Skill Mismatch Over The Technology Lifecycle

Ashwini K. Agrawal Daniel Kim Prasanna Tambe*

November 19, 2025

Abstract

Despite rapid advances in information technology (IT), technology diffusion remains slow and uneven across firms. This paper provides new empirical evidence that skill mismatch—the misalignment between firms' skill requirements and workers' capabilities—acts as a central friction in the diffusion process. Using large language models to construct granular measures of skill demand and supply from matched data on firm job postings and worker résumés, we test the predictions of a model in which skill mismatch evolves endogenously with firms' investments in different IT vintages. We document a U-shaped pattern: skill mismatch is highest when new technologies are introduced, declines as firms and workers adjust, and rises again as technologies become obsolete. These gaps are substantial not only for technical skills but also for IT-complementary nontechnical skills such as managerial and organizational capabilities. We further show that firms invest more in intangible capital and exhibit lower measured productivity early in the technology lifecycle when mismatch is greatest. The findings illustrate how skill mismatch plays a critical role in the diffusion of frontier technologies.

1 Introduction

Technological progress has accelerated in recent decades, yet the diffusion of new technologies across firms remains strikingly slow and uneven (Kalyani et al., 2025). A growing literature points to labor market frictions as a central factor behind this diffusion gap. In particular, many observers argue that the adoption of frontier technologies is constrained by the limited

*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.

supply of workers with the requisite specialized skills. Consistent with this view, the business press is replete with accounts of firms making extraordinary efforts—and facing substantial difficulties—to recruit workers with technology-complementary capabilities. Survey evidence echoes these concerns: the annual Duke/FRB CFO Survey consistently reports that one of the top challenges cited by CFOs is the difficulty of finding high-quality employees.

Despite the prominence of talent scarcity in discussions of technological change, however, we know relatively little about *how* labor-market frictions shape the technology diffusion process. Several distinct mechanisms are plausible. Early adopters of new technologies may be able to attract the most capable workers, while later adopters struggle to fill specialized positions. Alternatively, firms operating legacy systems may face fewer hiring frictions simply because their technologies have existed long enough for workers to develop familiarity with them. The nature of skill scarcity also remains poorly understood: are technical skills the primary bottleneck, or do shortages of complementary nontechnical skills—such as those required for organizational change and process redesign—play an equal or greater role? These open questions illustrate a large gap in our understanding of technology diffusion: we know little about how skill mismatch evolves over the technology lifecycle, which skills are most or least scarce, and how these frictions shape the productivity gains associated with IT adoption. Filling this gap is essential for understanding both the diffusion of new technologies and the growing dispersion of productivity across firms.

Progress on these questions has been limited by a combination of theoretical and empirical challenges. Theoretically, existing models of technology diffusion typically abstract from the rich, multidimensional nature of skill requirements and often treat workers' skills as exogenous, overlooking the incentives workers face when investing in new capabilities. Empirically, three obstacles have hindered progress. First, available data rarely provide sufficiently granular information on the skill attributes of firms and workers; measures based on education or occupation cannot distinguish between firms with broadly similar workforce characteristics but very different technology-specific needs. Second, identifying the precise skills that

firms demand is conceptually difficult, as many relevant capabilities are unobservable and not directly inferable from coarse job descriptions. A posting requiring "Java programming," for example, may reflect demand for a narrow technical task, broad problem-solving ability, or interpersonal skills needed for coordination with software engineers. Third, the computational complexity of processing thousands of heterogeneous skill attributes has historically made it challenging to construct consistent, comparable measures of skill mismatch across firms and technologies.

To address these challenges, we develop a simple model of technology diffusion with multidimensional skill mismatch. Our framework builds on the seminal work of Chari and Hopenhayn (1991), who provide a general equilibrium theory of technology diffusion and worker skill acquisition. We extend their model in three key ways. First, we introduce an explicit, multidimensional measure of skill mismatch—defined as the distance between the vector of skills required by firms and the vector of skills possessed by workers—capturing both technical and nontechnical dimensions of human capital. Second, we use skill mismatch to microfound the productivity dynamics assumed in the original model, and we connect these dynamics to firms' IT investments in tangible and intangible assets; these assumptions are grounded in a large literature on technology diffusion and productivity. Finally, we incorporate intuitive assumptions governing how firms and workers respond to mismatch, thereby yielding a set of sharp, testable predictions about how skill gaps evolve with technology vintage and how they shape investment dynamics and measured productivity. These predictions form the basis of our empirical analysis.

To test these predictions, we construct a novel matched employer–employee dataset that links U.S. firms' online job postings with workers' online résumés. The job postings provide detailed information on firms' demanded skills and their adoption of specific information technologies, extending data commonly used in studies such as Hershbein and Kahn (2018) and Deming and Kahn (2018). The résumé data allow us to observe workers' skill portfolios with comparable granularity. Linking the two sources enables us to measure the skill

attributes of new hires in close temporal proximity to firms' posted skill requirements.

We then apply large language models (LLMs) and natural language processing (NLP) techniques to convert unstructured text into high-dimensional representations of skill demand and supply. In particular, we use BERT (Bidirectional Encoder Representations from Transformers) (Devlin et al., 2019) to classify job requirements and worker skill portfolios into vectors spanning 25 skill categories—including managerial, analytical, engineering, financial, and operational capabilities. We quantify skill mismatch as the distance between firms' required skill vectors and the corresponding skill vectors of newly hired workers, using standard metrics from the NLP literature. This procedure yields firm-year measures of average skill mismatch that align with our conceptual framework, capture both technical and complementary nontechnical skills, and scale seamlessly across heterogeneous technologies, occupations, and time periods.

We document a number of empirical findings that confirm the predictions of our model. Our first main finding is that skill mismatch follows a U-shaped pattern over the technology lifecycle. Mismatch is high when firms initially adopt new technologies, declines as firms and workers gradually adjust, and rises again as technologies become obsolete. This pattern emerges consistently across both cross-sectional and within-firm analyses, and is especially pronounced among younger, smaller, and financially constrained firms—those that prior research has identified as struggling most with technology implementation. These results are robust to a wide range of specifications, including alternative LLMs, skill categorizations, and hiring windows around job postings.

Our second main finding concerns which skills are scarce and which are abundant as firms' IT systems evolve. We find persistent shortages in certain technical skills—such as advanced manufacturing and network security—as well as in complementary non-technical skills like management and financial analysis. In both cases, the associated gaps follow the same U-shaped pattern over the technology lifecycle. By contrast, skills tied to older or more routine tasks (e.g., hospitality, sales, or legacy programming) show little or no mismatch, and

in some cases appear oversupplied. These patterns indicate that mismatch is not confined to technical skills, but reflects a broader scarcity of capabilities that complement modern production systems. The evidence is consistent with recent observations that firms struggle to hire workers who combine managerial or financial expertise with sufficient technological knowledge to meet evolving job requirements.

Our third main finding is that firms respond to higher skill mismatch by increasing capital investments that are associated with worker skill acquisition. For example, we see that mismatch is associated with increased investment in intangible investments that encompass expenditures for training, software customization, and organizational work practices. This result aligns with recent arguments that general-purpose technologies require substantial complementary investments before yielding productivity gains (Brynjolfsson et al., 2021). Our evidence suggests that part of this investment reflects efforts to mitigate mismatch—for example, through worker training and other intangible assets that help firms adapt production when skill gaps are present. Consistent with these arguments, we further document that measured productivity increases with technology vintage, illustrating the time lag associated with payoffs from investments in intangible assets.

Our paper contributes to three strands of literature. First, we add to the literature on technology diffusion and productivity, which seeks to understand the forces shaping firms' adoption of different technologies. Our empirical findings provide novel microevidence that skill mismatch in the labor market is an important channel underlying the diffusion of technologies and the dispersion of productivity throughout the economy. Second, we advance the nascent but growing academic literature on skill mismatch, which examines the causes and consequences of misalignment between workers' skills and firms' needs——a concern increasingly emphasized by corporate managers and policymakers. Our paper offers a unique theoretical and empirical analysis that endogenously links firms' technology choices with worker skill mismatch. Third, our paper makes a methodological contribution to the literature that emphasizes the importance of viewing human capital as a multidimensional asset.

We develop a new empirical approach to characterizing worker and firm skill attributes, providing a tractable way to quantify dimensions of human capital that are otherwise difficult to capture using standard approaches.

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

2 Theoretical Framework

We develop a theoretical framework that incorporates multidimensional skill mismatch into the classic theory of technology diffusion and vintage human capital developed by Chari and Hopenhayn (1991). Theirs is a general equilibrium model of how the stock of unidimensional human capital—skilled vs. unskilled—in an economy endogenously evolves as new, more productive technologies continuously arise and become adopted by firms. This theory serves as a workhorse model that has been used in many other papers, such as xyz. We follow a similar approach to these papers by adapting the model to analyze how multidimensional skill mismatch endogenously relates to information technology diffusion.

Before presenting our adaptations, we briefly summarize the relevant aspects of Chari and Hopenhayn (1991) for our framework. 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, where 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); the 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. Unskilled workers and skilled workers are complements: the productivity of an unskilled

worker increases in the number of skilled workers who are simultaneously working in the same firm. 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. Specifically, we define skill mismatch between firms and workers as the discrepancy between firms' multidimensional skill requirements and workers' multidimensional skill attributes. This notion of skill mismatch has precedence in the academic literature (see for example, Lindenlaub and Postel-Vinay (2023)) 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 built around the technology. 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).[1]

[1]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).

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. Skills are imperfectly substitutable and include IT-complementary skills such as programming skills and managerial ability, as well as IT-neutral skills such as public speaking and event planning. This approach to modelling multidimensional human capital reflects the skill-weights approach of Lazear (2009).

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 define 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 i=1,2,...N). The skill mismatch $M^{\tau,s}$ of a firm-worker pairing of this type will thus be *zero*. We define 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*. In Lazear (2009), firm-specific human capital constitutes the unique weights that a firm places upon the skills that it requires of its workers. The skill-weights approach is well-suited to our model because it enables us to recast technology vintage-specific human capital from the original framework of Chari and Hopenhayn (1991) into multidimensional skill mismatch.

Our second adaptation is to use our definition of skill mismatch to microfound the production complementarity between unskilled and skilled workers in Chari and Hopenhayn (1991). We assume that as the firm employs more skilled workers s_τ, the firm becomes bet-

ter able to assign tasks and facilitate knowledge transfer to unskilled workers so that the skill requirements facing unskilled workers $F^{\tau,u}$ align more closely with unskilled workers' skill endowments $W^{\tau,u}$—thus reducing skill mismatch $M^{\tau,u}$. Formally, this means $\frac{\partial M^{\tau,u}}{\partial S_t} < 0$. We further 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 $\frac{\partial^2 f_U}{\partial (M^{\tau,u})^2} < 0$. Intuitively, an increase in the number of skilled employees allows firms to raise the marginal productivity of unskilled workers by altering skill requirements and reducing their skill mismatch.

Our third adaptation of Chari and Hopenhayn (1991) is to introduce capital investment as the mechanism through which firms respond to skill mismatch and facilitate workers' acquisition of IT-complementary skills. This approach is motivated by a large literature showing that firms' IT expenditures are key drivers of workers' IT-complementary human capital. Such expenditures include investments in intangible assets—training, organizational structures, and work practices—as well as tangible assets such as software and hardware systems. Our modification opens the original model's "black box" governing the transition of workers from unskilled to skilled and links this process, via mismatch, directly to firms' IT-related investment decisions.

For tractability, we model capital investment as a function of skill mismatch $M^{\tau,u}$, abstracting from other determinants, so that investment occurs only when mismatch is present: $I_t = Investment_t \times \mathbb{1}_{\{M^{\tau,u} \neq 0\}}$. This specification isolates the role of mismatch in shaping investment behavior. We further assume that investment follows standard q-theory under price taking and convex adjustment costs (Lucas, 1967; Treadway, 1969; Hayashi, 1982). The firm's production function is $Y_t = \gamma^{t-\tau} f(U_t, S_t) + \gamma^{t-\tau} g(K_t)$, where K_t includes both tangible and intangible capital; additive separability ensures that U_t and S_t do not depend on K_t. Capital evolves according to $\dot{K}_t = I_t - \delta K_t$, where δ is the depreciation rate. Firms face convex adjustment costs of the form $\Phi(I_t, K_{t-1}) = \frac{\phi}{2}\left(\frac{I_t}{K_{t-1}} - \delta\right)^2 K_{t-1}$.

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 lifecycle.

In equilibrium, our model predicts that the mismatch between firms and unskilled workers, $M^{\tau,u}$, follows a U-shaped pattern with respect to technology vintage τ. The mechanism operates through two channels. First, the production complementarity between skilled and unskilled workers implies that $\frac{\partial M^{\tau,u}}{\partial S_t} < 0$; that is, skill mismatch among unskilled workers decreases as the firm employs more skilled workers. Second, the equilibrium distribution of skilled labor, s_τ, is single-peaked across technology vintages. 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 and greater numbers of unskilled workers accumulate experience and become skilled, firms can assign tasks more efficiently, reducing mismatch for new entrants. 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 find it harder to reorganize tasks and skill requirements for new employees, causing skill mismatch to rise once again.

Prediction 2: IT-complementary skill requirements exceed worker capabilities

Although our framework can be broadly applied across a range of technologies, in this paper, we focus on the diffusion of information technologies. For the model to be consistent with observed patterns of IT adoption, we should observe firm skill requirements exceed

worker capabilities for IT-complementary skills, as these are the skills that are critical for IT diffusion. That is, we should observe positive skill mismatch, $M_c^{\tau,u} > 0$, for unskilled workers across IT-complementary skills $c \in \{1, 2, \ldots, N\}$. Conversely, we should observe negative mismatch for IT-neutral skills, since skill weights sum to one. Thus we predict that $M_n^{\tau,u} < 0$ for $n \neq c$, reflecting the relative over-supply of skills that are less relevant for explaining information technology diffusion.

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

Our assumption that the marginal productivity of unskilled workers is decreasing and concave in the degree of skill mismatch implies that productivity losses rise disproportionately with larger mismatches. Because of this concavity, firms have stronger incentives to close the largest, rather than the smallest, individual skill gaps $M_i^{\tau,j}$ between firm requirements and worker capabilities. Following Predictions 1 and 2, this reasoning implies that as s_τ increases, firms will prioritize adjustments in tasks and skill requirements for IT-complementary skills—those that initially 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. Consequently, the model predicts a U-shaped relationship between IT-complementary skill mismatch and technology vintage.

Prediction 4: Tobin's Q and firm investment decrease with technology vintage when skill mismatch $M^{\tau,u}$ is non-zero

Our model predicts that the firm's equilibrium marginal product of capital, and thus Tobin's q, are both decreasing in technology vintage τ. As τ increases, technologies become older and less productive than frontier technologies, as captured by the productivity term $\gamma^{t-\tau}$. Consequently, the expected marginal return on new investment declines. In the presence of convex adjustment costs $\Phi(I_t, K_t) = \frac{\phi}{2}\left(\frac{I_t}{K_{t-1}} - \delta\right)^2 K_t$ and non-zero skill mismatch $M^{\tau,u}$, firms optimally reduce investment I_t as q_t falls to avoid incurring high adjustment costs when

returns to new capital are low.[2]

Intuitively, firms using newer technologies (low τ) face high q and strong incentives to invest in both tangible and intangible capital in order to reduce skill mismatch and help new employees acquire IT-complementary skills. These incentives change with technology vintage, however. As technologies mature and τ increases, declining productivity and lower expected returns reduce the firm's incentives to make such investments. Firms gradually invest less in capital that facilitates worker adaptation, and workers have less incentive to join these firms. Eventually, these dynamics culminate in technologies becoming obsolete.

Implications for Firm Productivity and Performance

While our model does not deliver explicit predictions for productivity or profitability, it has sharp implications for interpreting empirical patterns in *measured* TFP and performance. The model implies that capital investment declines with technology vintage whenever skill mismatch remains positive. Because the accumulation of intangible assets—including training, organizational capital, and other complementary capabilities—is largely unobserved, and because these assets take time to translate into observable output, measured TFP will systematically understate true productivity early in the technology lifecycle. As mismatches narrow and intangible capital is gradually embodied in production, measured TFP should rise with technology vintage. Thus, the investment dynamics induced by skill mismatch and

[2]The related firm value at time t can be written as

$$V_t = \gamma^{t-1-\tau} f(U_{t-1}, S_{t-1}) + \gamma^{t-1-\tau} g(K_{t-1}) - I_t - \frac{\phi}{2}\left(\frac{I_t}{K_{t-1}} - \delta\right)^2 K_{t-1} + E[M_{t+1}V_{t+1}]$$

So, a Lagrangian is

$$\mathcal{L}_t = \gamma^{t-1-\tau} f(U_{t-1}, S_{t-1}) + \gamma^{t-1-\tau} g(K_{t-1}) - I_t - \frac{\phi}{2}\left(\frac{I_t}{K_{t-1}} - \delta\right)^2 K_{t-1} + E[M_{t+1}V_{t+1}] - q_t[K_t - (1-\delta)K_{t-1} - I_t]$$

FOC wrt I_t

$$q_t = 1 + \phi\left(\frac{I_t}{K_{t-1}} - \delta\right)$$

FOC wrt K_t

$$q_t = E_t\left[M_{t+1}\left\{\gamma^{t-\tau}\cdot\frac{\partial g(K_t)}{\partial K_t} + \frac{\phi}{2}\left(\frac{I_{t+1}}{K_t}\right)^2 - \frac{\phi\cdot\delta^2}{2} + q_{t+1}\cdot(1-\delta)\right\}\right]$$

the changing level of skill mismatch over the technology lifecycle provide a micro-founded interpretation of the J-curve patterns in productivity documented by Brynjolfsson et al. (2021).

3 Empirical Approach

3.1 Challenges

Our conceptual framework is helpful for motivating a theoretical notion of skill mismatch and seeing it as an endogenous equilibrium market outcome that is jointly shaped by firms' technology choices. However, empirically characterizing skill mismatch and its relationship with firm technologies remains 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 lifecycle.

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.

3.2 Empirical Methodology

We devise an empirical methodology that helps us overcome the aforementioned challenges and empirically characterize skill mismatch over the technology lifecycle. 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.[3] 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.

[3]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.

4 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 3 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 3]

5 Empirical Findings

5.1 Skill mismatch and technology vintage

Table 2 and Figure 3 present our first main finding: skill mismatch exhibits a pronounced U-shaped pattern over the technology life cycle—high when firms adopt frontier systems, declining as firms and workers adjust, and rising again as systems age and become obsolete. As shown in the regression estimates, the coefficients on the technology-vintage bins are robust to industry and firm fixed effects, indicating that this relationship holds both within firms over time and across firms in the cross section.

The observed pattern is consistent with the first prediction of our model in Section 2. As technology vintage increases and the share of skilled workers rises, average mismatch

among unskilled workers declines because they increasingly work alongside more skilled, complementary peers. Once the share of skilled workers reaches its peak, mismatch attains its minimum. As the share of skilled workers subsequently falls with further aging of the technology, mismatch increases again.

These findings rule out alternative hypotheses regarding the relationship between technology vintage and mismatch. For example, they are inconsistent with frictionless models featuring instantaneous and costless learning, under which no systematic variation in mismatch across vintages should emerge. They are also difficult to reconcile with frameworks that treat worker skills as exogenous to firms' technology choices. Instead, the results indicate that firms operating IT systems of differing vintages face systematically different levels of skill mismatch that evolve endogenously with technology adoption and adjustment dynamics.

Additional patterns in Figure 3 are also informative. The substantial duration represented by each technology bin indicates that technology adoption in our sample is slow, consistent with the motivating observation that diffusion is gradual relative to technological progress. The decline in mismatch across newer vintages follows the diminishing-marginal-adjustment mechanism in our model: as the stock of skilled workers expands, firms close their largest internal skill gaps first, yielding a decreasing rate of mismatch reduction over time. The figure further shows that mismatch reaches its lowest level in bin 9, corresponding to roughly 35 years in our data. This suggests that it can take up to three and a half decades for firms to reach their minimum mismatch level as they search for and integrate workers with the appropriate skills.

5.2 Heterogeneity across firms

We examine our measures of skill mismatch across industries and along key firm characteristics. Patterns across industries are presented in Figure 4. The results show that substantial mismatch arises precisely in sectors with well-documented heterogeneity in IT adoption. For example, healthcare and manufacturing exhibit both high average mismatch and pronounced

U-shapes over technology vintages. By contrast, industries such as retail and consumer nondurables display low and relatively flat mismatch profiles, consistent with the view that skill mismatch poses a limited constraint on technology diffusion in these sectors.

Figures 6, 7, and 8 show that average skill mismatch is systematically higher in younger, smaller, and more financially constrained firms—those with the fewest resources to smooth adjustment following IT adoption. The pattern is economically meaningful and statistically robust. It survives an extensive set of controls, including industry and firm fixed effects, calendar-year shocks, and alternative functional forms for technology vintage. The same qualitative patterns emerge when we use alternative mismatch measures (e.g., those based on occupational matching rather than firm–worker pairing). Together, these findings reinforce the interpretation of skill mismatch as a meaningful diffusion friction that interacts with other documented barriers to technology adoption.

In our framework, financing frictions constrain complementary investment in physical and organizational capital that makes worker learning productive—such as formal training programs, mentoring structures, process and data tooling, and managerial bandwidth. When these complements are scarce, learning-by-doing is less effective, and firms face greater difficulty reallocating tasks and adjusting skill requirements to align with the capabilities of new workers.

5.3 Mismatch by IT-complementary skills

A second key prediction of our model is that significant skill gaps should arise precisely for skills that are complementary to information technologies. In particular, firms' required skill levels should exceed workers' capabilities for IT-complementary attributes. As shown in Table A2, we classify the 25 skill groups in our analysis into three categories following prior work on IT and productivity (Bartel et al., 2007): IT-complementary technical skills, IT-complementary nontechnical skills, and IT-neutral (Peripheral) skills.

Consistent with this prediction, Figures 5 and 9 show that mismatch for IT-complementary

skills—and especially for IT-complementary nontechnical skills—is substantially larger than 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. Notably, 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" 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 for each of the three skill groups. Figure 10 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 the IT-complementary nontechnical skills.

5.4 Skill Mismatch, Investment, and Productivity

We next examine the model's predictions regarding firms' responses to skill mismatch and the implications of these responses for measured productivity and performance. As discussed in Section 2, the model predicts that both Tobin's Q and investment should decline as technology vintage increases. We proxy for Q using firms' market-to-book ratios. Figure **??** and Table 5 show that Q indeed decreases with technology vintage: as technologies age and become relatively less productive, the marginal return to investment falls accordingly.

Consistent with this pattern, Figure 12 documents declining investment rates as technology vintage rises. Notably, intangible investment (measured following Peters and Taylor (2017)) exceeds investment in tangible assets throughout the vintage distribution. These intangible expenditures—training, organizational redesign, changes in work practices, and related forms of complementary capital—are precisely the inputs emphasized in prior work on IT upgrading. As the model predicts, investment rates decline with technology vintage whenever skill mismatch is nonzero, a condition we established in earlier sections.

To assess the implications for total factor productivity (TFP) and profitability, we analyze standard measures of TFP and realized profitability across technology vintages. Table 6 (Panel A) and Figure 13 show that measured TFP increases with technology vintage. Coupled with the market-to-book evidence indicating that newer vintages are inherently more productive, this pattern suggests that the gap between measured and true TFP is largest for younger technologies—precisely when intangible investment is highest and skill mismatch is greatest. This follows from the well-known delay with which intangible investments translate into measured productivity gains (Brynjolfsson et al., 2021). In our setting, intangible investment peaks at the point of greatest mismatch, implying that measured TFP substantially understates true productivity during the early stages of adoption.

Taken together, these findings support the view that skill mismatch helps explain both the heavy intangible investment undertaken when adopting new technologies and the slug-

gish measured productivity associated with early IT diffusion. The results align closely with the "productivity J-curve" framework (Brynjolfsson et al., 2021), in which investments in intangible complements initially depress measured productivity before yielding long-run gains.

6 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.

6.1 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 7, 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 7]

6.2 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.

6.3 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.

6.4 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.

6.5 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.

7 Conclusion

This paper provides theory and evidence showing that skill gaps in the labor market constitute a central friction in the diffusion of new technologies. Focusing on information-technology adoption, 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.

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.

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.

Brynjolfsson, E., Rock, D., and Syverson, C. (2021). The productivity j-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1).

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.

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

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.

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.

Hayashi, F. (1982). Tobin's marginal q and average q: A neoclassical interpretation. *Econometrica*, 50(1):213–224.

Hershbein, B. and Kahn, L. (2018). Do recessions accelerate routine-biased technological change? evidence from vacancy postings. *American Economic Review*, 108(7):1732–1772.

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.

Lucas, R. E. (1967). Adjustment costs and the theory of supply. *Journal of Political Economy*, 75(4):321–334.

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.

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.

Treadway, A. B. (1969). On rational entrepreneurial behavior and the demand for investment. *Review of Economic Studies*, 36(2):227–239.

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.

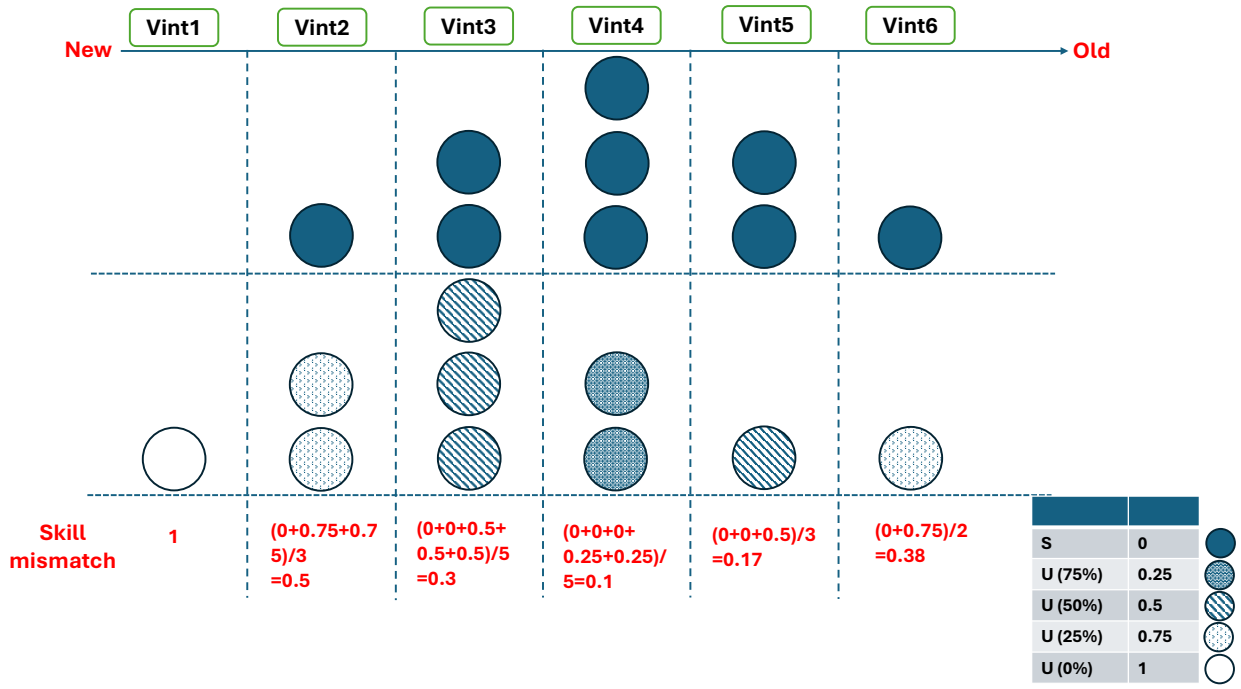


Figure 1: Model Overview

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 vintage 4. At that time, fewer unskilled workers join the firm, subsequently increasing the mismatch that new workers exhibit each period, until the technology become obsolete.

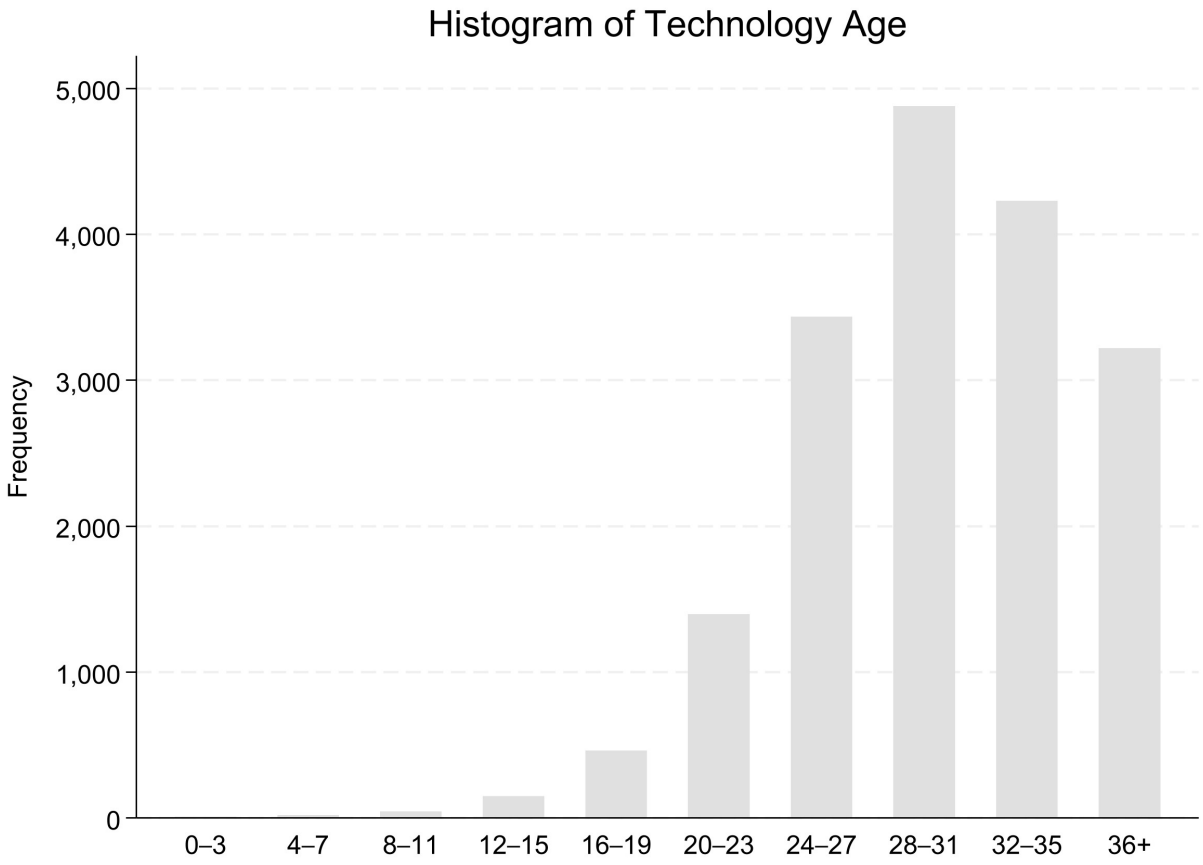


Figure 2: Skill Mismatch by Firm Technology Vintage
This figure presents a histogram of technology vintage in our sample. Each data point corresponds to a firm-year observation. The x-axis depicts the average technology vintage, computed as the average information technology age in the firm's job postings as of the year of observation.

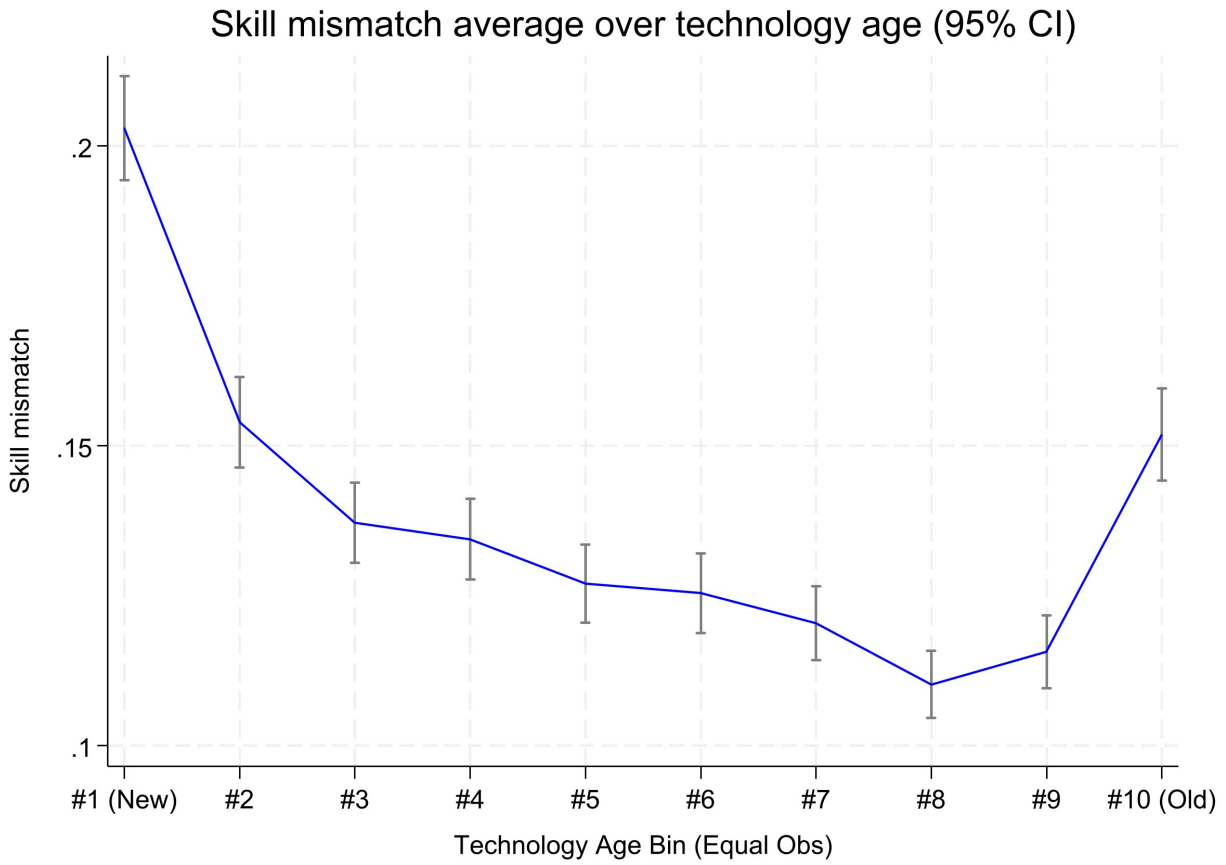


Figure 3: Skill Mismatch Over The Technology Lifecycle
This figure presents the empirical relationship between firm technology vintage and average skill mismatch between the firm and its new employees. Skill mismatch is defined as the Euclidean distance between the firm's skill requirements and new employees' 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. The figure demonstrates that firms using newer technologies tend to experience larger skill mismatch, consistent with the rapid evolution of skill requirements in technology-intensive industries.

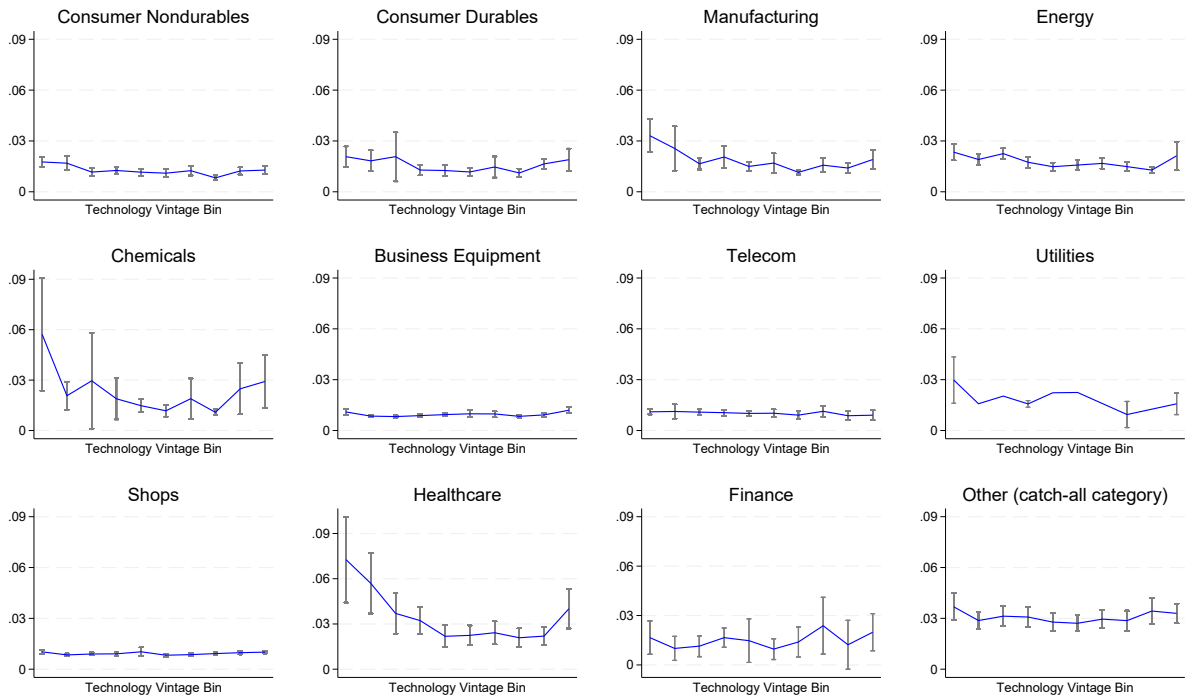


Figure 4: Distribution of Skill Mismatch by Industries

This figure presents the distribution of skill mismatch across technology vintage for each of 12 Fama-French-defined industries. Variable definitions are the same as in Figure 3.

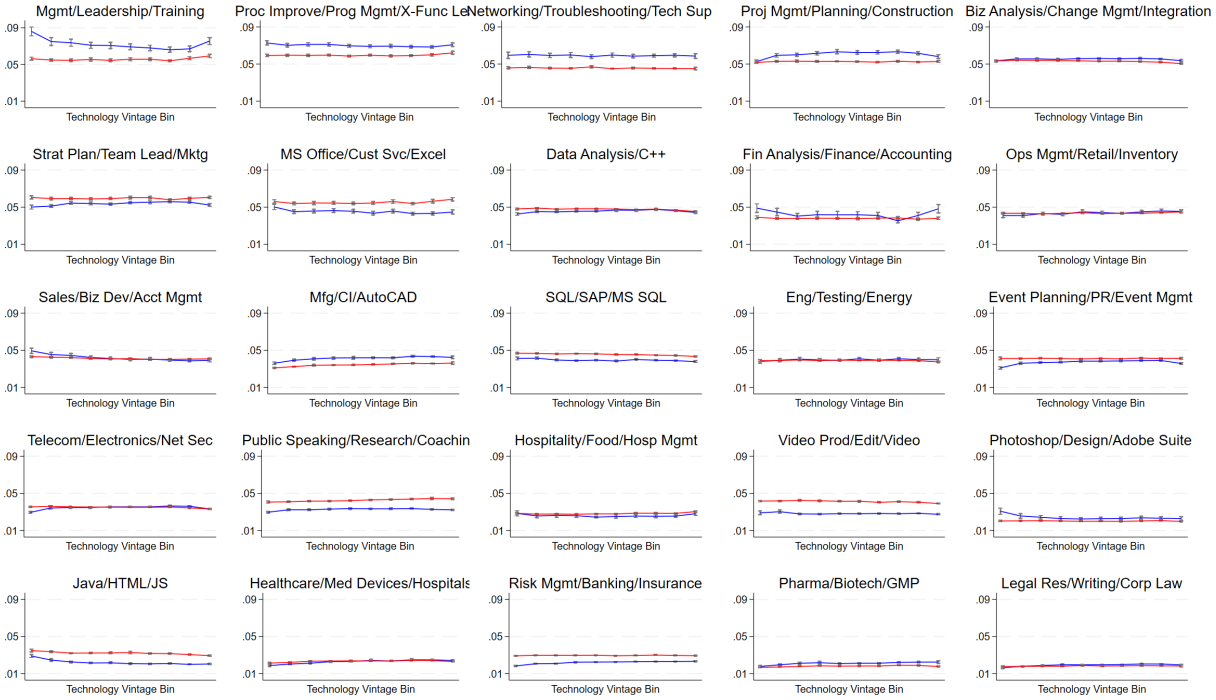


Figure 5: Distribution of Skill Mismatch by Skill Category: Demand vs. Supply
This figure presents the distribution of skill demand weights for job requirements (blue) and skill supply weights for workers (red) across each of the 25 skill categories used in our analysis (defined by Revelio Labs). The Y-axis depicts the average weight score for each skill group, the X-axis depicts technology vintage bins as defined in Figure 3.

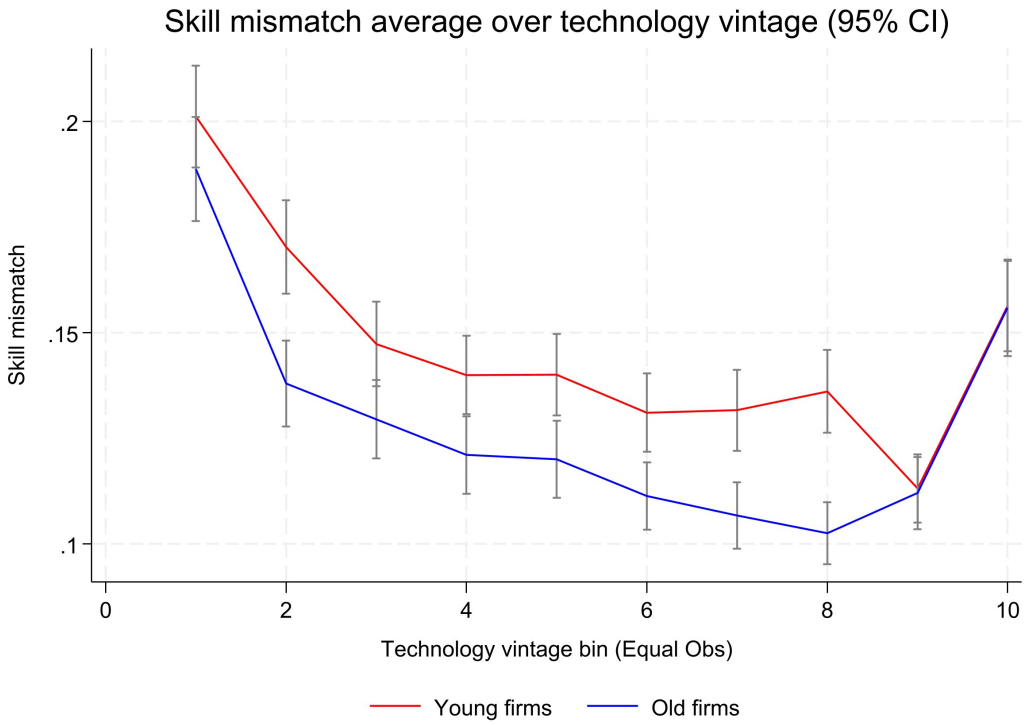


Figure 6: Young vs. Old: Skill Mismatch Over The Technology Lifecycle
This figure shows the relationship between firm technology vintage and average skill mismatch for young vs old firms. Technology vintage is measured as the vintage of the primary technologies used by the firm, with newer technologies corresponding to lower technology vintage values. The figure demonstrates that firms using newer technologies tend to experience larger skill mismatch, consistent with the rapid evolution of skill requirements in technology-intensive industries.

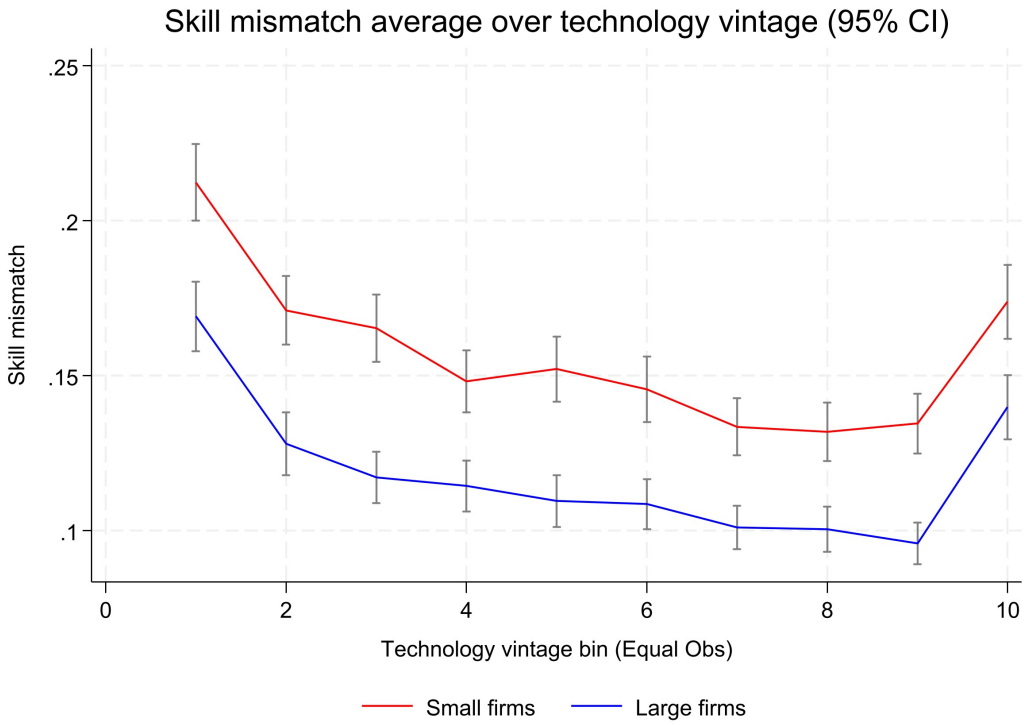


Figure 7: Small vs. Large: Skill Mismatch Over The Technology Lifecycle
This figure shows the relationship between firm technology vintage and average skill mismatch for small vs large firms. Technology vintage is measured as the vintage of the primary technologies used by the firm, with newer technologies corresponding to lower technology vintage values. The figure demonstrates that firms using newer technologies tend to experience larger skill mismatch, consistent with the rapid evolution of skill requirements in technology-intensive industries.

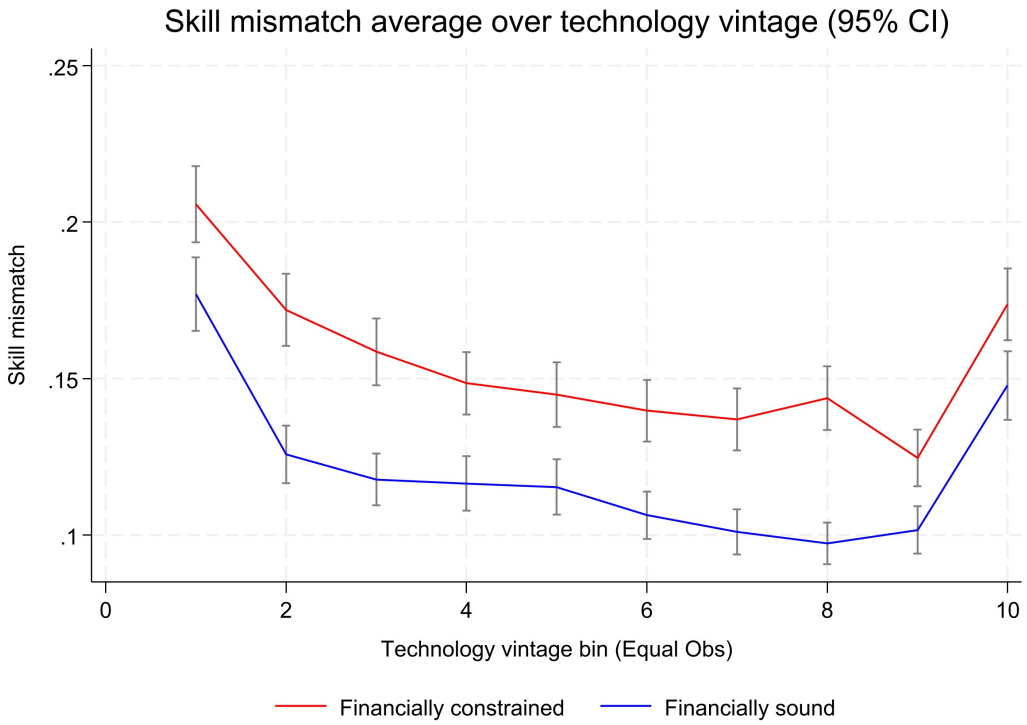


Figure 8: Financially constrained vs. sound: Skill Mismatch Over The Technology Lifecycle
This figure shows the relationship between firm technology vintage and average skill mismatch for financially constrained vs sound firms. Technology vintage is measured as the vintage of the primary technologies used by the firm, with newer technologies corresponding to lower technology vintage values. The figure demonstrates that firms using newer technologies tend to experience larger skill mismatch, consistent with the rapid evolution of skill requirements in technology-intensive industries.

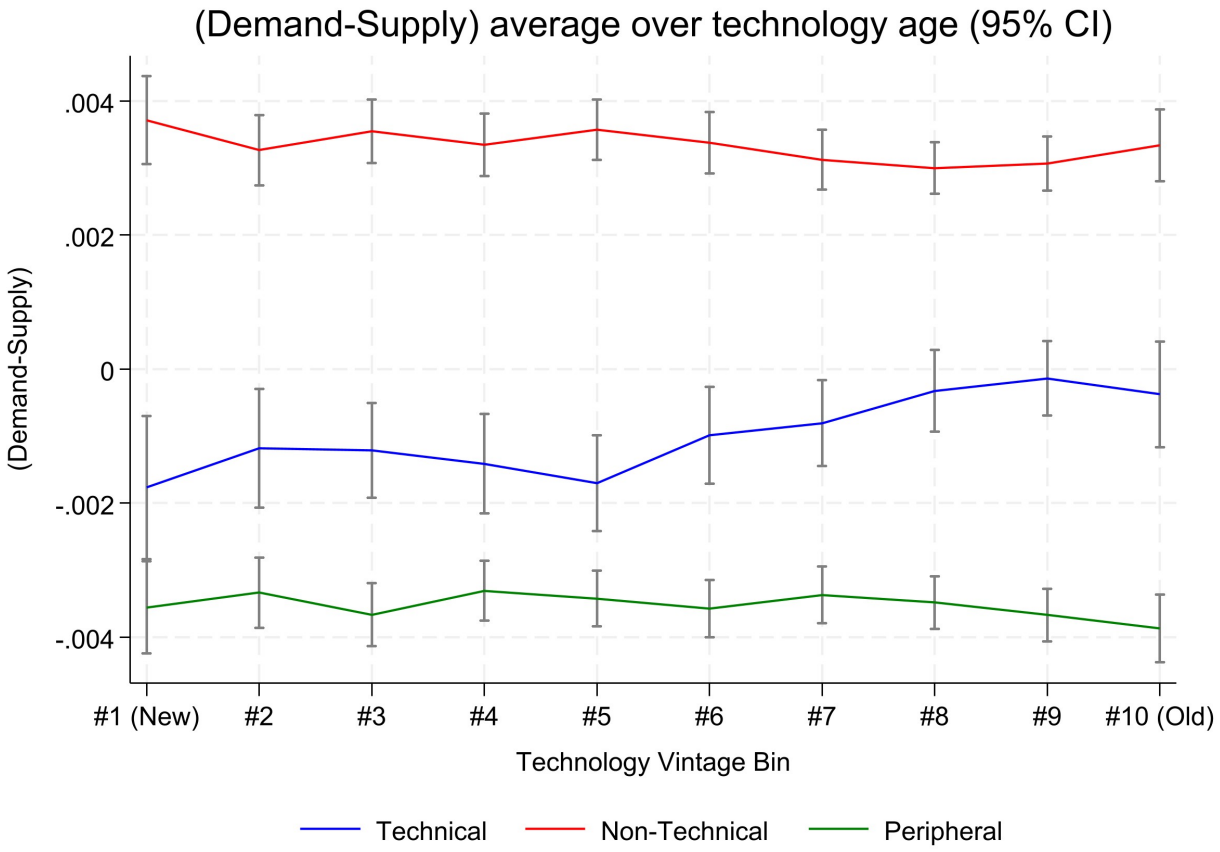


Figure 9: XXX

This figure presents differences between firm skill requirements and worker skill attributes for three groups of skills: IT-complementary technical, IT-complementary nontechnical, and IT-neutral skills. The Y-axis depicts the average score for each skill group, the X-axis depicts technology vintage bins as defined in Figure 3.

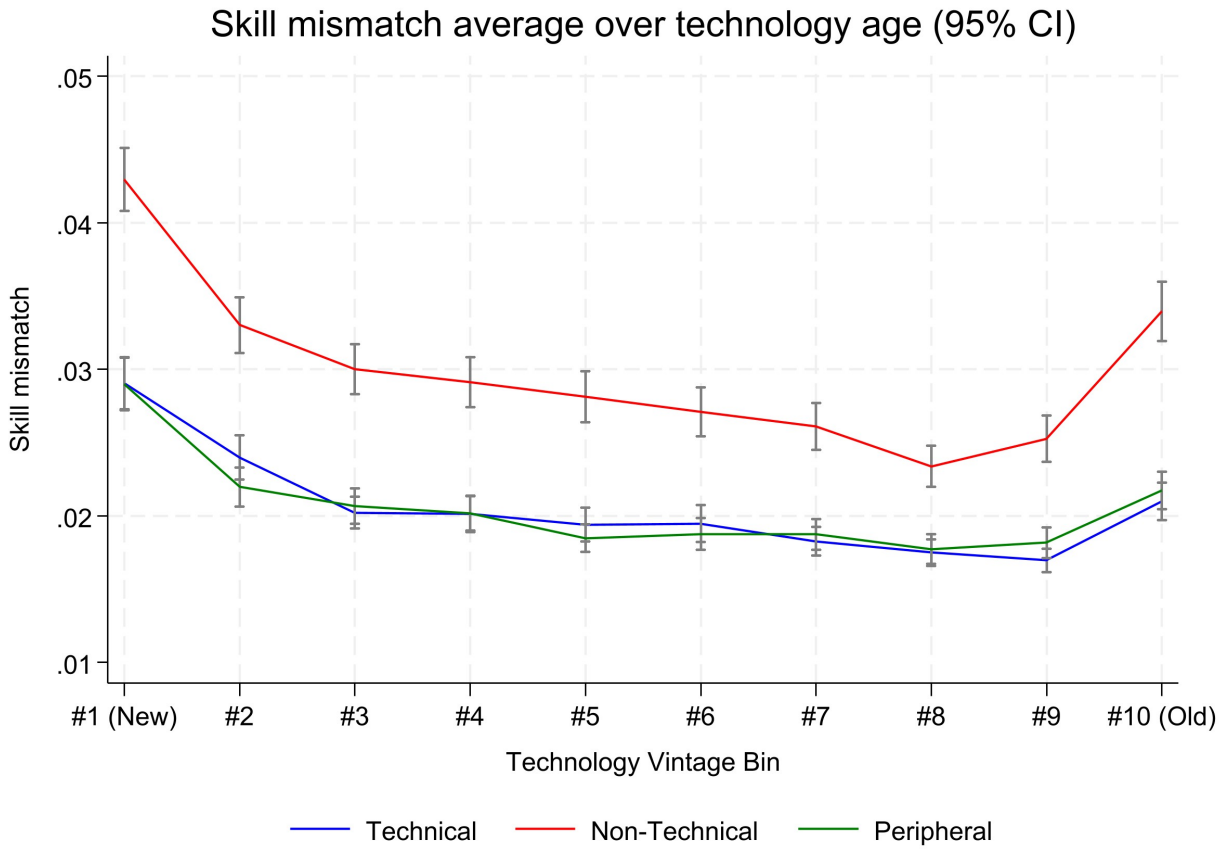


Figure 10: XXX

This figure presents skill mismatch (Euclidean distance measures) between firm skill requirements and worker skill attributes for three groups of skills: IT-complementary technical, IT-complementary nontechnical, and IT-neutral skills. The Y-axis depicts the average mismatch for each skill group, the X-axis depicts technology vintage bins as defined in Figure 3.

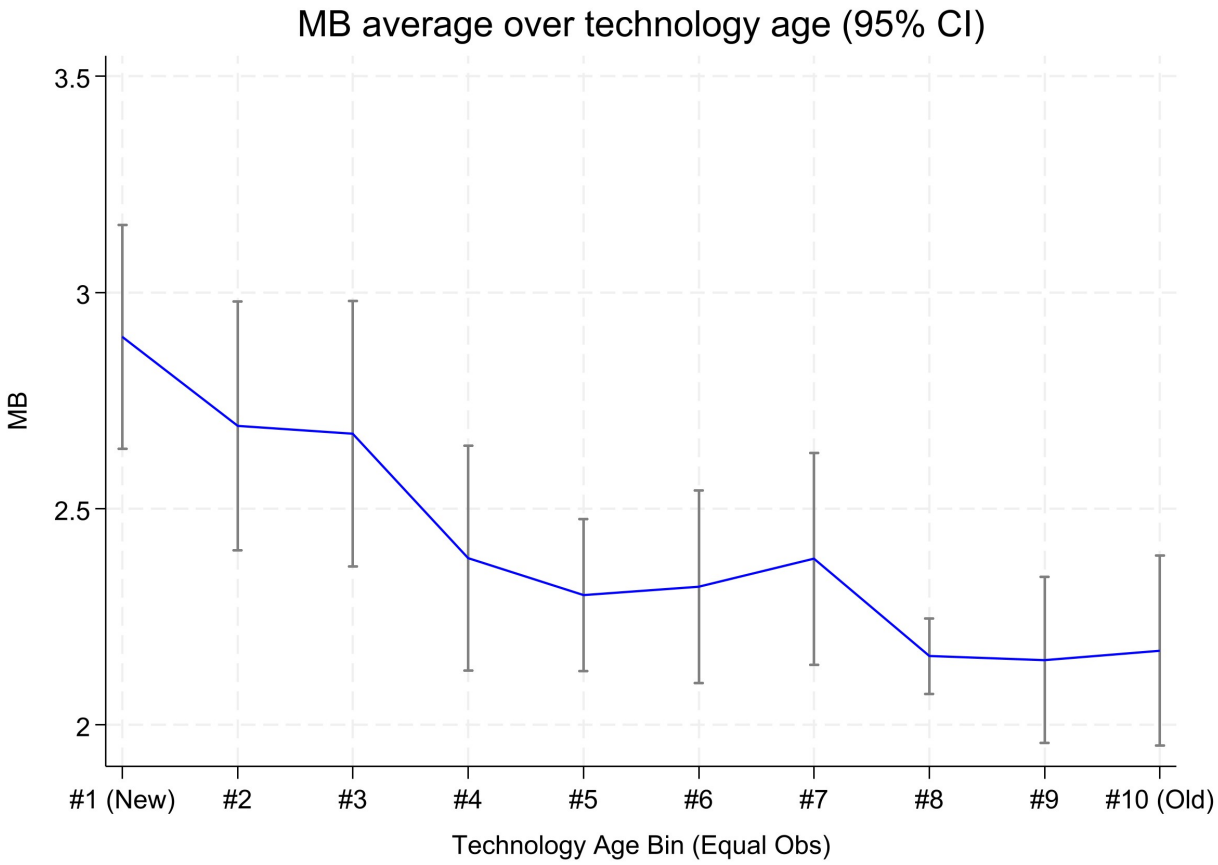


Figure 11: Market-to-Book Ratios over The Technology Lifecycle

This figure depicts firm market-to-book ratios across technology vintages. The Y-axis depicts the average market-to-book ratio for sample firms; the X-axis depicts technology vintage bins as defined in Figure 3.

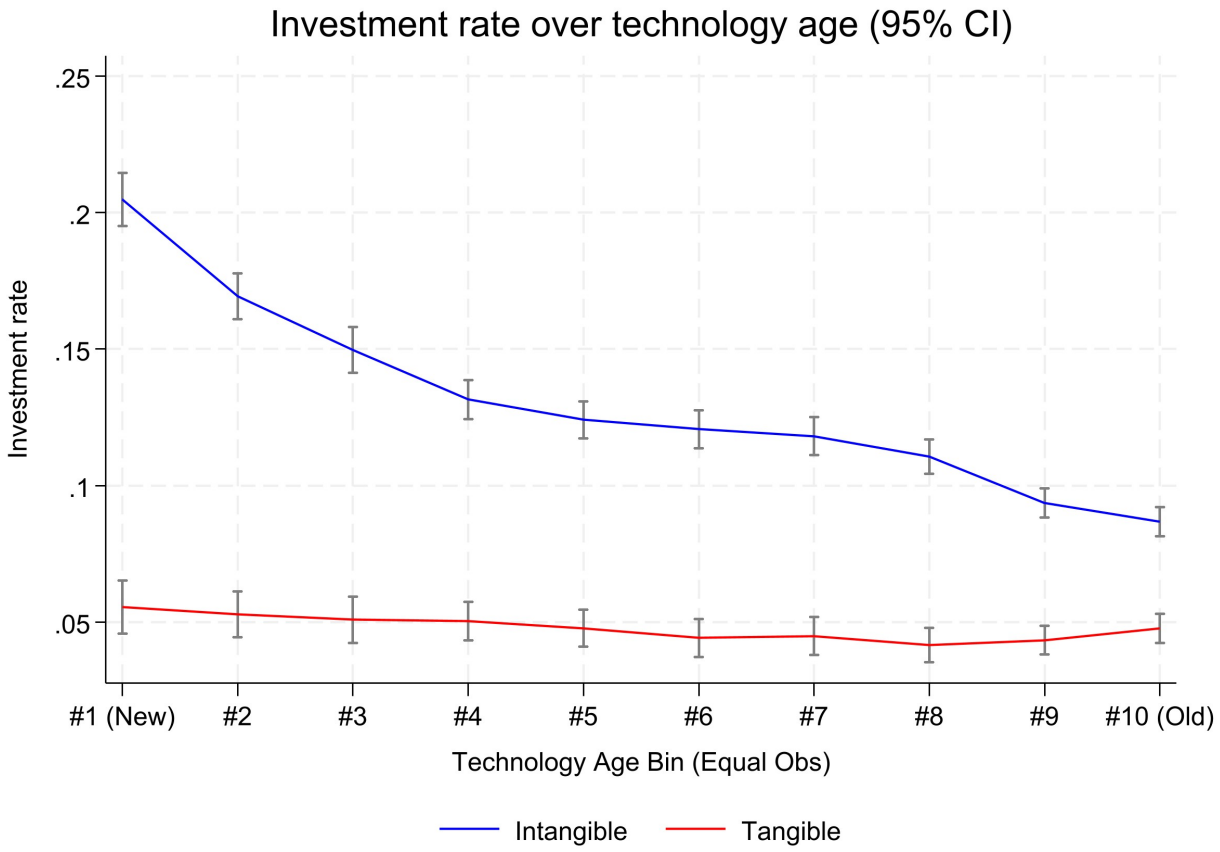


Figure 12: Investment over The Technology Lifecycle

This figure depicts firm investment across technology vintages. The Y-axis depicts the average investment rate for sample firms; the X-axis depicts technology vintage bins as defined in Figure 3. The red line corresponds to investment in tangible assets, the blue line corresponds to investment in intangible assets. Please see appendix for variable construction.

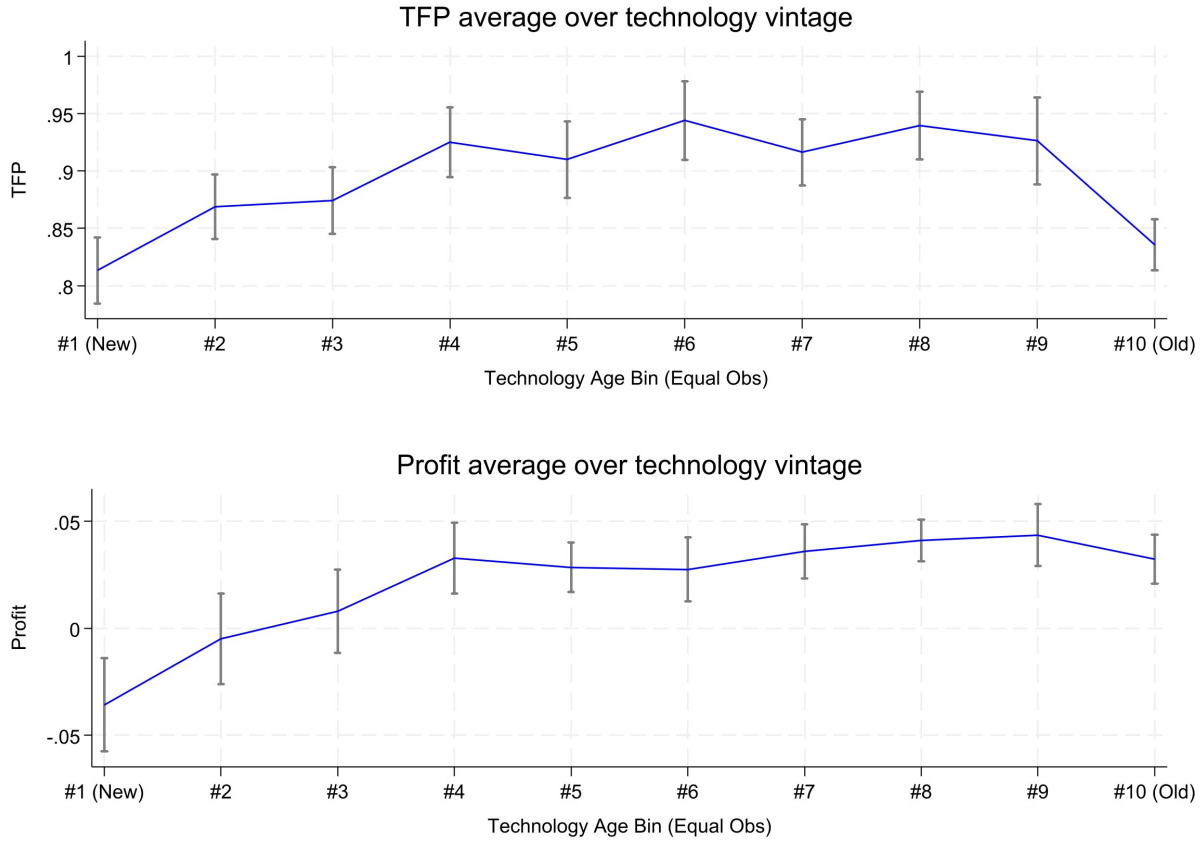


Figure 13: Relationship Between Skill Mismatch and Firm Performance

This figure depicts total factor productivity (top panel) and operating performance (bottom panel) across technology vintages. The X-axis depicts technology vintage, as defined in Figure 3. The Y-axis depicts average TFP and average profitability for each firm-year observation in the sample. Variable definitions for TFP and operating performance are provided in the appendix.

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 vintage	17,998	30.85	6.195	26.94	30.75	34.64
TFP	14,039	-0.242	0.507	-0.486	-0.246	0.0204
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
Intangible investment rate	19,011	0.104	0.112	0.0318	0.0743	0.132
Tangible investment rate	18,990	0.0439	0.0479	0.0156	0.0292	0.0527
Panel B: Compustat Population						
TFP	33,953	-0.317	0.576	-0.580	-0.296	-0.0113
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
Intangible investment rate	96,436	0.157	0.274	0.0187	0.0747	0.174
Tangible investment rate	95,830	0.0598	0.127	0.00622	0.0233	0.0576

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 our measure of firm-level skill mismatch. Technology vintage is defined as per Figure 3. Panel A shows statistics for firms in our sample with both job posting and employee skill data. Panel B shows statistics for all Compustat firms during the same time period. All dollar amounts are inflation-adjusted to 2010 dollars. All variables are winsorized at the 1st and 99th percentiles.

Table 2: Main Result: Skill Mismatch vs Technology Vintage

	(1)	(2)	(3)	(4)	(5)	(6)
TechVint	-2.292***		-2.317***		-2.222***	
	(0.173)		(0.169)		(0.168)	
TechVint^2	3.464***		3.479***		3.342***	
	(0.271)		(0.262)		(0.260)	
TechVint2		-0.0445***		-0.0441***		-0.0413***
		(0.00650)		(0.00638)		(0.00616)
TechVint3		-0.0617***		-0.0621***		-0.0592***
		(0.00641)		(0.00629)		(0.00616)
TechVint4		-0.0608***		-0.0616***		-0.0604***
		(0.00654)		(0.00642)		(0.00625)
TechVint5		-0.0632***		-0.0652***		-0.0622***
		(0.00670)		(0.00659)		(0.00648)
TechVint6		-0.0625***		-0.0648***		-0.0632***
		(0.00700)		(0.00706)		(0.00692)
TechVint7		-0.0637***		-0.0654***		-0.0633***
		(0.00687)		(0.00687)		(0.00673)
TechVint8		-0.0713***		-0.0735***		-0.0682***
		(0.00707)		(0.00721)		(0.00713)
TechVint9		-0.0637***		-0.0660***		-0.0623***
		(0.00743)		(0.00766)		(0.00750)
TechVint10		-0.0298***		-0.0335***		-0.0307***
		(0.00812)		(0.00830)		(0.00815)
Observations	17,978	17,978	17,978	17,978	17,975	17,975
R-squared	0.045	0.040	0.062	0.056	0.115	0.110
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 investigates how skill mismatch vary across different technology vintage. TechVint is technology vintage, normalized by 100. Columns (1) and (4) show linear regression. Columns (2) and (5) show regression with linear and quadratic terms, revealing a U-shaped relationship. Columns (3) and (6) show piece-wise linear regression. The results consistently show that skill mismatch first decrease with technology vintage (as firms adopt newer technologies) but then increase again for the newest technologies, reflecting the particularly challenging skill requirements of cutting-edge innovations. Two-digit SIC industry fixed effects included. Standard errors are heteroskedasticity consistent and clustered by firm.

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

Table 3: Summary Statistics: Skill Mismatch

#	Skill group	Variable	N	Mean	Std Dev
#1	Mgmt / Leadership / Training	Demand	22,467	0.0766	0.0865
		Supply	22,467	0.0570	0.0374
		Mismatch	22,467	0.0400	0.0872
#2	Proc Improve / Prog Mgmt / X-Func Lead	Demand	22,467	0.0720	0.0440
		Supply	22,467	0.0602	0.0279
		Mismatch	22,467	0.0283	0.0447
#3	Networking / Troubleshooting / Tech Support	Demand	22,467	0.0583	0.0559
		Supply	22,467	0.0452	0.0204
		Mismatch	22,467	0.0263	0.0543
#4	Proj Mgmt / Planning / Construction	Demand	22,467	0.0581	0.0447
		Supply	22,467	0.0527	0.0210
		Mismatch	22,467	0.0255	0.0425
#5	Biz Analysis / Change Mgmt / Integration	Demand	22,467	0.0543	0.0216
		Supply	22,467	0.0524	0.0138
		Mismatch	22,467	0.0158	0.0193
#6	Strat Plan / Team Lead / Mktg	Demand	22,467	0.0524	0.0356
		Supply	22,467	0.0606	0.0314
		Mismatch	22,467	0.0248	0.0416
#7	MS Office / Cust Svc / Excel	Demand	22,467	0.0468	0.0476
		Supply	22,467	0.0565	0.0413
		Mismatch	22,467	0.0296	0.0560
#8	Data Analysis / C++	Demand	22,467	0.0442	0.0250
		Supply	22,467	0.0470	0.0132
		Mismatch	22,467	0.0157	0.0230
#9	Fin Analysis / Finance / Accounting	Demand	22,467	0.0441	0.0861
		Supply	22,467	0.0378	0.0279
		Mismatch	22,467	0.0299	0.0850
#10	Ops Mgmt / Retail / Inventory	Demand	22,467	0.0441	0.0427
		Supply	22,467	0.0436	0.0153
		Mismatch	22,467	0.0193	0.0398
#11	Sales / Biz Dev / Acct Mgmt	Demand	22,467	0.0434	0.0460
		Supply	22,467	0.0416	0.0199
		Mismatch	22,467	0.0211	0.0442
#12	Mfg / CI / AutoCAD	Demand	22,467	0.0407	0.0299
		Supply	22,467	0.0338	0.0163
		Mismatch	22,467	0.0170	0.0285
#13	SQL / SAP / MS SQL	Demand	22,467	0.0394	0.0245
		Supply	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	Demand	22,467	0.0388	0.0389
		Supply	22,467	0.0383	0.0169
		Mismatch	22,467	0.0181	0.0372
#15	Event Planning / PR / Event Mgmt	Demand	22,467	0.0351	0.0253
		Supply	22,467	0.0409	0.0183
		Mismatch	22,467	0.0178	0.0256
#16	Telecom / Electronics / Net Sec	Demand	22,467	0.0329	0.0196
		Supply	22,467	0.0345	0.0113
		Mismatch	22,467	0.0138	0.0169
#17	Public Speaking / Research / Coaching	Demand	22,467	0.0319	0.0191
		Supply	22,467	0.0426	0.0192
		Mismatch	22,467	0.0176	0.0228
#18	Hospitality / Food / Hosp Mgmt	Demand	22,467	0.0294	0.0596
		Supply	22,467	0.0287	0.0157
		Mismatch	22,467	0.0174	0.0560
#19	Video Prod / Edit / Video	Demand	22,467	0.0277	0.0229
		Supply	22,467	0.0412	0.0205
		Mismatch	22,467	0.0197	0.0264
#20	Photoshop / Design / Adobe Suite	Demand	22,467	0.0244	0.0476
		Supply	22,467	0.0203	0.0148
		Mismatch	22,467	0.0147	0.0474
#21	Java / HTML / JS	Demand	22,467	0.0227	0.0217
		Supply	22,467	0.0317	0.0212
		Mismatch	22,467	0.0165	0.0266
#22	Healthcare / Med Devices / Hospitals	Demand	22,467	0.0225	0.0200
		Supply	22,467	0.0232	0.0136
		Mismatch	22,467	0.0114	0.0194
#23	Risk Mgmt / Banking / Insurance	Demand	22,467	0.0211	0.0115
		Supply	22,467	0.0295	0.0119
		Mismatch	22,467	0.0132	0.0128
#24	Pharma / Biotech / GMP	Demand	22,467	0.0206	0.0224
		Supply	22,467	0.0178	0.0121
		Mismatch	22,467	0.00963	0.0221
#25	Legal Res / Writing / Corp Law	Demand	22,467	0.0185	0.0121
		Supply	22,467	0.0179	0.0166
		Mismatch	22,467	0.00956	0.0180

Table 4: MB ratio vs. vs 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)
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 investigates how skill mismatch varies across different technology vintage. TechVint is technology vintage, normalized by 100. Columns (1) and (4) show linear regression. Columns (2) and (5) show regression with linear and quadratic terms, revealing a U-shaped relationship. Columns (3) and (6) show piece-wise linear regression. The results consistently show that skill mismatch first decreases with technology vintage (as firms adopt newer technologies) but then increase again for the newest technologies, reflecting the particularly challenging skill requirements of cutting-edge innovations. Two-digit SIC industry fixed effects included. Standard errors are heteroskedasticity consistent and clustered by firm.

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

Table 5: Firm Investment vs Technology Vintages

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Intangible Investment (Peters and Taylor, 2017)						
TechVint	-0.404***		-0.275***		-0.182***	
	(0.0370)		(0.0286)		(0.0263)	
TechVint2		-0.0274***		-0.0246***		-0.0195***
		(0.00764)		(0.00646)		(0.00604)
TechVint3		-0.0486***		-0.0378***		-0.0300***
		(0.00801)		(0.00664)		(0.00611)
TechVint4		-0.0572***		-0.0464***		-0.0360***
		(0.00845)		(0.00694)		(0.00625)
TechVint5		-0.0642***		-0.0501***		-0.0380***
		(0.00853)		(0.00682)		(0.00617)
TechVint6		-0.0691***		-0.0526***		-0.0434***
		(0.00874)		(0.00689)		(0.00634)
TechVint7		-0.0715***		-0.0551***		-0.0411***
		(0.00895)		(0.00711)		(0.00641)
TechVint8		-0.0763***		-0.0580***		-0.0431***
		(0.00902)		(0.00703)		(0.00641)
TechVint9		-0.0851***		-0.0622***		-0.0443***
		(0.00921)		(0.00717)		(0.00662)
TechVint10		-0.0960***		-0.0660***		-0.0423***
		(0.00906)		(0.00721)		(0.00659)
Observations	17,567	17,567	17,567	17,567	17,564	17,564
R-squared	0.052	0.058	0.271	0.275	0.425	0.429
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)
Panel B: Tangible Investment						
TechVint	-0.00182		-0.0167		-0.0233**	
	(0.0111)		(0.0108)		(0.00990)	
TechVint2		-0.00311		-0.00321		-0.00325
		(0.00225)		(0.00210)		(0.00203)
TechVint3		-0.00525**		-0.00562**		-0.00494**
		(0.00247)		(0.00222)		(0.00216)
TechVint4		-0.00553**		-0.00559**		-0.00582***
		(0.00246)		(0.00229)		(0.00221)
TechVint5		-0.00523**		-0.00645***		-0.00679***
		(0.00246)		(0.00222)		(0.00215)
TechVint6		-0.00648***		-0.00708***		-0.00824***
		(0.00242)		(0.00221)		(0.00213)

Continued on next page

Table 5 – continued from previous page

	(1)	(2)	(3)	(4)	(5)	(6)
TechVint7		-0.00516**		-0.00545**		-0.00735***
		(0.00248)		(0.00232)		(0.00220)
TechVint8		-0.00591**		-0.00622***		-0.00717***
		(0.00242)		(0.00226)		(0.00216)
TechVint9		-0.00418		-0.00474**		-0.00641***
		(0.00255)		(0.00240)		(0.00228)
TechVint10		6.21e-05		-0.00410		-0.00539**
		(0.00274)		(0.00264)		(0.00242)
Observations	17,549	17,549	17,549	17,549	17,546	17,546
R-squared	0.033	0.035	0.172	0.173	0.270	0.271
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 investigates how firm investment varies across technology vintage. TechVint is technology vintage, normalized by 100. Columns (1) and (3) show intangible investment normalized by lagged total capital (intangible + tangible) as defined by Peters and Taylor (2017). Columns (2) and (4) show tangible investment (CAPX) normalized by lagged total capital. The results reveal negative relationship between technology vintage and firm investment. Two-digit SIC industry fixed effects included. Standard errors are heteroskedasticity consistent and clustered by firm.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: Performance vs Technology Vintage

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: TFP (Imrohoroglu and Tuzel, 2014)						
TechVint	2.918***		3.001***		2.515***	
	(0.722)		(0.705)		(0.682)	
TechVint^2	-4.263***		-4.554***		-3.715***	
	(1.042)		(1.030)		(0.988)	
TechVint2		0.0627**		0.0612**		0.0564**
		(0.0258)		(0.0254)		(0.0254)
TechVint3		0.0715**		0.0696**		0.0703***
		(0.0282)		(0.0279)		(0.0271)
TechVint4		0.116***		0.113***		0.110***
		(0.0289)		(0.0281)		(0.0270)
TechVint5		0.0919***		0.0845***		0.0800***
		(0.0308)		(0.0296)		(0.0283)
TechVint6		0.127***		0.122***		0.112***
		(0.0325)		(0.0312)		(0.0297)
TechVint7		0.118***		0.114***		0.104***
		(0.0324)		(0.0316)		(0.0303)
TechVint8		0.122***		0.111***		0.101***
		(0.0338)		(0.0328)		(0.0311)
TechVint9		0.0961***		0.0770**		0.0741**
		(0.0352)		(0.0345)		(0.0325)
TechVint10		0.0259		0.00246		0.0164
		(0.0355)		(0.0348)		(0.0320)
Observations	13,104	13,104	13,103	13,103	13,101	13,101
R-squared	0.007	0.009	0.038	0.041	0.138	0.141
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)
Panel B: Profitability						
TechVint	0.639***		0.714***		0.641***	
	(0.163)		(0.153)		(0.146)	
TechVint^2	-0.785***		-0.928***		-0.881***	
	(0.236)		(0.224)		(0.210)	
TechVint2		0.0180***		0.0179***		0.0151***
		(0.00546)		(0.00521)		(0.00502)
TechVint3		0.0222***		0.0214***		0.0187***
		(0.00578)		(0.00553)		(0.00531)
TechVint4		0.0296***		0.0287***		0.0242***
		(0.00604)		(0.00579)		(0.00550)

Continued on next page

Table 6 – continued from previous page

	(1)	(2)	(3)	(4)	(5)	(6)
TechVint5		0.0295***		0.0288***		0.0237***
		(0.00622)		(0.00594)		(0.00564)
TechVint6		0.0307***		0.0300***		0.0260***
		(0.00630)		(0.00596)		(0.00583)
TechVint7		0.0358***		0.0360***		0.0282***
		(0.00647)		(0.00623)		(0.00592)
TechVint8		0.0336***		0.0341***		0.0259***
		(0.00659)		(0.00623)		(0.00605)
TechVint9		0.0342***		0.0335***		0.0250***
		(0.00687)		(0.00649)		(0.00613)
TechVint10		0.0323***		0.0274***		0.0171***
		(0.00699)		(0.00675)		(0.00623)
Observations	17,978	17,978	17,978	17,978	17,975	17,975
R-squared	0.014	0.014	0.055	0.056	0.144	0.145
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 investigates how skill mismatch varies across different technology vintage. TechVint is technology vintage, normalized by 100. Columns (1) and (4) show linear regression. Columns (2) and (5) show regression with linear and quadratic terms, revealing a U-shaped relationship. Columns (3) and (6) show piece-wise linear regression. The results consistently show that skill mismatch first decreases with technology vintage (as firms adopt newer technologies) but then increase again for the newest technologies, reflecting the particularly challenging skill requirements of cutting-edge innovations. Two-digit SIC industry fixed effects included. Standard errors are heteroskedasticity consistent and clustered by firm.

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

Table 7: Alternative Skill Mismatch Measure

	(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)
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 uses five alternative skill mismatch measures for year 2020. Column (1) shows our benchmark result. Column (2) uses skill mismatch measure constructed by aggregating over occupation. Column (3) uses 10 skill classifications used by Deming and Kahn (2018). Column (4) uses Modern Bert. Column (5) uses Cosine simiarlity. Column (6) uses Hellinger. The results are qualitatively similar to our main findings, confirming that skill mismatches negatively affect firm performance. Standard errors are clustered by firm. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Appendix for "Skill Mismatch Over The Technology Lifecycle"

A Variable definitions

- *Skill mismatch*: As described in Section 3.2
- *Technology vintage*: Following the literature (eg. 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)
- TFP: TFP, as explained in Imrohoroglu and Tuzel (2014). The extended sample for TFP is provided https://sites.google.com/usc.edu/selale-tuzel/home?authuser=2

B Additional Tables

Table A1: Top 7 Software Used in Each TechVint Bins

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: 25 skills into three groups

Skill groups	25 Skill categories
Technical	java / html / javascript
	structured query language / sap / microsoft sql server
	analysis / data analysis / c++
	networking / troubleshooting / technical support
	telecommunications / electronics / network security
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
Peripheral	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

Table A3: Non-Technical Skills: Skill Mismatch vs Technology Vintage

	(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 investigates how skill mismatch vary across different technology vintage. TechVint is technology vintage, normalized by 100. Columns (1) and (4) show linear regression. Columns (2) and (5) show regression with linear and quadratic terms, revealing a U-shaped relationship. Columns (3) and (6) show piece-wise linear regression. The results consistently show that skill mismatch first decrease with technology vintage (as firms adopt newer technologies) but then increase again for the newest technologies, reflecting the particularly challenging skill requirements of cutting-edge innovations. Two-digit SIC industry fixed effects included. Standard errors are heteroskedasticity consistent and clustered by firm.

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

Table A4: Technical Skills: Skill Mismatch vs Technology Vintage

	(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 investigates how skill mismatch vary across different technology vintage. TechVint is technology vintage, normalized by 100. Columns (1) and (4) show linear regression. Columns (2) and (5) show regression with linear and quadratic terms, revealing a U-shaped relationship. Columns (3) and (6) show piece-wise linear regression. The results consistently show that skill mismatch first decrease with technology vintage (as firms adopt newer technologies) but then increase again for the newest technologies, reflecting the particularly challenging skill requirements of cutting-edge innovations. Two-digit SIC industry fixed effects included. Standard errors are heteroskedasticity consistent and clustered by firm.

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

Table A5: Peripheral Skills: Skill Mismatch vs Technology Vintage

	(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 investigates how skill mismatch vary across different technology vintage. TechVint is technology vintage, normalized by 100. Columns (1) and (4) show linear regression. Columns (2) and (5) show regression with linear and quadratic terms, revealing a U-shaped relationship. Columns (3) and (6) show piece-wise linear regression. The results consistently show that skill mismatch first decrease with technology vintage (as firms adopt newer technologies) but then increase again for the newest technologies, reflecting the particularly challenging skill requirements of cutting-edge innovations. Two-digit SIC industry fixed effects included. Standard errors are heteroskedasticity consistent and clustered by firm.

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