiers: The origins and content of new work, 1940–2018. *Quarterly Journal of Economics*, 139(3):1399–1465.

Autor, D. H., Levy, F., and Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, 118(4):1279–1333.

Baley, I., Figueiredo, A., and Ulbricht, R. (2022). Mismatch cycles. *Journal of Political Economy*, 130(11):2943–2984.

Bartel, A., Ichniowski, C., and Shaw, K. (2007). How does information technology affect productivity? plant-level comparisons of product innovation, process improvement, and worker skills. *Quarterly Journal of Economics*, 122(4):1721–1758.

Beaudry, P., Doms, M., and Lewis, E. (2010). Should the personal computer be considered a technological revolution? evidence from u.s. metropolitan areas. *Journal of Political Economy*, 118(5):988–1036.

Bresnahan, T. F., Brynjolfsson, E., and Hitt, L. M. (2002). Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. *Quarterly Journal of Economics*, 117(1):339–376.

Brunello, G. and Wruuck, P. (2021). Skill shortages and skill mismatch: A review of the literature. *Journal of Economic Surveys*, 35(4):1145–1167.

Brynjolfsson, E. and Hitt, L. M. (2000). Beyond computation: Information technology, organizational transformation and business performance. *Journal of Economic Perspectives*, 14(4):23–48.

Brynjolfsson, E., Hitt, L. M., and Yang, S. (2002). Intangible assets: Computers and organizational capital. *Brookings Papers on Economic Activity*, 2002(1):137–181.

Caroli, E. and vanReenen, J. (2001). Skill-biased organizational change? evidence from a panel of british and french establishments. *Quarterly Journal of Economics*, 116(4):1449–

1492.

Chari, V. V. and Hopenhayn, H. (1991). Vintage human capital, growth, and the diffusion of new technolog. *Journal of Political Economy*, 99(6):1142–1165.

Coraggio, L., Pagano, M., Scognamiglio, A., and Tåg, J. (2025). Jaq of all trades: Job mismatch, firm productivity and managerial quality. *Journal of Financial Economics*, 164:103992.

Dell, M. (2025). Deep learning for economists. *Journal of Economic Literature*, 63(1):5–58.

Deming, D. and Kahn, L. (2018). Skill requirements across firms and labor markets: Evidence from job postings. *Journal of Labor Economics*, 36(S1).

Deming, D. J. (2023). *Multidimensional human capital and the wage structure*, volume 7.

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). Bert: Pre-training of deep bidirectional transformers for language understanding. *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186.

Goldin, C. and Katz, L. F. (2008). *The Race Between Education and Technology.* Belknap Press for Harvard University Press.

Goos, M. and Manning, A. (2007). Lousy and lovely jobs: The rising polarization of work in britain. *Review of Economics and Statistics*, 89(1):118–133.

Goos, M., Manning, A., and Salomons, A. (2014). Explaining job polarization: Routine-biased technological change and offshoring. *American Economic Review*, 104(8):2509–2526.

Guvenen, F., Kuruscu, B., Tanaka, S., and Wiczer, D. (2020). Multidimensional skill mismatch. *American Economic Journal: Macroeconomics*, 12(1):210–244.

Imrohoroglu, A. and Tuzel, S. (2014). Firm-level productivity, risk, and return. *Management Science*, 60:2073–2090.

Kalyani, A., Bloom, N., Carvalho, M., Hassan, T., Lerner, J., and Tahoun, A. (2025). The diffusion of new technologies. *Quarterly Journal of Economics*, 140(2):1299–1365.

Lazear, E. P. (2009). Firm-specific human capital: A skill-weights approach. *Journal of Political Economy*, 117(5):914–940.

Lindenlaub, I. and Postel-Vinay, F. (2023). Multidimensional sortingunder random search. *Journal of Political Economy*, 131(12):3267–3540.

Lise, J. and Postel-Vinay, F. (2020). Multidimensional skills, sorting, and human capital accumulation. *American Economic Review*, 110(8):2328–2376.

ManpowerGroup (2024). 2024 global talent shortage: Infographic. https://go.manpowergroup.com/hubfs/Talent%20Shortage/Talent%20Shortage%202024/MPG_TS_2024_GLOBAL_Infographic.pdf. ManpowerGroup surveyed over 40,000 employers in 41 countries on difficulty finding needed talent, October 2023.

Michaels, G., Natraj, A., and Reenen, J. V. (2014). Has ICT poliazied skill demand? Evidence from elevent countries over twenty-five years. *Review of Economics and Statistics*, 96(1):60–77.

MIT Technology Review Insights (2023). New approaches to the tech talent shortage. Technical report, MIT Technology Review Insights. Executive briefing paper, sponsored by HCLTech.

Monaghan, B. D. and Bass, J. M. (2020). Redefining legacy: A technical debt perspective. In Morisio, M., Torchiano, M., and Jedlitschka, A., editors, *Product-Focused Software Process Improvement*, pages 254–269, Cham. Springer International Publishing.

Moscarini, G. (2001). Excess worker reallocation. *Review of Economic Studies*, 68(3):593–612.

OECD, editor (2024). *Understanding Skill Gaps in Firms: Results of the PIAAC Employer Module*. OECD Skills Studies. OECD Publishing, Paris. Accessed: March 16, 2026.

Peters, R. and Taylor, L. (2017). Intangible capital and the investment-q relation. *Journal of Financial Economics*, 123:251–272.

Tambe, P. (2014). Big data investment, skills, and firm value. *Management Science*, 60(6):1452–1469.

Yin, W., Hay, J., and Roth, D. (2019). Benchmarking zero-shot text classification: Datasets, evaluation and entailment approach. *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3914—-3923.

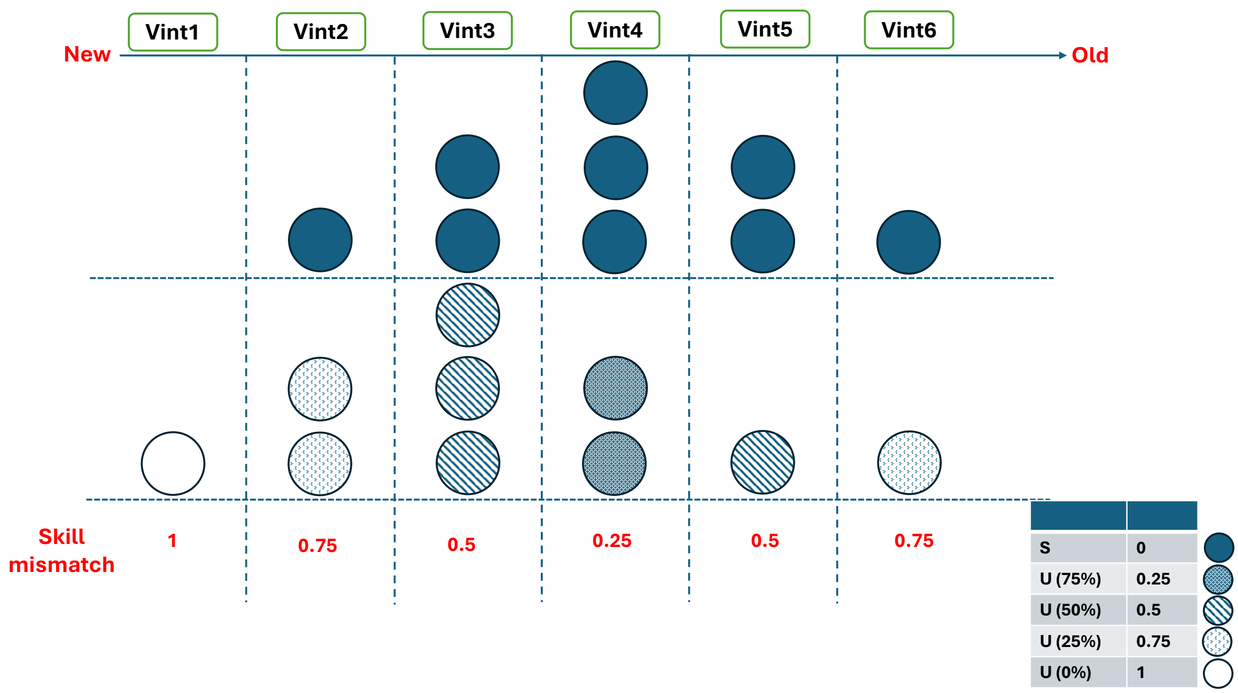

Figure 1: Model Overview

This figure provides a graphic illustration of our model. Each column corresponds to a technology vintage, increasing from the youngest to oldest vintage. In the first period, a single unskilled worker joins a firm employing the newest technology. Because the firm has no existing skilled employees, the level of mismatch between the new worker and the firm is normalized to 1. After one period of experience, this worker becomes skilled, and then two new unskilled workers join the firm. The skilled worker has 0 mismatch with the firm, while the two unskilled workers each reflect mismatch 0.25, since the skilled worker enables the firm to rearrange tasks and skill requirements that more closely align with unskilled workers' capabilities. The evolution of mismatch and employment growth of the firm continues each period until the technology reaches middle vintage 4. At this time (and moving forward), fewer unskilled workers will join the firm, since these workers face less incentive to work with sufficiently older technologies that are less productive than the newest technologies. The firm thus gets fewer skilled workers over time, thereby increasing the mismatch realized between the firm and its new employees This process continues until the technology becomes obsolete.

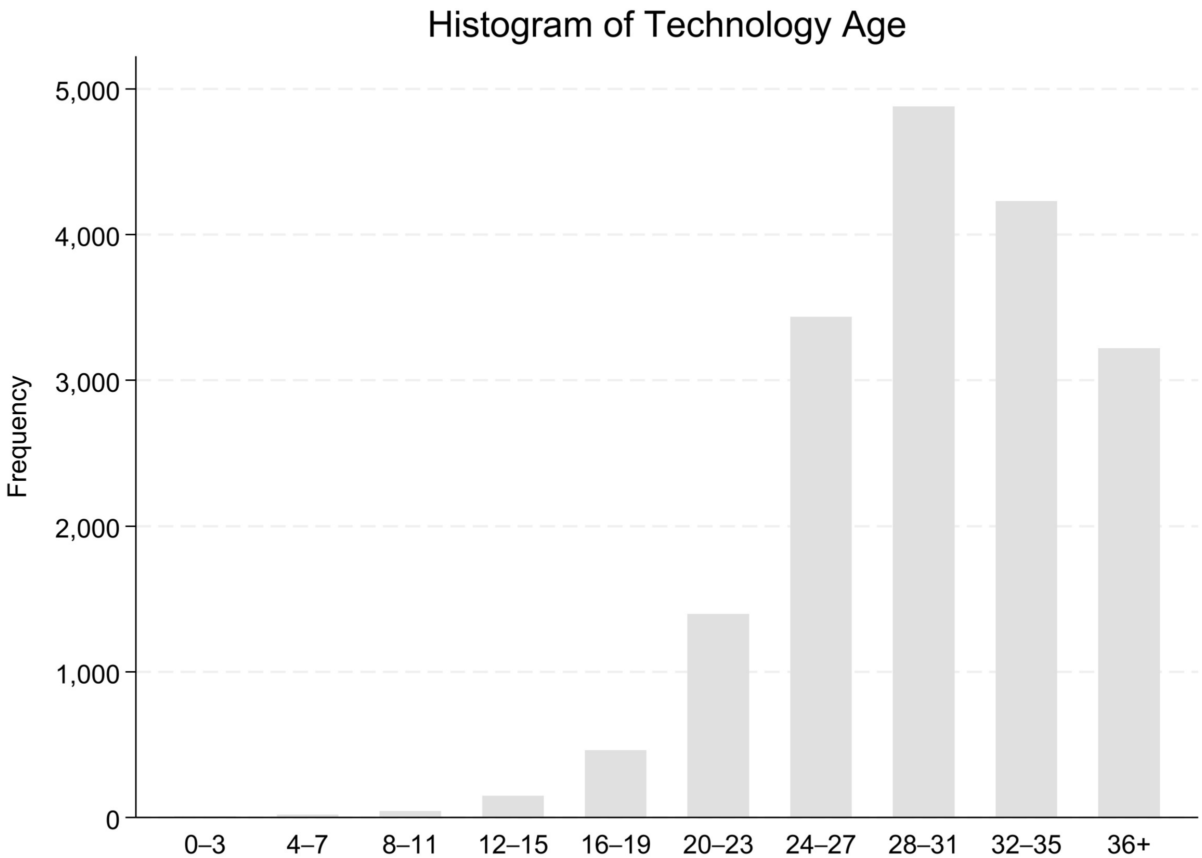

Figure 2: Technology Age Distribution

This figure presents a histogram of technology age observed for firms in our sample. Each data point corresponds to a firm-year observation. The x-axis depicts bins of technology age, computed as the average age of all information technologies listed in the firm's job postings in a given year. The y-axis depicts the frequency (count) of firm-year observations in the sample that belong to each technology age bin.

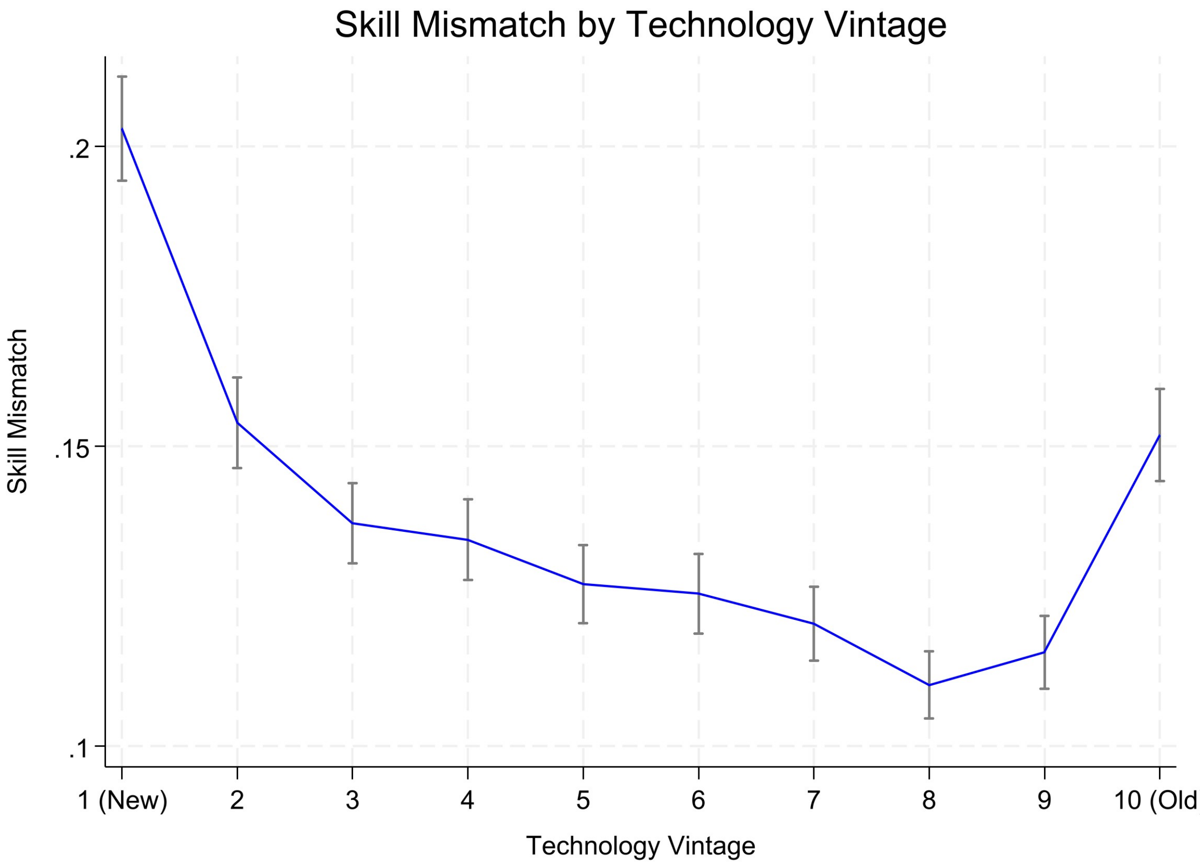

Figure 3: Skill Mismatch Across Technology Vintages
This figure presents the empirical relationship between firm technology vintage and average skill mismatch between the firm and its new employees. Skill mismatch is defined as the Euclidean distance between the firm's skill requirements and new employees' skill portfolios. Technology vintage is measured as the relative age of the technologies used by the firm; newer technologies corresponding to lower technology vintage values. Bin#1 corresponds to firm-year observations whose technology vintage between 0 and 23.4 years. Bin#2 corresponds to firm-year observations whose technology vintage between 23.4 and 26.1 years. Bin#3 corresponds to firm-year observations whose technology vintage between 26.1 and 27.8 years. Bin#4 corresponds to firm-year observations whose technology vintage between 27.8 and 29.3 years. Bin#5 corresponds to firm-year observations whose technology vintage between 29.3 and 30.8 years. Bin#6 corresponds to firm-year observations whose technology vintage between 30.8 and 32.2 years. Bin#7 corresponds to firm-year observations whose technology vintage between 32.2 and 33.8 years. Bin#8 corresponds to firm-year observations whose technology vintage between 33.8 and 35.6 years. Bin#9 corresponds to firm-year observations whose technology vintage between 35.6 and 38.2 years. Bin#10 corresponds to firm-year observations whose technology vintage between 38.2 and 91.5 years.

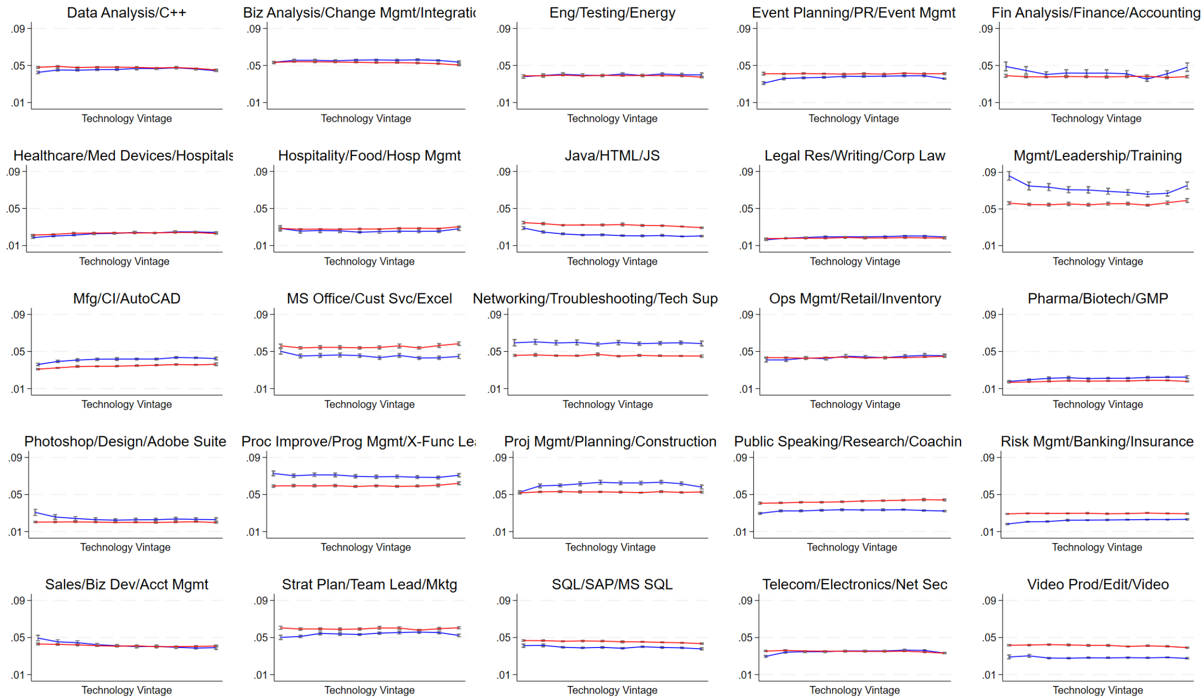

Figure 4: Skill Vector Weights by Skill Category

This figure presents the distribution of skill vector weights for job requirements (blue) and skill vector weights for newly hired employees (red) across each of the 25 skill categories used in our analysis (skill definitions provided in Appendix A2). The Y-axis depicts the average skill weight for each skill category, the X-axis depicts technology vintage bins as defined in Figure 3.

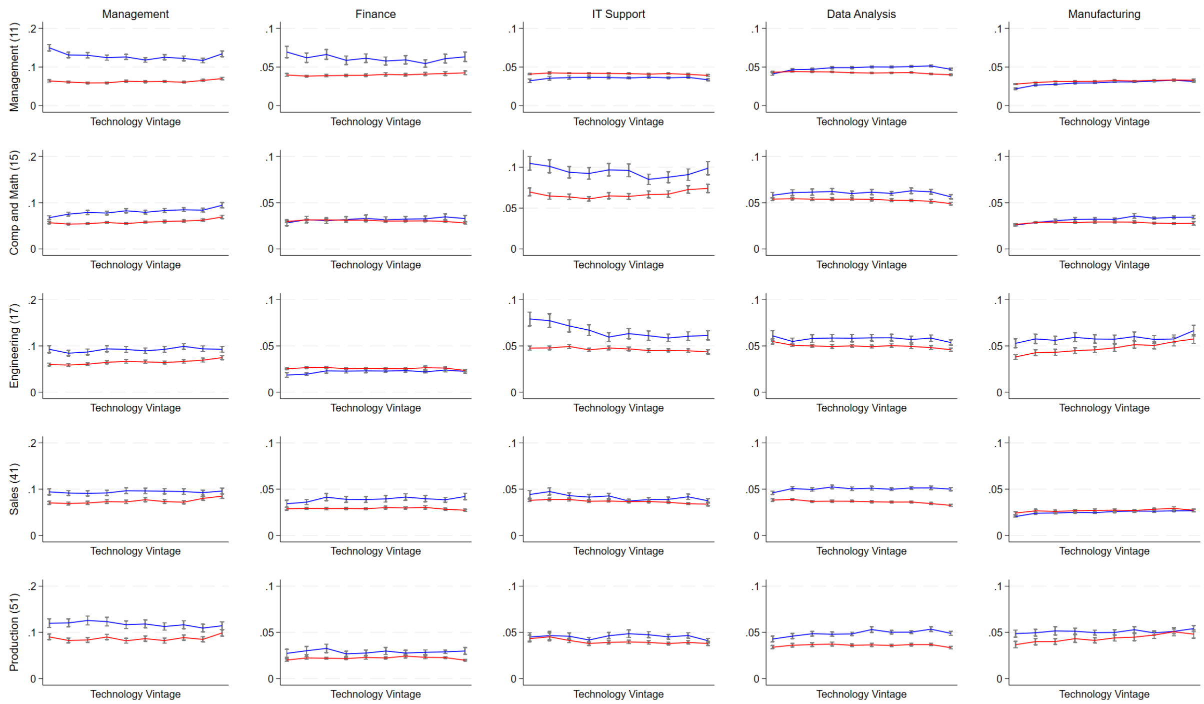

Figure 5: Skill Vector Weights for {Occupation x Skill} Pairs
This figure presents the distribution of skill vector weights for firm job requirements (blue) and skill vector weights for newly hired employees (red), across five rows of 2-digit major standard occupational codes (SOC): Management (11), Computer and Mathematical (15), Architecture and Engineering (17), Sales (41), and Production (51). Each column represents an IT-complementary skill category (skill definitions provided in Appendix A2). The Y-axis depicts the average weight score for each skill group, the X-axis depicts technology vintage bins as defined in Figure 3.

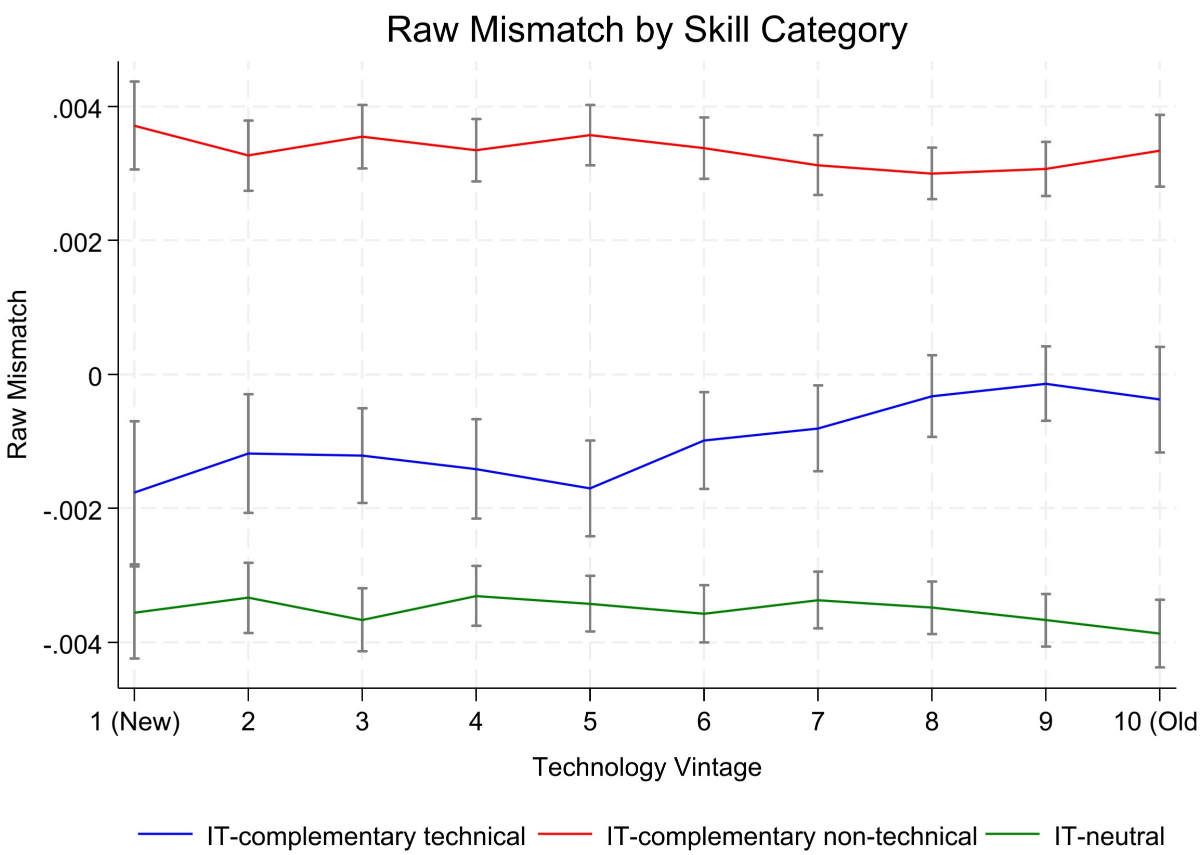

Figure 6: Raw Mismatch Scores by Skill Category

This figure presents the sum of raw differences between firm skill requirements and worker skill attributes, for three aggregate categories of skills: IT-complementary technical, IT-complementary nontechnical, and IT-neutral skills (skill definitions and groupings are provided in Appendix A2). The Y-axis depicts the average difference for each skill group, the X-axis depicts technology vintage bins as defined in Figure 3.

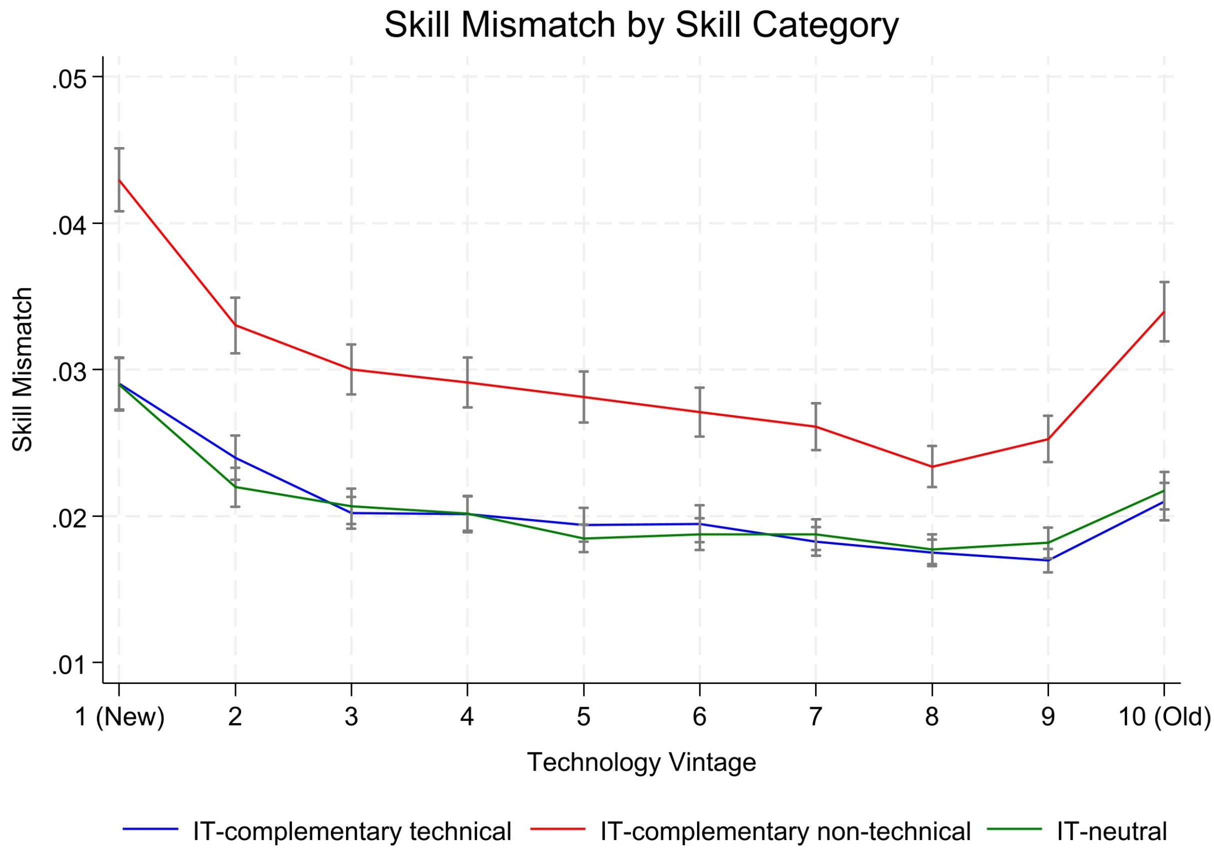

Figure 7: Skill Mismatch Across Skill Categories
This figure presents skill mismatch (Euclidean distance measures) between firm skill requirements and worker skill attributes across three categories of aggregated skills: IT-complementary technical, IT-complementary nontechnical, and IT-neutral skills (skill definitions and groupings are provided in Appendix A2). The Y-axis depicts the average difference for each skill group, the X-axis depicts technology vintage bins as defined in Figure 3.

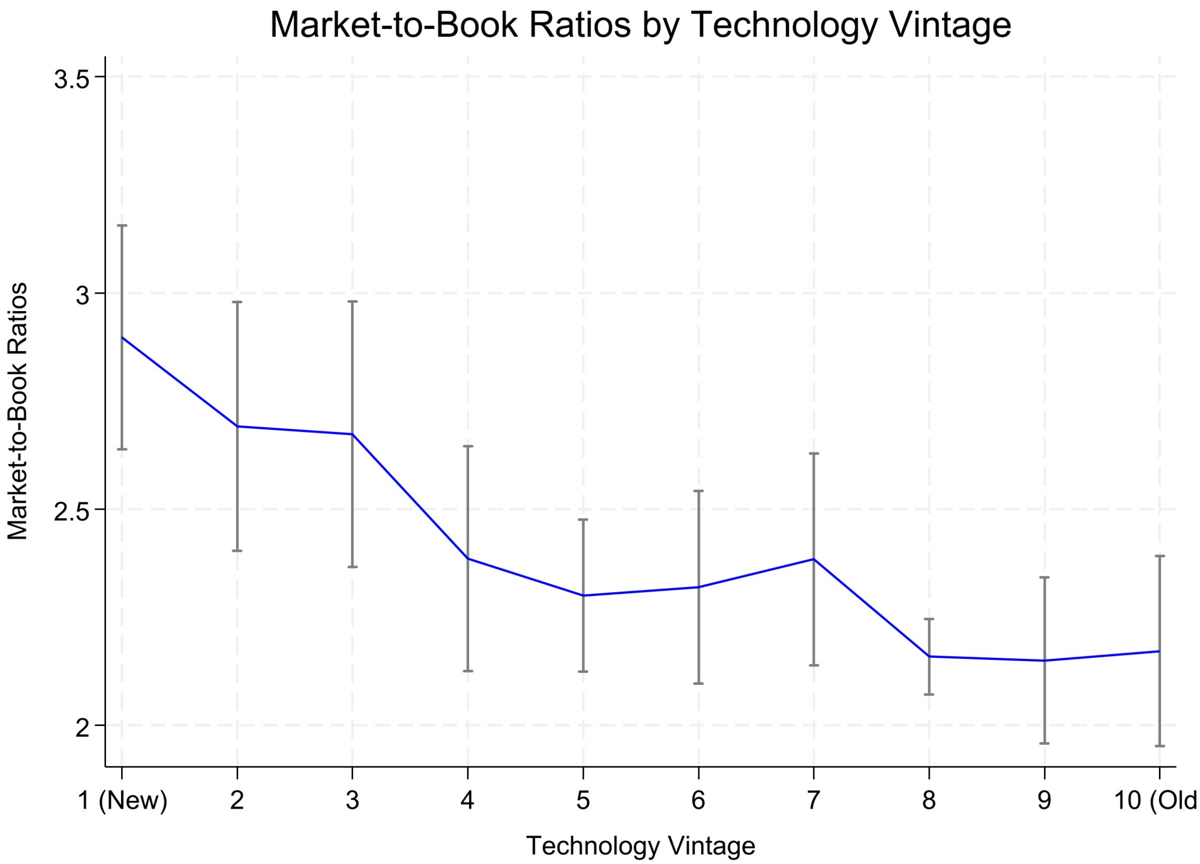

Figure 8: Market-to-Book Ratios across Technology Vintages
This figure depicts average market-to-book ratios for firms that deploy information technologies of different vintages. The X-axis depicts technology vintage bins as defined in Figure 3; the Y-axis depicts the average market-to-book ratio for sample firms in a given bin.

Table 1: Summary Statistics: Sample vs Population

Variable	N	Mean	Std Dev	p25	Median	p75
Panel A: Our Sample						
Skill mismatch	22,467	0.159	0.166	0.0624	0.0762	0.201
Technology age	17,998	30.85	6.195	26.94	30.75	34.64
Sale (millions)	19,350	6,408	16,166	416.8	1,395	4,620
Profitability	19,461	0.105	0.114	0.0659	0.111	0.159
Firm Age	19,461	22.72	11.33	14	23	35
Total asset (millions)	19,351	8,665	22,677	497.6	1,678	6,000
Investment rate	18,990	0.0439	0.0479	0.0156	0.0292	0.0527
Panel B: Compustat Population						
Sale (millions)	124,905	2,630	8,451	6.367	126.9	1,089
Profitability	175,342	-0.238	1.347	-0.00332	0	0.0797
Firm Age	175,342	12.87	11.35	3	9	20
Total asset (millions)	125,471	7,435	28,560	35.36	376.0	2,475
Investment rate	95,830	0.0598	0.127	0.00622	0.0233	0.0576

Notes: This table compares summary statistics between our analysis sample and the broader Compustat population. The sample period runs from year 2010 through year 2024. Skill mismatch is the Euclidean distance between firm and worker skill vectors in a given year. Technology age is defined as per Figure 3. Panel A shows statistics for firms in our sample for which we observe job postings and employee skill data. Panel B shows statistics for all Compustat firms during the sample period. All dollar amounts are inflation-adjusted to 2010 dollars. All variables are winsorized at the 1st and 99th percentiles.

Table 2: Summary Statistics: Skill Mismatch

#	Skill group	Variable	N	Mean	Std Dev
#1	Mgmt / Leadership / Training	Firms	22,467	0.0766	0.0865
		Workers	22,467	0.0570	0.0374
		Mismatch	22,467	0.0400	0.0872
#2	Proc Improve / Prog Mgmt / X-Func Lead	Firms	22,467	0.0720	0.0440
		Workers	22,467	0.0602	0.0279
		Mismatch	22,467	0.0283	0.0447
#3	Networking / Troubleshooting / Tech Support	Firms	22,467	0.0583	0.0559
		Workers	22,467	0.0452	0.0204
		Mismatch	22,467	0.0263	0.0543
#4	Proj Mgmt / Planning / Construction	Firms	22,467	0.0581	0.0447
		Workers	22,467	0.0527	0.0210
		Mismatch	22,467	0.0255	0.0425
#5	Biz Analysis / Change Mgmt / Integration	Firms	22,467	0.0543	0.0216
		Workers	22,467	0.0524	0.0138
		Mismatch	22,467	0.0158	0.0193
#6	Strat Plan / Team Lead / Mktg	Firms	22,467	0.0524	0.0356
		Workers	22,467	0.0606	0.0314
		Mismatch	22,467	0.0248	0.0416
#7	MS Office / Cust Svc / Excel	Firms	22,467	0.0468	0.0476
		Workers	22,467	0.0565	0.0413
		Mismatch	22,467	0.0296	0.0560
#8	Data Analysis / C++	Firms	22,467	0.0442	0.0250
		Workers	22,467	0.0470	0.0132
		Mismatch	22,467	0.0157	0.0230
#9	Fin Analysis / Finance / Accounting	Firms	22,467	0.0441	0.0861
		Workers	22,467	0.0378	0.0279
		Mismatch	22,467	0.0299	0.0850
#10	Ops Mgmt / Retail / Inventory	Firms	22,467	0.0441	0.0427
		Workers	22,467	0.0436	0.0153
		Mismatch	22,467	0.0193	0.0398
#11	Sales / Biz Dev / Acct Mgmt	Firms	22,467	0.0434	0.0460
		Workers	22,467	0.0416	0.0199
		Mismatch	22,467	0.0211	0.0442
#12	Mfg / CI / AutoCAD	Firms	22,467	0.0407	0.0299
		Workers	22,467	0.0338	0.0163
		Mismatch	22,467	0.0170	0.0285
#13	SQL / SAP / MS SQL	Firms	22,467	0.0394	0.0245
		Workers	22,467	0.0449	0.0139
		Mismatch	22,467	0.0166	0.0227

Continued on next page

Summary Statistics: Skill Mismatch (continued)

#	Skill group	Variable	N	Mean	Std Dev
#14	Eng / Testing / Energy	Firms	22,467	0.0388	0.0389
		Workers	22,467	0.0383	0.0169
		Mismatch	22,467	0.0181	0.0372
#15	Event Planning / PR / Event Mgmt	Firms	22,467	0.0351	0.0253
		Workers	22,467	0.0409	0.0183
		Mismatch	22,467	0.0178	0.0256
#16	Telecom / Electronics / Net Sec	Firms	22,467	0.0329	0.0196
		Workers	22,467	0.0345	0.0113
		Mismatch	22,467	0.0138	0.0169
#17	Public Speaking / Research / Coaching	Firms	22,467	0.0319	0.0191
		Workers	22,467	0.0426	0.0192
		Mismatch	22,467	0.0176	0.0228
#18	Hospitality / Food / Hosp Mgmt	Firms	22,467	0.0294	0.0596
		Workers	22,467	0.0287	0.0157
		Mismatch	22,467	0.0174	0.0560
#19	Video Prod / Edit / Video	Firms	22,467	0.0277	0.0229
		Workers	22,467	0.0412	0.0205
		Mismatch	22,467	0.0197	0.0264
#20	Photoshop / Design / Adobe Suite	Firms	22,467	0.0244	0.0476
		Workers	22,467	0.0203	0.0148
		Mismatch	22,467	0.0147	0.0474
#21	Java / HTML / JS	Firms	22,467	0.0227	0.0217
		Workers	22,467	0.0317	0.0212
		Mismatch	22,467	0.0165	0.0266
#22	Healthcare / Med Devices / Hospitals	Firms	22,467	0.0225	0.0200
		Workers	22,467	0.0232	0.0136
		Mismatch	22,467	0.0114	0.0194
#23	Risk Mgmt / Banking / Insurance	Firms	22,467	0.0211	0.0115
		Workers	22,467	0.0295	0.0119
		Mismatch	22,467	0.0132	0.0128
#24	Pharma / Biotech / GMP	Firms	22,467	0.0206	0.0224
		Workers	22,467	0.0178	0.0121
		Mismatch	22,467	0.00963	0.0221
#25	Legal Res / Writing / Corp Law	Firms	22,467	0.0185	0.0121
		Workers	22,467	0.0179	0.0166
		Mismatch	22,467	0.00956	0.0180

Notes: This table presents sample statistics describing skill vector weights for each of the 25 skill groups in the sample (skill definitions and groupings are provided in Appendix A2). For each skill group, skill weights are computed for firm job postings in a given year (demand) and for workers who join a firm in a given year (supply). Mismatch is the squared difference in skill vector weights between a firm and its new employees (i.e. demand minus supply) in a given year.

Table 3: Skill Mismatch Across Technology Vintages

	(1)	(2)	(3)	(4)	(5)	(6)
TechVint	-2.292***		-2.317***		-0.811***	
	(0.173)		(0.169)		(0.165)	
TechVint^2	3.464***		3.479***		1.241***	
	(0.271)		(0.262)		(0.254)	
TechVint2		-0.0445***		-0.0441***		-0.0120**
		(0.00650)		(0.00638)		(0.00573)
TechVint3		-0.0617***		-0.0621***		-0.0248***
		(0.00641)		(0.00629)		(0.00557)
TechVint4		-0.0608***		-0.0616***		-0.0220***
		(0.00654)		(0.00642)		(0.00585)
TechVint5		-0.0632***		-0.0652***		-0.0239***
		(0.00670)		(0.00659)		(0.00591)
TechVint6		-0.0625***		-0.0648***		-0.0208***
		(0.00700)		(0.00706)		(0.00636)
TechVint7		-0.0637***		-0.0654***		-0.0263***
		(0.00687)		(0.00687)		(0.00630)
TechVint8		-0.0713***		-0.0735***		-0.0246***
		(0.00707)		(0.00721)		(0.00651)
TechVint9		-0.0637***		-0.0660***		-0.0207***
		(0.00743)		(0.00766)		(0.00693)
TechVint10		-0.0298***		-0.0335***		-0.0126*
		(0.00812)		(0.00830)		(0.00732)
Observations	17,978	17,978	17,978	17,978	17,800	17,800
R-squared	0.045	0.040	0.062	0.056	0.439	0.439
Year FE	YES	YES	YES	YES	YES	YES
Industry/Firm FE	NO	NO	Ind(2-digit)	Ind(2-digit)	Firm	Firm

Notes: This table presents regression coefficient estimates of skill mismatch as a function of technology vintage. TechVint is defined as technology vintage, normalized by 100, with a quadratic term included in odd-numbered columns. Even-numbered columns depict piece-wise linear regression coefficients for indicators of technology vintage bins (described in Figure 3). Fixed effects for the year of observation, industry (2-digit SIC code), and firm are included. Standard errors are heteroskedasticity consistent and clustered by firm.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Skill Mismatch and Technology Vintage — Occupation-Level Analysis

ONET	Constant	SE	TechVint	SE	TechVint^2	SE	Observations	R-squared
Panel A: Occupations with Positive and Significant Skill Mismatch Curvature								
11	0.442***	0.044	-1.417***	0.270	2.127***	0.412	15221	0.454
13	0.461***	0.049	-1.434***	0.302	2.229***	0.464	14403	0.456
15	0.362***	0.046	-0.876***	0.279	1.271***	0.426	13752	0.487
17	0.420***	0.065	-0.897**	0.393	1.177**	0.593	10397	0.512
31	0.640***	0.151	-1.669*	0.872	2.438*	1.283	1597	0.565
37	0.673***	0.122	-1.474**	0.741	2.070*	1.132	2491	0.501
41	0.370***	0.047	-0.797***	0.282	1.198***	0.423	12184	0.505
43	0.411***	0.048	-0.992***	0.291	1.512***	0.435	13788	0.461
51	0.425***	0.062	-0.780**	0.362	1.008*	0.533	9651	0.511
Panel B: Occupations with Negative or Insignificant Skill Mismatch Curvature								
19	0.399***	0.084	-0.368	0.490	0.357	0.717	7506	0.463
21	0.435***	0.119	-0.490	0.706	1.003	1.054	3212	0.499
23	0.535***	0.119	-0.903	0.716	1.447	1.074	5152	0.482
25	0.341**	0.141	0.121	0.851	0.003	1.271	4410	0.423
27	0.471***	0.081	-0.804	0.495	1.011	0.754	9699	0.453
29	0.331***	0.090	-0.016	0.536	0.061	0.801	5209	0.553
33	0.405***	0.111	-0.373	0.642	0.658	0.924	3359	0.512
35	0.498***	0.124	-0.709	0.688	1.159	0.945	3658	0.492
39	0.356***	0.117	-0.045	0.654	0.151	0.913	2376	0.492
45	0.100	0.326	1.326	1.799	-1.074	2.492	527	0.521
47	0.167	0.110	1.153*	0.645	-1.892**	0.950	4792	0.474
49	0.299***	0.072	-0.045	0.413	0.141	0.593	8154	0.528
53	0.328***	0.066	-0.241	0.383	0.332	0.559	8462	0.539

Notes: This table presents regression coefficient estimates of skill mismatch on a quadratic function of technology vintage for each 2-digit SOC occupational code (row) in the sample. Panel A (B) depicts the results for occupational codes where the quadratic term is positive and statistically significant (negative and/or insignificant). Standard errors are clustered by firm.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 5: Skill Mismatch and Technology Vintage: IT-Complementary Non-Technical Skills

	(1)	(2)	(3)	(4)	(5)	(6)
TechVint	-0.496***		-0.518***		-0.497***	
	(0.0417)		(0.0409)		(0.0407)	
TechVint^2	0.766***		0.790***		0.757***	
	(0.0663)		(0.0644)		(0.0639)	
TechVint2		-0.00899***		-0.00908***		-0.00862***
		(0.00154)		(0.00151)		(0.00146)
TechVint3		-0.0121***		-0.0126***		-0.0120***
		(0.00147)		(0.00145)		(0.00143)
TechVint4		-0.0120***		-0.0126***		-0.0125***
		(0.00151)		(0.00148)		(0.00145)
TechVint5		-0.0119***		-0.0129***		-0.0122***
		(0.00153)		(0.00152)		(0.00150)
TechVint6		-0.0124***		-0.0136***		-0.0133***
		(0.00159)		(0.00160)		(0.00156)
TechVint7		-0.0123***		-0.0135***		-0.0132***
		(0.00155)		(0.00156)		(0.00153)
TechVint8		-0.0143***		-0.0156***		-0.0145***
		(0.00158)		(0.00161)		(0.00159)
TechVint9		-0.0118***		-0.0133***		-0.0127***
		(0.00169)		(0.00174)		(0.00170)
TechVint10		-0.00385**		-0.00550***		-0.00527***
		(0.00190)		(0.00193)		(0.00190)
Observations	17,978	17,978	17,978	17,978	17,975	17,975
R-squared	0.037	0.032	0.053	0.047	0.103	0.099
Year FE	YES	YES	YES	YES	YES	YES
Industry/Firm FE	NO	NO	Ind(2-digit)	Ind(2-digit)	Ind(4-digit)	Ind(4-digit)

Notes: This table presents regression coefficient estimates of skill mismatch as a function of technology vintage, where skill mismatch is measured only across IT-complementary non-technical skills (skill definitions and groupings are provided in Appendix A2). TechVint is defined as technology vintage, normalized by 100, with a quadratic term included in odd-numbered columns. Even-numbered columns depict piece-wise linear regression coefficients for indicators of technology vintage bins (described in Figure 3). Fixed effects for the year of observation, industry (2-digit and 4-digit SIC code) are included. Standard errors are heteroskedasticity consistent and clustered by firm.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: Skill Mismatch and Technology Vintage: IT-Complementary Technical Skills

	(1)	(2)	(3)	(4)	(5)	(6)
TechVint	-0.231***		-0.220***		-0.209***	
	(0.0253)		(0.0253)		(0.0255)	
TechVint^2	0.335***		0.327***		0.314***	
	(0.0370)		(0.0365)		(0.0365)	
TechVint2		-0.00414***		-0.00380***		-0.00335***
		(0.00102)		(0.00100)		(0.000983)
TechVint3		-0.00714***		-0.00669***		-0.00601***
		(0.000978)		(0.000970)		(0.000958)
TechVint4		-0.00733***		-0.00670***		-0.00617***
		(0.000995)		(0.000998)		(0.000987)
TechVint5		-0.00729***		-0.00663***		-0.00607***
		(0.00102)		(0.00102)		(0.00101)
TechVint6		-0.00707***		-0.00630***		-0.00564***
		(0.00105)		(0.00106)		(0.00106)
TechVint7		-0.00735***		-0.00642***		-0.00580***
		(0.00102)		(0.00103)		(0.00102)
TechVint8		-0.00751***		-0.00648***		-0.00563***
		(0.00105)		(0.00110)		(0.00110)
TechVint9		-0.00765***		-0.00636***		-0.00544***
		(0.00102)		(0.00106)		(0.00105)
TechVint10		-0.00458***		-0.00330***		-0.00241**
		(0.00107)		(0.00112)		(0.00112)
Observations	17,978	17,978	17,978	17,978	17,975	17,975
R-squared	0.031	0.030	0.046	0.045	0.082	0.080
Year FE	YES	YES	YES	YES	YES	YES
Industry/Firm FE	NO	NO	Ind(2-digit)	Ind(2-digit)	Ind(4-digit)	Ind(4-digit)

Notes: This table presents regression coefficients of skill mismatch as a function of technology vintage, restricted to IT-complementary technical skills (skill group definitions in Appendix A2). TechVint is normalized by 100. Odd-numbered columns include a quadratic term; even-numbered columns show piece-wise linear regression using technology vintage bins (Figure 3). Year, industry (2-digit and 4-digit SIC) fixed effects included. Standard errors are heteroskedasticity-consistent and clustered by firm.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 7: Skill Mismatch and Technology Vintage: IT-Neutral Skills

	(1)	(2)	(3)	(4)	(5)	(6)
TechVint	-0.252***		-0.242***		-0.232***	
	(0.0260)		(0.0258)		(0.0257)	
TechVint^2	0.366***		0.347***		0.332***	
	(0.0391)		(0.0383)		(0.0381)	
TechVint2		-0.00605***		-0.00601***		-0.00559***
		(0.00110)		(0.00108)		(0.00107)
TechVint3		-0.00753***		-0.00746***		-0.00732***
		(0.00109)		(0.00106)		(0.00105)
TechVint4		-0.00762***		-0.00763***		-0.00760***
		(0.00110)		(0.00109)		(0.00108)
TechVint5		-0.00869***		-0.00881***		-0.00854***
		(0.00111)		(0.00108)		(0.00109)
TechVint6		-0.00861***		-0.00867***		-0.00864***
		(0.00112)		(0.00113)		(0.00114)
TechVint7		-0.00814***		-0.00810***		-0.00794***
		(0.00116)		(0.00116)		(0.00117)
TechVint8		-0.00923***		-0.00925***		-0.00880***
		(0.00119)		(0.00120)		(0.00122)
TechVint9		-0.00897***		-0.00897***		-0.00870***
		(0.00123)		(0.00126)		(0.00128)
TechVint10		-0.00568***		-0.00617***		-0.00593***
		(0.00131)		(0.00135)		(0.00137)
Observations	17,978	17,978	17,978	17,978	17,975	17,975
R-squared	0.018	0.018	0.039	0.039	0.083	0.083
Year FE	YES	YES	YES	YES	YES	YES
Industry/Firm FE	NO	NO	Ind(2-digit)	Ind(2-digit)	Ind(4-digit)	Ind(4-digit)

Notes: This table presents regression coefficients of skill mismatch as a function of technology vintage, restricted to IT-neutral skills (skill group definitions in Appendix A2). TechVint is normalized by 100. Odd-numbered columns include a quadratic term; even-numbered columns show piece-wise linear regression using technology vintage bins (Figure 3). Year, industry (2-digit and 4-digit SIC) fixed effects included. Standard errors are heteroskedasticity-consistent and clustered by firm.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 8: Firm Market-to-Book Ratio and Technology Vintage

	(1)	(2)	(3)	(4)	(5)	(6)
TechVint	-2.800***		-1.717***		-1.236***	
	(0.508)		(0.435)		(0.404)	
TechVint2		-0.0237		0.00456		0.000222
		(0.0871)		(0.0794)		(0.0790)
TechVint3		-0.220**		-0.105		-0.0677
		(0.0863)		(0.0768)		(0.0752)
TechVint4		-0.249***		-0.129		-0.0987
		(0.0958)		(0.0868)		(0.0825)
TechVint5		-0.339***		-0.190**		-0.139*
		(0.0978)		(0.0863)		(0.0816)
TechVint6		-0.316***		-0.130		-0.118
		(0.103)		(0.0901)		(0.0869)
TechVint7		-0.374***		-0.198**		-0.154*
		(0.107)		(0.0922)		(0.0876)
TechVint8		-0.346***		-0.163*		-0.135
		(0.112)		(0.0963)		(0.0926)
TechVint9		-0.492***		-0.249**		-0.166*
		(0.118)		(0.0994)		(0.0930)
TechVint10		-0.619***		-0.365***		-0.237**
		(0.117)		(0.103)		(0.0949)
Constant	3.018***	2.452***	2.682***	2.304***	2.533***	2.262***
	(0.167)	(0.0900)	(0.139)	(0.0745)	(0.129)	(0.0713)
Observations	16,337	16,337	16,337	16,337	16,333	16,333
R-squared	0.021	0.022	0.144	0.144	0.244	0.244
Year FE	YES	YES	YES	YES	YES	YES
Industry/Firm FE	NO	NO	Ind(2-digit)	Ind(2-digit)	Ind(4-digit)	Ind(4-digit)

Notes: This table presents regression coefficient estimates of firm market-to-book ratios as a function of technology vintage. *TechVint* is technology vintage, normalized by 100. Fixed effects for year and industry (two-digit and four-digit SIC codes) are included where specified. Standard errors are heteroskedasticity consistent and clustered by firm.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 9: Robustness Analysis for Skill Mismatch and Technology Vintage

	(1)	(2)	(3)	(4)	(5)	(6)
TechVint2	-0.0107	0.00357	0.0124	-0.00780	0.0148	-0.0164
	(0.0166)	(0.0161)	(0.0160)	(0.0106)	(0.0210)	(0.0161)
TechVint3	-0.0252	-0.0120	-0.0180	-0.0142	0.0351*	-0.0284*
	(0.0164)	(0.0159)	(0.0158)	(0.0105)	(0.0208)	(0.0160)
TechVint4	-0.0189	0.000330	-0.00633	-0.0118	0.0293	-0.0240
	(0.0168)	(0.0163)	(0.0162)	(0.0108)	(0.0213)	(0.0164)
TechVint5	-0.0357**	-0.00636	-0.0296*	-0.0285***	0.0559***	-0.0428***
	(0.0164)	(0.0159)	(0.0158)	(0.0105)	(0.0208)	(0.0160)
TechVint6	-0.0334**	-0.0260*	-0.0260*	-0.0268***	0.0492**	-0.0417***
	(0.0162)	(0.0157)	(0.0156)	(0.0104)	(0.0205)	(0.0158)
TechVint7	-0.0315*	-0.0260*	-0.0213	-0.0202*	0.0479**	-0.0386**
	(0.0162)	(0.0156)	(0.0156)	(0.0104)	(0.0205)	(0.0157)
TechVint8	-0.0387**	-0.0388**	-0.0423***	-0.0316***	0.0602***	-0.0512***
	(0.0160)	(0.0155)	(0.0154)	(0.0102)	(0.0203)	(0.0156)
TechVint9	-0.0279*	-0.0237	-0.0142	-0.0268***	0.0380*	-0.0345**
	(0.0160)	(0.0156)	(0.0154)	(0.0102)	(0.0202)	(0.0155)
TechVint10	-0.00519	0.0164	-0.000590	-0.00955	0.0137	-0.0133
	(0.0162)	(0.0157)	(0.0156)	(0.0104)	(0.0205)	(0.0158)
Constant	0.140***	0.335***	0.180***	0.275***	0.847***	0.202***
	(0.0123)	(0.0119)	(0.0118)	(0.00785)	(0.0155)	(0.0119)
Observations	1,417	1,431	1,417	1,417	1,417	1,417
R-squared	0.009	0.017	0.015	0.014	0.013	0.013

Notes: This table presents regression coefficient estimates of skill mismatch as a function of technology vintage bins (using a piecewise linear specification), using sample data from 2020. Column (1) presents our benchmark results from Table 3. Column (2) uses a measure of firm-year skill mismatch that is computed by first measuring skill mismatch between firms and workers at the occupation level in a given year, and then aggregating these occupation-level measures of skill mismatch to the firm-year level. Column (3) uses a measure of skill mismatch that is based on the 10-skill group classifications defined by Deming and Kahn (2018), instead of the 25-skill group classification otherwise used in this paper. Column (4) uses a measure of skill mismatch that is estimated using an alternative large language model—ModernBERT—instead of BERT. Column (5) uses a measure of skill mismatch that is based on cosine similarity between firm and worker skill vectors, rather than Euclidean distance, as per our main specifications. Column (6) presents results using skill mismatch defined as the Hellinger distance between firm and worker skill vectors. Standard errors are clustered by industry.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix

A Variable definitions

- *Skill mismatch*: As described in Section 4.2
- *Technology vintage*: Following the literature (e.g., Tambe (2014)), we compute the average (and median) vintage year of all the technologies that appear in the firms' job postings. This measure captures the stage of the technology cycle that corresponds to the information technologies that are being adopted by firms and listed in job postings to convey worker skill requirements.
- Market-to-book ratio: (AT+PRCC*CSHO-CEQ-TXDB)/AT
- Total asset: AT in 2010 dollar
- Sale: SALE in 2010 dollar
- Profitability: Operating Income Before Depreciation (OIBDP) scaled by total asset (AT)

B Additional Tables

Table A1: Most Frequently Listed Information Technologies by Technology Vintage Bin

Rank	TechVint1: less than 23.4	TechVint2: (23.4, 26.1)
1	Microsoft Outlook	JavaScript (Programming Language)
2	Java (Programming Language)	Microsoft Access
3	Salesforce	HyperText Markup Language (HTML)
4	Amazon Web Services	Windows Servers
5	Microsoft Azure	PostgreSQL
6	C# (Programming Language)	Visual Basic (Programming Language)
7	Microsoft SharePoint	SolidWorks (CAD)
Rank	TechVint3: (26.1, 27.8)	TechVint4: (27.8, 29.3)
1	Linux	Python (Programming Language)
2	Microsoft Visio	Microsoft SQL Servers
3	R (Programming Language)	Information Technology Infrastructure Library
4	PL/SQL	Learning Management Systems
5	Applicant Tracking Systems	Business Intelligence Tools
6	SAP ERP	Operational Data Store
7	Web Servers	Web Browsers
Rank	TechVint5: (29.3, 30.8)	TechVint6: (30.8, 32.2)
1	Microsoft Office	Microsoft PowerPoint
2	Dashboard	Adobe Illustrator
3	Perl (Programming Language)	Transact-SQL
4	Firewall	LabVIEW
5	Adobe Photoshop	MicroStation (CAD Design Software)
6	PeopleSoft Applications	Information Technology Architecture
7	Bash (Scripting Language)	Presentation Software
Rank	TechVint7: (32.2, 33.8)	TechVint8: (33.8, 35.6)
1	Microsoft Excel	Microsoft Word
2	C++ (Programming Language)	Microsoft Project
3	Microsoft Windows	MATLAB
4	Field-Programmable Gate Array (FPGA)	IBM DB2
5	Customer Relationship Management (CRM) Software	Verilog
6	IBM Maximo	Primavera (Software)
7	Objective-C (Programming Language)	OSI Models
Rank	TechVint9: (35.6, 38.2)	TechVint10: greater than 38.2
1	AutoCAD	SQL (Programming Language)
2	TurboTax	SAP Applications
3	Oracle Databases	Operating Systems
4	Mac OS	Application Programming Interface (API)
5	Bug Tracking And Management	Spreadsheets
6	Oracle Hyperion	Unix
7	QuickBooks (Accounting Software)	C (Programming Language)

Table A2: Skill Labels and Groupings

Skill groups	**25 Skill categories**
IT-Complementary Technical	java / html / javascript, structured query language / sap / microsoft sql server, analysis / data analysis / c++, networking / troubleshooting / technical support, telecommunications / electronics / network security
IT-Complementary Non-technical	business analysis / change management / integration, process improvement / program management / cross-functional team leadership, management / leadership / training, project management / project planning / construction, strategic planning / team leadership / marketing, manufacturing / continuous improvement / autocad, financial analysis / finance / accounting, operations management / retail / inventory management, risk management / banking / insurance, engineering / testing / energy, sales / business development / account management
IT-Neutral	microsoft office / customer service / microsoft excel, healthcare / medical devices / hospitals, pharmaceutical industry / biotechnology / GMP, public speaking / research / coaching, photoshop / graphic design / adobe creative suite, hospitality / food / hospitality management, video production / video editing / video, event planning / public relations / event management, legal research / legal writing / corporate law

C Skill Adjustment and Mismatch Dynamics

This appendix provides a simple microfoundation for the assumption that skill mismatch between the firm and its unskilled workers decreases as the number skilled employees increases.

C.1 Skill Vectors and Adding-Up Constraints

There are $N = 25$ skill dimensions indexed by $i = 1, \ldots, N$.

Firm f at time t has a vector of skill requirements

$$F_t = (F_{1,t}, \ldots, F_{N,t}),$$

and workers hired at time t have skill endowments

$$W_t = (W_{1,t}, \ldots, W_{N,t}).$$

Both vectors are normalized and consist of skill weights that sum to one:

$$\sum_{i=1}^{N} F_{i,t} = 1, \tag{1}$$

$$\sum_{i=1}^{N} W_{i,t} = 1. \tag{2}$$

All components are nonnegative; i.e. all elements satisfy $F_{i,t} \geq 0$ and $W_{i,t} \geq 0$.

C.2 Mismatch Measure

Skill mismatch is defined as the Euclidean distance:

$$\widetilde{M}_t = \left(\sum_{i=1}^{N} (F_{i,t} - W_{i,t})^2 \right)^{1/2}.$$

For analytical convenience, we work with squared mismatch:

$$\mathcal{M}_t \equiv \sum_{i=1}^{N} (F_{i,t} - W_{i,t})^2. \tag{3}$$

Since the square root is monotone, minimizing $\mathcal{M}_t$ is equivalent to minimizing $\widetilde{M}_t$.

C.3 Firm Adjustment Problem

Between t and $t+1$, firms adjust skill requirements. Let $S_t > 0$ denote the number of skilled workers employed by the firm and let κ represent an adjustment efficiency parameter .

The firm chooses $\{F_{i,t+1}\}_{i=1}^{N}$ to solve:

$$\min_{\{F_{i,t+1}\}} \quad \underbrace{\sum_{i=1}^{N}(F_{i,t+1}-W_{i,t+1})^2}_{\text{mismatch}} + \frac{1}{2\kappa S_t}\underbrace{\sum_{i=1}^{N}(F_{i,t+1}-F_{i,t})^2}_{\text{adjustment costs}} \tag{4}$$

$$\text{s.t.} \ \sum_{i=1}^{N} F_{i,t+1} = 1.$$

Adjustment costs decline in S_t, reflecting greater internal capacity to reorganize skill requirements.

C.4 Lagrangian

The Lagrangian is:

$$\mathcal{L} = \sum_{i=1}^{N}(F_{i,t+1}-W_{i,t+1})^2 + \frac{1}{2\kappa S_t}\sum_{i=1}^{N}(F_{i,t+1}-F_{i,t})^2 + \lambda\left(\sum_{i=1}^{N} F_{i,t+1} - 1\right). \tag{5}$$

C.5 First-Order Conditions

For each i:

$$\frac{\partial \mathcal{L}}{\partial F_{i,t+1}} = 2(F_{i,t+1}-W_{i,t+1}) + \frac{1}{\kappa S_t}(F_{i,t+1}-F_{i,t}) + \lambda = 0. \tag{6}$$

Rearranging:

$$\left(2+\frac{1}{\kappa S_t}\right)F_{i,t+1} = 2W_{i,t+1} + \frac{1}{\kappa S_t}F_{i,t} - \lambda. \tag{7}$$

C.6 Solving for the Multiplier

Summing (7) over i and using $\sum_i F_{i,t+1} = 1$, $\sum_i F_{i,t} = 1$, and $\sum_i W_{i,t+1} = 1$:

$$\left(2+\frac{1}{\kappa S_t}\right) = 2 + \frac{1}{\kappa S_t} - N\lambda,$$

which implies

$$\lambda = 0.$$

C.7 Adjustment Rule

Substituting $\lambda = 0$:

$$F_{i,t+1} = \frac{2}{2+\frac{1}{\kappa S_t}}W_{i,t+1} + \frac{\frac{1}{\kappa S_t}}{2+\frac{1}{\kappa S_t}}F_{i,t}.$$

Define

$$\alpha(S_t) = \frac{2\kappa S_t}{1+2\kappa S_t}, \quad 1-\alpha(S_t) = \frac{1}{1+2\kappa S_t}.$$

Then

$$F_{i,t+1} = (1-\alpha(S_t))F_{i,t} + \alpha(S_t)W_{i,t+1}. \tag{8}$$

C.8 Mismatch After Adjustment

Subtracting $W_{i,t+1}$:

$$F_{i,t+1} - W_{i,t+1} = (1-\alpha(S_t))(F_{i,t} - W_{i,t+1}).$$

Squaring and summing across skills:

$$\begin{aligned}
\mathcal{M}_{t+1} &= \sum_{i=1}^{N}(F_{i,t+1} - W_{i,t+1})^2 && (9)\\
&= \sum_{i=1}^{N}(1-\alpha(S_t))^2(F_{i,t} - W_{i,t+1})^2 && (10)\\
&= (1-\alpha(S_t))^2\sum_{i=1}^{N}(F_{i,t} - W_{i,t+1})^2. && (11)
\end{aligned}$$

The inner summation is strictly positive unless skill requirements exactly match worker skills.

C.9 Comparative Statics with Respect to Skilled Workers

Using $1-\alpha(S_t) = \frac{1}{1+2\kappa S_t}$, rewrite (11) as:

$$\mathcal{M}_{t+1} = \frac{1}{(1+2\kappa S_t)^2}\sum_{i=1}^{N}(F_{i,t} - W_{i,t+1})^2.$$

First derivative.

$$\frac{\partial \mathcal{M}_{t+1}}{\partial S_t} = -\frac{4\kappa}{(1+2\kappa S_t)^3}\sum_{i=1}^{N}(F_{i,t} - W_{i,t+1})^2 < 0, \tag{12}$$

since the summation term is strictly positive.

Thus mismatch between the firm and its unskilled workers is decreasing in the number of skilled workers S_t employed by the firm.

Second derivative.

$$\frac{\partial^2 \mathcal{M}_{t+1}}{\partial S_t^2} = \frac{24\kappa^2}{(1+2\kappa S_t)^4} \sum_{i=1}^{N} (F_{i,t} - W_{i,t+1})^2 > 0. \tag{13}$$

Thus mismatch decreases in S_t at a diminishing rate.