From Mainframes to Machine Learning: Skill Gaps over the Technology Life Cycle

Ashwini K. Agrawal Daniel Kim Prasanna Tambe*

March 26, 2026

Abstract

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

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

1 Introduction

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

Yet despite the prominence of skill gaps in discussions of technological change, we have a limited understanding of how skill gaps arise and evolve as technologies diffuse. Firstly, much of the existing literature treats skill gaps as exogenous, abstracting from firms' endogenous adjustments of skill requirements and workers' incentives to acquire new skills as technologies mature. As a consequence, we know little about whether observed skill gaps reflect fixed barriers to technology adoption or whether they reflect temporary adjustment dynamics of firms and workers. Secondly, prior empirical work on technology diffusion typically relies on univariate measures of skill gaps, obscuring the multiple dimensions of human capital that constrain diffusion. Identifying the various skills demanded by firms and supplied by workers is inherently challenging, as many relevant capabilities are unobservable and difficult to infer from job descriptions. For example, a job posting requiring "Java programming" may reflect demand for narrow technical expertise, broad problem-solving ability, and/or social skills needed for collaborative software development. Disentangling these dimensions is crucial for identifying which skills—and *combinations* of skills—impede technology diffusion.

In this paper, we study how skill gaps endogenously evolve and shape the diffusion of information technologies. We develop a conceptual framework that builds on the classic work

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

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

The key departure is the introduction of skill gaps——defined as the distance between the vector of skill weights required by firms and the vector of skill weights possessed by workers (Lazear, 2009). These skill vectors capture both technical and nontechnical dimensions of human capital that relate to technology adoption. Differences in skill requirements across firms stem from variation in the tasks embodied in the information technologies and complementary organizational structures that firms adopt (Acemoglu and Autor, 2011; Brynjolfsson and Hitt, 2000). We use skill gaps to microfound the productivity dynamics of newly hired workers that are otherwise assumed in the baseline model; we postulate that larger gaps reduce worker productivity, consistent with prior empirical evidence.[2] Our model adaptations generate novel predictions about three key issues: (i) how skill gaps endogenously evolve over the technology life cycle, (ii) what skill attributes serve as technology diffusion bottlenecks, and (iii) how skill gaps can help explain the persistence of legacy technologies with limited growth opportunities.

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

[2]See for example, Bandiera et al. (2025), Coraggio et al. (2025), Lise and Postel-Vinay (2020), and Guvenen et al. (2020).

technologies across firms over time, following Tambe (2014).

We extract firm skill requirements and worker skill attributes from the raw, unstructured text of job postings and résumés by using large language models (LLMs) in combination with standard natural language (NLP) processing techniques. Specifically, we implement a zero-shot classification approach (Yin et al., 2019) based on Bidirectional Encoder Representations from Transformers (BERT; Devlin et al. (2019)) to map each set of observed skill descriptions into a vector of weights across various skill categories, spanning managerial, financial, and technical competencies.[3] This procedure yields a consistent and comparable representation of the skills that firms demand and the skills that workers supply in our sample.

A key advantage of our approach is that LLMs recover latent skill content that extends beyond what is explicitly encoded in raw skill labels, addressing important limitations of both keyword-based methods and direct use of the raw data (Dell, 2025). Job postings and résumés contain thousands of distinct and inconsistently defined skills, making it difficult to construct parsimonious and comparable measures of skill demand and supply without imposing strong assumptions on how skills relate to one another. More fundamentally, individual skill labels only imperfectly reflect the underlying capabilities they signal: for example, a skill requirement such as "Java programming" implicitly conveys a broader set of technical and analytical competencies that are not readily captured by keyword matching or manually constructed taxonomies. Because LLMs are trained on a vast corpus of text, they are able to encode the semantic meaning of skills as high-dimensional embeddings that capture the information understood and conveyed by firms and workers in practice. This allows us to systematically map heterogeneous and evolving skill labels into a unified representation, mitigating measurement error and avoiding the limitations of ad hoc skill definitions.

Our model yields predictions about the extent to which the skill attributes of newly hired workers align with firms' skill requirements. In the data, we measure skill gaps as the

[3]In our main specifications, we rely on 25 skill categories identified and verified by the résumé data provider as skill labels that efficiently and parsimoniously span the observed skill space. Our results are robust to alternative skill categorizations used in the literature, such as those developed by Deming and Kahn (2018) and Autor et al. (2003), which define categories such as routine, cognitive, manual, etc.

Euclidean distance between the skill supply vector of a firm's new hires and the firm's skill demand vector, as inferred from its job postings.[4] To approximate the mapping between hires and vacancies, we implement a range of sampling windows and matching criteria. The resulting distance provides an intuitive and transparent summary statistic of skill gaps: larger values indicate greater divergence between worker skills and firm needs. Because skill weights sum to one, this measure captures the degree to which workers' comparative advantages across skills align with firms' needs. Our vector-based approach allows us to compare skill gaps consistently across firms and over time while incorporating a broad range of skill dimensions.

We document several empirical findings that support the predictions of the model. The first finding is that skill gaps between firms and newly hired workers follow a U-shaped pattern over the technology life cycle: skill gaps are high when new technologies are adopted, decline as firms and workers adjust, and rise again as technologies mature into legacy systems. Importantly, this process is slow, unfolding over several decades. Consistent with the model, skill gaps are initially high because workers lack the specific combinations of skills required to work with new technologies. As adoption expands, workers acquire relevant skills and firms adjust their skill requirements, leading to closer alignment between firm skill demand and worker skill supply. Over time, however, as newer and more productive technologies emerge, workers have weaker incentives to join firms that continue to rely on older technologies, causing skill gaps at these firms to increase. Thus, in contrast to the commonly held assumption that skill gaps are fixed exogenous barriers to technology adoption, our analysis shows that skill gaps endogenously vary over the technology diffusion process in a manner that can be explained by fundamental economic forces.

Our second set of results identifies the specific skill attributes that serve as bottlenecks to information technology diffusion. We show that diffusion is constrained not by technical skill gaps alone, but by shortages in combinations of technical and non-technical skills. We draw

[4]Euclidean distance has been used Bandiera et al. (2025) provide another application of Euclidean distance to measure skill gaps.

this conclusion from several complementary findings. First, we document a U-shaped pattern in skill gaps both across technical occupations (such as engineers and data scientists) and across non-technical occupations (such as managers and business analysts) that are central to the organizational changes that accompany IT adoption (Bresnahan et al., 2002; Bartel et al., 2007). Second, when examining gaps along individual skill dimensions, we find that skill gaps are most pronounced among non-technical IT-complementary skills, particularly managerial and coordination capabilities. Moreover, these gaps appear alongside substantial shortfalls in technical skills such as advanced manufacturing and network security knowledge. By contrast, skills associated with tasks less central to IT diffusion (e.g., hospitality or public speaking) are relatively oversupplied on average. Taken together, these patterns indicate that skill gaps reflects scarcities in combinations of technical and non-technical skills that underpin the effective use of information technologies, rather than narrow deficits in technical expertise alone.

Our third set of findings shows how skill gaps contribute to the persistence of legacy technologies with limited growth opportunities. In the model, firms operating older technologies remain viable because they attract workers whose skill portfolios are poorly aligned with frontier systems. Skill gaps at the technological frontier therefore allow firms deploying legacy systems to continue operating despite inferior performance. Consistent with this mechanism, we document that firms operating legacy systems exhibit lower market-to-book ratios than firms that adopt new technologies (market-to-book ratios are commonly used measures of investors' valuations of firm growth opportunities). Both types of firms are nevertheless prevalent and continue to operate and hire despite the skill gaps they face in the labor market. These results complement existing explanations for the persistence of legacy technologies, which otherwise emphasize adjustment costs, switching frictions, or technological lock-in. Our empirical evidence shows that skill gaps explain the coexistence of frontier technologies and legacy systems despite vast differences in the capital market's valuation of their future growth prospects.

We conduct an extensive set of robustness analyses. Our findings are robust to alternative skill taxonomies proposed in prior work, including those developed by Deming and Kahn (2018), and to alternative choices of job-posting and résumé sampling windows used to construct skill gap measures. The results are also consistent across alternative measures of technology vintage. We also analyze sampling rates of job postings and online resumes for firms and workers across technology vintages, and find that sampling rates are unable to account for our findings. Finally, we note that the results are even stronger when we use alternative LLM's such as ModernBERT to estimate firm and worker skill vectors, suggesting that our baseline results are conservative.

We also discuss two important limitations of our analysis. First, our methodology captures skill gaps in workers' comparative advantages across skill dimensions rather than gaps in absolute skill levels — a distinction that may matter empirically if firms care about the depth of workers' capabilities in specific areas rather than their relative profile across skill categories alone. Second, our sample skews toward higher-skilled workers who are more likely to appear on online job platforms and have LinkedIn profiles, which limits our ability to generalize to lower-skilled workers for whom comparable skills data are relatively scarce. This limitation, however, is partially mitigated by the nature of our research question: technology adoption decisions and the skill gaps that shape them are likely to be concentrated among the higher-skilled workers and technology-intensive firms that our well represented by our sample.

The main contribution of this paper is to show how skill gaps endogenously evolve and shape the technology diffusion process. We advance the literature on technological change and human capital by relaxing two standard assumptions—that skill gaps are exogenous and sufficiently characterized by one-dimensional measures of skills. These innovations yield several new insights into labor market dynamics that are otherwise obscured in existing frameworks. First, skill gaps follow a U-shaped pattern over the technology life cycle and are pronounced not only for firms adopting new technologies, but also for firms deploying legacy systems. Second, diffusion bottlenecks arise from shortages of specific combinations of

nontechnical and technical skills, not from technical skill deficits alone. Third, labor market segmentation induced by skill gaps helps explain the persistence of legacy systems despite their limited growth prospects. Together, these findings recast skill gaps not as a fixed barrier to technology adoption but as a dynamic feature of the diffusion process itself.

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

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

2 Related Literature

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

pelling evidence of skill gaps as a central determinant of technology diffusion and aggregate productivity.

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

Understanding the forces that shape skill gaps is crucial for assessing how these gaps dissipate or persist as technologies diffuse, thereby influencing long-run adoption patterns and the coexistence of new technologies with legacy systems. Absent a clear account of the mechanisms that shape skill gaps, it is difficult to interpret observed skill shortages and distinguish whether skill gaps reflect a persistent barrier to adoption or temporary adjustment frictions. Moreover, skill gaps are not confined to firms adopting new technologies, as firms operating legacy technologies also frequently report difficulty hiring workers with desirable skills (Monaghan and Bass, 2020). The nature of skill gaps facing these firms, however, may reflect fundamentally different economic forces. Finally, understanding the endogenous formation of skill gaps is a prerequisite for evaluating policies aimed at workforce development, as effective policy design requires separating market-driven adjustment dynamics from exogenous skill constraints (Dillon et al., 2025; Katz et al., 2022; Card et al., 2018).

Progress on these issues also requires understanding the multiple dimensions of human capital that differ between firm requirements and worker capabilities. Empirical studies of technological change and human capital have typically examined skill gaps through a univariate lens, focusing on broad categorizations of skill or education (Deming, 2023). This abstraction is sufficient for existing work to explain aggregate patterns of technology diffusion and productivity, however, it is insufficient for identifying the specific skill requirements that reflect technology differences across firms. For example, we have a limited understand-

ing of the extent to which technology diffusion is constrained by technical skill gaps alone or by gaps in IT-complementary nontechnical skills such as managerial and organizational capabilities. Moreover, there is little evidence on the *combinations* of these skills as diffusion bottlenecks (Tambe (2026) is one exception), although there are numerous accounts of hiring managers reporting difficulties in finding employees with the right "mix" of skills (OECD, 2024).

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

We depart from this framework to allow for heterogeneous skill gaps between firms and newly hired workers, to more closely reflect the rich nature of human capital emphasized in contemporary discussions of technology adoption. In our model, firms differ in their skill requirements across multiple dimensions, and newly hired workers vary in the extent to which their skill portfolios align with these requirements. This extension allows skill gaps to vary systematically across firms, technologies, and stages of the technology life cycle. The model thus generates new empirical predictions about the joint evolution of skill gaps and technology diffusion that go beyond the original scope of Chari and Hopenhayn (1991) and speak directly to the central questions of this paper.

3 Theoretical Framework

We present our conceptual framework by first summarizing the relevant aspects of Chari and Hopenhayn (1991). The model is an overlapping-generations economy with an infinite horizon (time t = 0,1,2...), where workers of two types (young and old, each of mass 1) live for two periods. In each time period t, a new technology is introduced that is more productive than all previous vintages; a specific vintage is denoted by τ (which means that the technology was first introduced at time t-τ). Firms adopt a technology of a particular vintage and hire U unskilled workers (type u) and S skilled workers (type s), where skill is a one-dimensional measure of human capital that is specific to a given technology vintage; there are no search or matching frictions in the model. A firm's production function $Y = \gamma^{t-\tau} f(U, S)$ is constant returns to scale, strictly increasing and concave in the number of skilled workers, and the productivity parameter γ is greater than 1.

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

Our first adaptation is to introduce an intuitive and well-accepted definition of skill gaps

to provide a richer characterization of skilled and unskilled workers in the baseline model. Specifically, we define a skill gap between firms and workers as the discrepancy between firms' multidimensional skill requirements and workers' multidimensional skill attributes, each measured using skill weight vectors following Lazear (2009). This measure has precedence in the academic literature (Bandiera et al., 2025) and the practitioner sphere (International Labour Organization, see https://www.ilo.org/skills-mismatches). In the context of our model, when a firm adopts information technology of a particular vintage, the firm designs a production system with complementary assets built around the technology (Brynjolfsson et al., 2002; Brynjolfsson and Hitt, 2000). Production entails the completion of various tasks; the skills required of workers to perform these tasks represent the firm's skill requirements. This connection between technology, tasks, and skill requirements is rooted in the task-based approach to human capital (Acemoglu and Autor, 2011).[5]

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

We denote the skill gap for individual skill i between a firm using technology of vintage

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

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

Our second adaptation is to assume that technological change is skill-biased and that the demand for IT-complementary skills rises faster than the speed by which workers acquire these skills. This assumption is grounded in a large literature that examines changes in the wage structure through the lens of a race between technological change and education (Goldin and Katz, 2008). A number of papers find that changes in information and communication technologies during the recent decades have been associated with relative wage and employment changes that reflect rapid growth in the demand for IT-complementary skills that outpaces their supply (see for example Deming and Noray (2020); Hershbein and Kahn (2018); Acemoglu and Autor (2011)). Formally, we assume that new technology vintages have greater IT-complementary skill requirements than older technologies: $\frac{\partial F_i^{\tau,u}}{\partial \tau} < 0$ for IT-complementary skills $i \in \mathcal{C}$. Because skill weights must add up to 1, this assumption implies that $\frac{\partial F_i^{\tau,u}}{\partial \tau} \geq 0$ for IT-neutral skills $i \in \mathcal{N}$. We model the relatively slower speed of workers' IT-complementary skill acquisition by assuming $|\frac{\partial W_i^{\tau,u}}{\partial \tau}| < |\frac{\partial F_i^{\tau,u}}{\partial \tau}|$ for IT-complementary skills $i \in \mathcal{C}$.

Our third adaptation is to use our definition of skill gaps to microfound the production

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

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

Taken together, the adaptations that we introduce microfound key primitives of Chari and Hopenhayn (1991) without altering its mechanics.[6] As a result, we are able to generate

[6]Following Chari and Hopenhayn (1991), we assume frictionless labor markets in which workers lacking technology-specific human capital are indifferent across firms. This assumption simplifies the matching process between firms and workers: since all new hires lack vintage-specific human capital upon arrival — a condition we verify in the data — there is no scope for differential sorting across skill dimensions, and our framework does not require the richer machinery developed to study multidimensional matching. A growing literature develops such models; see Lindenlaub (2017); Lise and Postel-Vinay (2020); Guvenen et al. (2020); Lindenlaub and Postel-Vinay (2023); Bandiera et al. (2025) for prominent contributions. Our paper takes a complementary approach, focusing on how firm and worker skill attributes evolve over the technology life cycle rather than on the process by which multidimensional firms and workers form matches.

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 skill gaps 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 gaps over the technology life cycle.

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

Intuitively, skill gaps are high for newly introduced technologies because few workers initially possess the capabilities required to utilize them effectively. With time, as more unskilled workers gain experience working with specific technologies, they become skilled and allow firms to adjust tasks so that firm skill requirements more closely match the skills of incoming unskilled workers. However, once maturing technologies become sufficiently old and less productive than new technologies, unskilled workers have dampened incentives to join firms that continue to deploy these older systems. These firms then face higher adjustment costs when reorganizing tasks and skill requirements for new hires, causing skill gaps to rise once again.

Prediction 2: IT-complementary skill requirements exceed worker capabilities

The model predicts that firms' demand for IT-complementary skills exceeds incoming unskilled workers' supply of these capabilities. That is, the model predicts positive skill gaps, $G_c^{\tau,u} > 0$, for unskilled workers across IT-complementary skills $i \in \mathcal{C}$. This prediction stems

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

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

The model predicts that skill gaps along IT-complementary skill dimensions should be U-shaped across technology vintages. The logic for this prediction stems from the arguments used for predictions 1 and 2. As s_τ increases, firms will prioritize adjustments in IT-complementary skill requirements, as these skill requirements exhibit the largest gaps between firm demands and worker capabilities due to skill-biased technical change. 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 skill gaps, leading to a rise in misalignment. The model thus predicts a U-shaped relationship between IT-complementary skill gaps and technology vintage that drives the patterns we observe for the full array of skills in the sample.

Prediction 4: Firm growth opportunities decrease with technology vintage

In the model, newer technologies are inherently more productive than older vintages, implying that firms operating legacy technologies face weaker growth opportunities due to their lower productive capacity. Nevertheless, skill gaps allow these firms to continue at-

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

4 Empirical Approach

4.1 Challenges

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

There are numerous challenges with implementing this procedure in practice. First, there are limited data of sufficient granularity that adequately describe the skill attributes of firms and workers. Prior studies of skill gaps 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 gaps 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 de-

mand, 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 specific coding expertise, 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 gaps.

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 gaps that are computationally feasible and comparable across firms.

4.2 Empirical Methodology

We devise an empirical methodology that helps us overcome the aforementioned challenges and empirically characterize skill gaps over the technology life cycle. First, we construct a novel matched employer-employee panel dataset that combines information from U.S. firm

job postings with information from workers' online professional profiles. 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. From these data, we construct new firm-worker matches each year by identifying workers who join a firm within a given year.

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 skill gaps. 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.[7] We adopt Revelio's skill taxonomy rather than define our own skill taxonomy, under the assump-

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

tion 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 skill gaps.

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 gaps. Specifically, for each firm-year observation, we define the skill gap 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 $\mathbf{w} = (w_1, w_2, \ldots, w_{25})$:

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

for skill attributes j=1, 2,...25. Intuitively, our skill gap 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 smaller skill gap. 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 skills gap.

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.

5 Data

5.1 Data Sources and Sample Construction

We construct a matched employer-employee dataset that combines information from three different sources. First, we collect data on job postings for U.S. firms over the period 2010–2024 from Lightcast. Job postings contain detailed information on the skill requirements firms specify when hiring, making them well suited to measuring firm-level skill demand (Hershbein and Kahn, 2018; Deming and Kahn, 2018). Second, we collect online resume data for individual workers from Revelio Labs. The Revelio data allow us to observe workers' skill attributes and employment histories, including the firms they join and the years in which they do so. Third, we draw standard firm-level financial and accounting variables from Compustat.

We merge these sources at the firm-year level. The Lightcast and Revelio data are the primary inputs for constructing our skill gap measures: Lightcast captures what firms require, and Revelio captures what workers supply. The matched structure of the dataset allows us to approximate, for each firm and year, the gap between the skill profile demanded in job postings and the skill profile of workers hired into the firm. Specifically, we identify all the

job postings that correspond to a given firm-year, and then identify all workers who join the firm within different sampling windows around this period (6 months, one-year, 2-year, etc.). We use these worker-firm pairs to estimate the degree of skill gaps observed over a given period of time.

5.2 Sample Characteristics

Table 1 compares our analysis sample with the broader Compustat population. Our sample consists of 22,467 firm-year observations spanning 2010–2024. Relative to the Compustat population, sample firms are substantially larger, with median sales of $1.4 billion compared to $127 million in the population, and considerably more profitable, with a median profitability ratio of 0.111 versus essentially zero in the Compustat universe. Sample firms are also older, with a median age of 23 years compared to 9 years in the full Compustat population. These differences reflect selection toward larger and more established firms that are more likely to post jobs online and employ workers with online profiles. We discuss the implications of our sample composition for the interpretation of the findings in Section 7.

5.3 Technology Vintage

A central variable in our analysis is the age of the information technologies that firms deploy, which we refer to as technology vintage. We identify the technologies listed in firm job postings and assign each a vintage date based on the year of its commercial introduction. Figure 2 plots the distribution of technology vintage across firm-year observations in our sample. The modal vintage corresponds to technologies approximately 28 to 31 years old, and a substantial share of observations involve technologies older than 36 years — consistent with the widespread persistence of legacy systems documented in the broader literature on technology diffusion.

To illustrate the range of technologies captured by this classification, Table A1 lists the most frequently cited technologies within each vintage bin. Frontier technologies — those in

the youngest bins — include cloud platforms such as Amazon Web Services and Microsoft Azure, and modern programming languages such as Python and C. Technologies in the oldest bins include legacy systems and lower-level languages associated with mainframe-era infrastructure, such as COBOL, Unix, SQL, and C. The progression across bins tracks recognizable waves of enterprise technology adoption, lending face validity to our vintage classification.

5.4 Skill Measurement

We identify 25 skill groups in accordance with Revelio's classification scheme, and then further organize these skills into three categories based on their complementarity with information technology deployment. The first category comprises IT-complementary technical skills, including programming languages, database management, networking, and related capabilities. The second comprises IT-complementary nontechnical skills — managerial, organizational, and coordination capabilities that are necessary for firms to effectively deploy and operate IT systems (Bartel et al., 2007; Brynjolfsson et al., 2002). The third category comprises IT-neutral kills that are unrelated to information technology, such as healthcare, legal, and hospitality competencies. Table A2 in the appendix provides the full list of skill groups and their classification.

For each skill group, we compute an entailment score that captures the degree to which a given job posting or worker résumé reflects that skill category. Firm-level skill demand is then measured as the distribution of entailment scores across job postings filed by the firm in a given year; worker-level skill supply is measured as the corresponding distribution across the résumés of workers who join the firm in that year. The skill gap for a given firm-year observation is the Euclidean distance between the firm's skill demand vector and its workers' skill supply vector across the 25 skill groups.

Table 2 presents summary statistics for skill demand, supply, and gaps across all 25 skill groups. The largest gaps are concentrated in IT-complementary nontechnical categories:

management and leadership (gap mean of 0.040), process improvement and program management (0.028), and project management (0.026). Among IT-complementary technical skills, networking and troubleshooting exhibits the largest gap (0.026), while gaps in programming-intensive categories such as SQL and Java are more modest. IT-neutral skill groups exhibit the smallest gaps on average, consistent with their weaker connection to technology deployment. Across all skill groups, the mean aggregate skill gap is 0.159, with substantial dispersion (standard deviation of 0.166), indicating significant heterogeneity in the degree of skill misalignment across firms and years.

6 Empirical Findings

6.1 Skill gaps over the technology life cycle

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

Several additional features of Figure 3 shed light on the relationship between skill gaps and the technology diffusion process. The long time-series of observed technology vintages in the sample reflects the slow pace of adoption in the data, consistent with studies of gradual diffusion relative to the rate of technological progress (Kalyani et al., 2025). The

declining slope of skill gaps among younger technologies aligns with the model's diminishing-adjustment mechanism: as the stock of skilled workers expands, firms first close their largest skill gaps, generating progressively smaller reductions in skill gaps over time. Finally, skill gaps reach their minimum roughly 35 years after technology inception, suggesting that it can take several decades for information technologies to fully diffuse and for skill gaps to be minimized.

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

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

These findings rule out several plausible alternative priors. First, skill gaps are not a fixed barrier to adoption, nor are they confined to firms adopting new technologies: they fluctuate

with technology vintage and serve as a binding constraint for firms operating legacy systems. This pattern is consistent with practitioner surveys and business-press accounts documenting hiring difficulties among firms reliant on older technologies. An important implication of the findings is that similar observed levels of skill gaps can reflect fundamentally different economic forces. We observe comparable gaps among firms deploying frontier technologies and among firms using legacy technologies. According to the model, skill gaps in the former case reflects a scarcity of workers who possess the skills required to operate new technologies, whereas skill gaps in the latter case partly reflect weak worker incentives to join firms that rely on technologies that are relatively less productive. Distinguishing between these mechanisms is essential for interpreting observed skill shortages and for evaluating policies aimed at workforce development and technology adoption.

6.2 Characterizing Skill Bottlenecks to Technology Diffusion

The model also makes predictions about the key skill attributes that serve as bottlenecks to information technology diffusion. Prediction 2 states that firm requirements should exceed worker capabilities for IT-complementary skills. These skills may include technical skills such as programming and troubleshooting, as well as non-technical skills such as management and organizational capabilities (Brynjolfsson and Hitt, 2000). To test this prediction, we examine skill gaps across each of the individual skill categories in our data. We graphically depict average skill vector weights for firm requirements and worker attributes in Figure 4. We also aggregate the skills in our sample into three groups following prior work on IT and productivity (Bartel et al., 2007): IT-complementary technical skills, IT-complementary nontechnical skills, and IT-neutral skills, and we present these aggregated skills in Figure 5.

Consistent with this model, Figures 4 and 5 show that skill gaps for IT-complementary skills—and especially for IT-complementary nontechnical skills—are substantially larger than skill gaps for IT-neutral skills across all technology vintages. These results indicate that firms' skill requirements exceed workers' capabilities precisely in the dimensions most

relevant for IT adoption. It is also notable that the largest bottlenecks to diffusion arise not only in technical competencies but also in nontechnical, IT-complementary skills that the literature identifies as critical for organizational and process changes accompanying new technologies (Bartel et al., 2007; Bresnahan et al., 2002).

For example, the most pronounced gaps 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, and have been highlighted by the literature as being extremely important for firm productivity (see for example, Bloom et al. (2013) and Weidmann et al. (2024). 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 deficits and are sometimes supplied in excess. Taken together, these findings demonstrate that overall skill gaps do not simply reflect a shortfall of "STEM skills" alone, but rather a targeted scarcity in high-weight, IT-complementary attributes, including managerial and finance-related capabilities. This finding underscores recent calls for further analysis of the determinants of the demand for multidimensional human capital (Deming, 2023).

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

In a second analysis, we examine skill gaps and technology vintage separately for each 2-digit occupational code in our sample. We estimate our main regression specification of skill gaps on technology vintage (using both a quadratic and a piecewise linear specification) for each occupational code, and present the results in Table 7. The findings indicate that we observe significant skill gaps that decrease over time for occupations such as management, computer and math intensive jobs, engineers, sales, and production worker roles (the individual skill vector weights for these workers and their employers are depicted in Figure 7). In contrast, Table 7 shows that occupations that are typically less related to information technologies do not show meaningful variation in skill gaps across technology vintages. The results support the view that skill gaps for occupations that are central to IT-enabled changes within the firm are critical for explaining IT-diffusion rates and exhibit patterns that are consistent with the model mechanism.

6.3 Skill Gaps and the Persistence of Legacy Technologies with Poor Returns

The model also predicts that skill gaps can help explain the persistence of legacy systems with relatively poor growth opportunities. Firms operating older technologies remain viable because they attract workers whose skill portfolios are poorly aligned with frontier systems. Skill gaps at the technological frontier therefore allow firms deploying legacy systems to continue operating despite inferior performance.

To test this prediction, we estimate sample firms' market-to-book ratios—a commonly used measure of investors' valuations of firm growth prospects—and study how these ratios vary with technology vintage. We graph the relationship between market-to-book ratios and technology vintage in Figure 8. Consistent with the model mechanism, we document that firms operating legacy systems exhibit lower market-to-book ratios than firms that adopt new technologies. Both types of firms are nevertheless prevalent, and continue to operate and hire workers despite the skill gaps that they both face in the labor market. The graphical results

are further reinforced by the regression analysis presented in Table 8. Across specifications that control for year and industry fixed effects, the estimated coefficient on technology vintage is negative and statistically significant, ranging from -1.2 to -2.8.

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

7 Robustness Analysis

We present a number of analyses that illustrate the robustness of our empirical findings. We also discuss some of the limitations of our analysis.

7.1 Choice of language model

Our main results are based on skill vectors that we construct using BERT, a large language model that is well known for its accuracy and computational feasibility. BERT is a workhorse model for text classification because it is bidirectional (in contrast to large language models

that only read text in one direction), and therefore better suited for capturing contextual information than many other language models. BERT is appropriate for our skill classification task because it can capture the semantic meaning of explicit raw skill data and extract the skills that are implied but not explicitly mentioned in observable skill labels.

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. Nevertheless, to explore the robustness of our results to the choice of large language model, we repeat our analysis for firm-worker matches in the year 2020 using ModernBert. We find results that are even stronger than what we find using Bert. For example, as shown in Table 9, column 4, the U-shape of skill gaps 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.

7.2 Technology vintage measure

In our main specifications, we use the average vintage for all information technologies listed in a firm's job postings as an estimate of the IT capital vintage of the firm. The average IT vintage that we estimate might be skewed by software that is particularly old and so commonplace that it may not be particularly relevant for characterizing changes in production systems that are empirically most relevant in the context of skill gaps. 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 gaps that we present are not significantly related to the particular

measures of IT capital vintage that we employ.

7.3 Absolute advantage in skill

Our measure of skill gaps between workers and firms is centered around differences in comparative advantage across skills — not absolute advantage. For example, when two workers list "Python programming" on their professional profiles, we cannot distinguish between the two workers based on their absolute 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.

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. Our measure of skill gaps is thus informative about relative differences in skill attributes between firms and workers, assuming a common baseline level of skill proficiency that is similar across firms and workers. This assumption is potentially strengthened by the observation that the level of absolute advantage in a particular skill is likely to be correlated with other skills that appear in a job posting or worker resume, leading to reduced measurement error overall when considering skill weights across all labels in our analysis.

7.4 Measurement error: misidentifying firm-worker matches and sample selection

We show that our results are robust to using alternative time windows and other criteria to define worker-firm matches that we use to measure skill gaps between firms and workers. For example, we show that our results are similar when we vary the window length in which we match newly hired workers with firm job postings. We also show that our results are unaffected by whether we match workers and firms by occupation, and whether we weight observations by sampling rates of occupations across industries.

7.5 Alternative skill labels

Our results are robust to alternative skill classifications, including the categories proposed by Deming and Kahn (2018) and Autor et al. (2003). These frameworks are well suited to studying the effects of technological change and other market-wide forces on aggregate shifts in skill demand, and we find similar results across all of our tests when we adopt them in place of our baseline skill labels provided by Revelio Labs. Our findings therefore do not hinge on any particular skill taxonomy; they reflect skill gaps consistent with the predictions of our model.

In our setting, we prefer the Revelio skill labels for two reasons. First, they capture granular differences in skill content within broad categories — such as cognitive or technical skills — that would otherwise appear similar across firms adopting new technologies and those deploying older legacy systems. Second, they are identified and verified by Revelio Labs to efficiently span the full skill space represented in the underlying resume data. We therefore present our main results using the Revelio skill labels as our baseline throughout.

8 Conclusion

This paper studies how skill gaps evolve during the technology diffusion process. The existing literature largely treats skill gaps as an exogenous, one-dimensional barrier to adoption. We show instead that skill gaps evolve endogenously as technologies move through their life cycle, reflecting simultaneous shortages across multiple skill dimensions. Using matched data on firm job postings and worker résumés, we construct firm-level measures of skill demand and supply that reveal how gaps in both technical and complementary nontechnical capabilities develop and shape technology adoption.

Our analysis yields three main findings. Skill gaps follow a U-shaped pattern over the technology life cycle: high at initial adoption, declining as firms and workers adjust, and rising again as technologies mature into legacy systems. Diffusion bottlenecks arise not only from

technical skill deficits but also from shortages of complementary nontechnical capabilities — particularly managerial and coordination skills — that enable firms to effectively deploy new technologies. We also present evidence that skill gaps contribute to the persistence of legacy technologies, helping explain why many firms continue operating older systems despite their limited growth prospects. Together, these results recast skill gaps as an endogenous, dynamic feature of the diffusion process rather than a static constraint on adoption.

Our findings point to two directions for future research. The first concerns the joint determination of technology adoption and human capital investment. Workers decide whether to acquire skills that complement emerging technologies while firms decide whether to adopt new systems given the available skill base. Understanding this interaction more deeply is critical for further explaining the speed and pattern of technology diffusion, and integrating these forces into a unified framework remains an important theoretical challenge. The second concerns measurement. The approach we introduce — applying large language models to job postings and résumés — opens broader opportunities to study how skills, technologies, and organizational capabilities co-evolve over time. Leveraging these methods at scale can enhance our understanding of the microeconomic foundations of technological diffusion and its implications for productivity and growth.

References

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

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

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

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

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

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

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

Bandiera, O., Kotia, A., Lindenlaub, I., Moser, C., and Prat, A. (2025). Meritocracy across countries. Working Paper 12208, CESifo. Available at SSRN: https://ssrn.com/abstract=5637579.

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.

Bloom, N., Eifert, B., Mahajan, A., McKenzie, D., and Roberts, J. (2013). Does management matter? evidence from india. *Quarterly Journal of Economics*, 128(1):1–51.

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

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

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

Brynjolfsson, E., Hitt, L. M., and Yang, S. (2002). Intangible assets: Computers and orga-

nizational capital. *Brookings Papers on Economic Activity*, 2002(1):137–181.

Card, D., Kluve, J., and Weber, A. (2018). What works? a meta analysis of recent active labor market program evaluations. *Journal of the European Economic Association*, 16(3):894–931.

Caroli, E. and vanReenen, J. (2001). Skill-biased organizational change? evidence from a panel of british and french establishments. *Quarterly Journal of Economics*, 116(4):1449–1492.

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

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

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

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

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

Deming, D. J. and Noray, K. (2020). Earnings dynamics, changing job skills, and stem careers. *Quarterly Journal of Economics*, 135(4):1965–2005.

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.

Dillon, E. W., Kahn, L. B., Venator, J., and Dalton, M. (2025). Do workforce development programs bridge the skills gap? Technical report, Microsoft Research and University of Rochester and NBER and Boston College and Bureau of Labor Statistics. Working paper.

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

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

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

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

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

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.

Katz, L. F., Roth, J., Hendra, R., and Schaberg, K. (2022). Why do sectoral employment programs work? lessons from workadvance. *Journal of Labor Economics*, 40(S1):S249–S291.

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

Lindenlaub, I. (2017). Sorting multidimensional types: Theory and application. *The Review of Economic Studies*, 84(2):718–789.

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

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

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

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

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

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

OECD, editor (2024). *Understanding Skill Gaps in Firms: Results of the PIAAC Employer Module*. OECD Skills Studies. OECD Publishing, Paris. Accessed: March 26, 2026.

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

Tambe, P. B. (2026). Reskilling the workforce for AI: Domain expertise and algorithmic literacy. *Management Science*, 72(1):515–537.

Weidmann, B., Vecci, J., Said, F., Deming, D. J., and Bhalotra, S. R. (2024). How do you find a good manager? Working Paper 32699, National Bureau of Economic Research.

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.

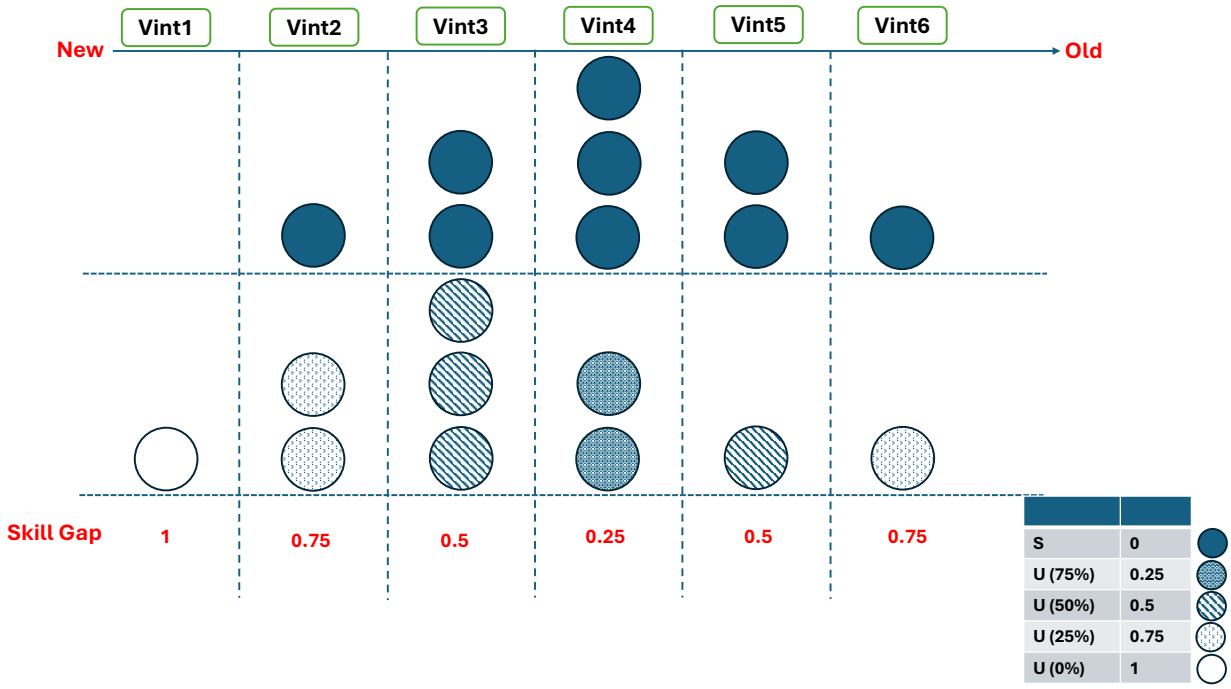


Figure 1: Model Overview

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 gap 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 gap with the firm, while the two unskilled workers each reflect gap 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 gap and employment growth of the firm continues each period until the technology reaches middle vintage 4. At this time (and moving forward), fewer unskilled workers will join the firm, since these workers face less incentive to work with sufficiently older technologies that are less productive than the newest technologies. The firm thus gets fewer skilled workers over time, thereby increasing the gap realized between the firm and its new employees This process continues until the technology becomes obsolete.

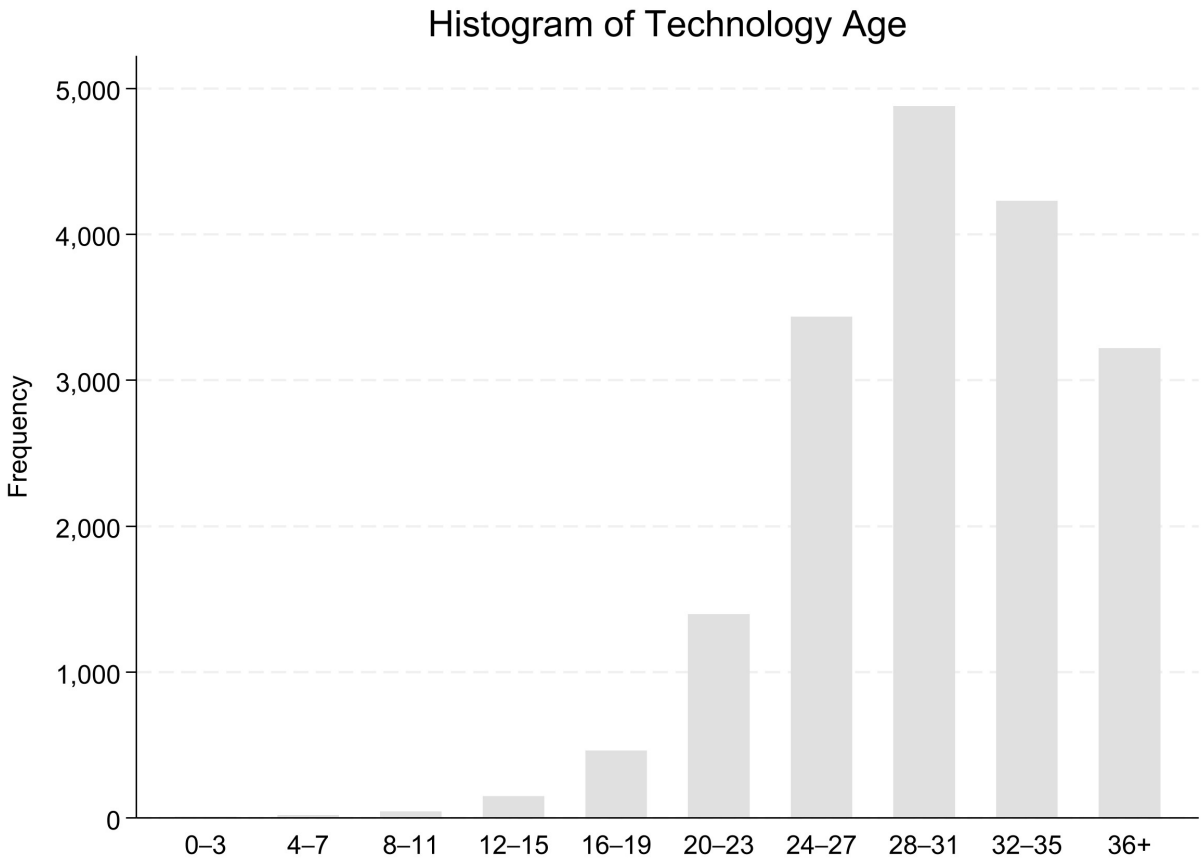


Figure 2: Technology Age Distribution

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

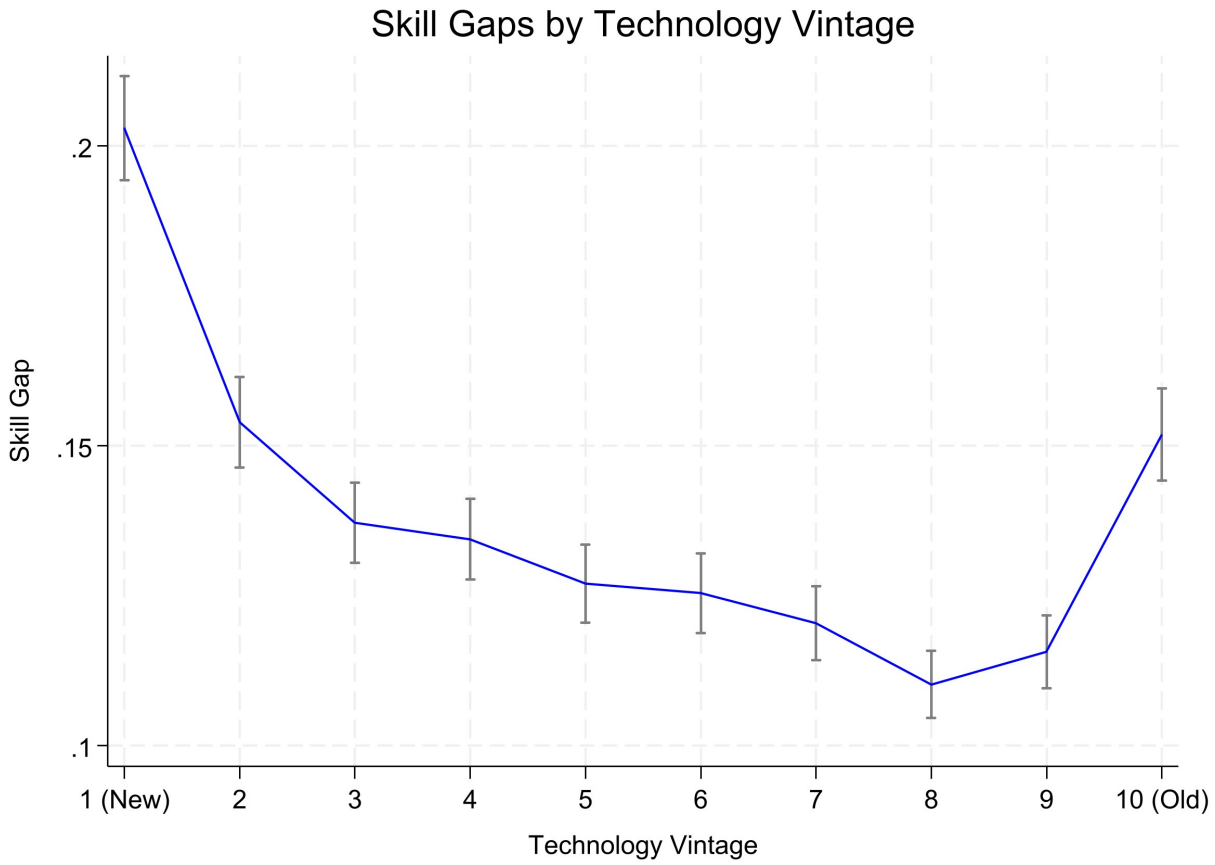


Figure 3: Skill Gaps Across Technology Vintages

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

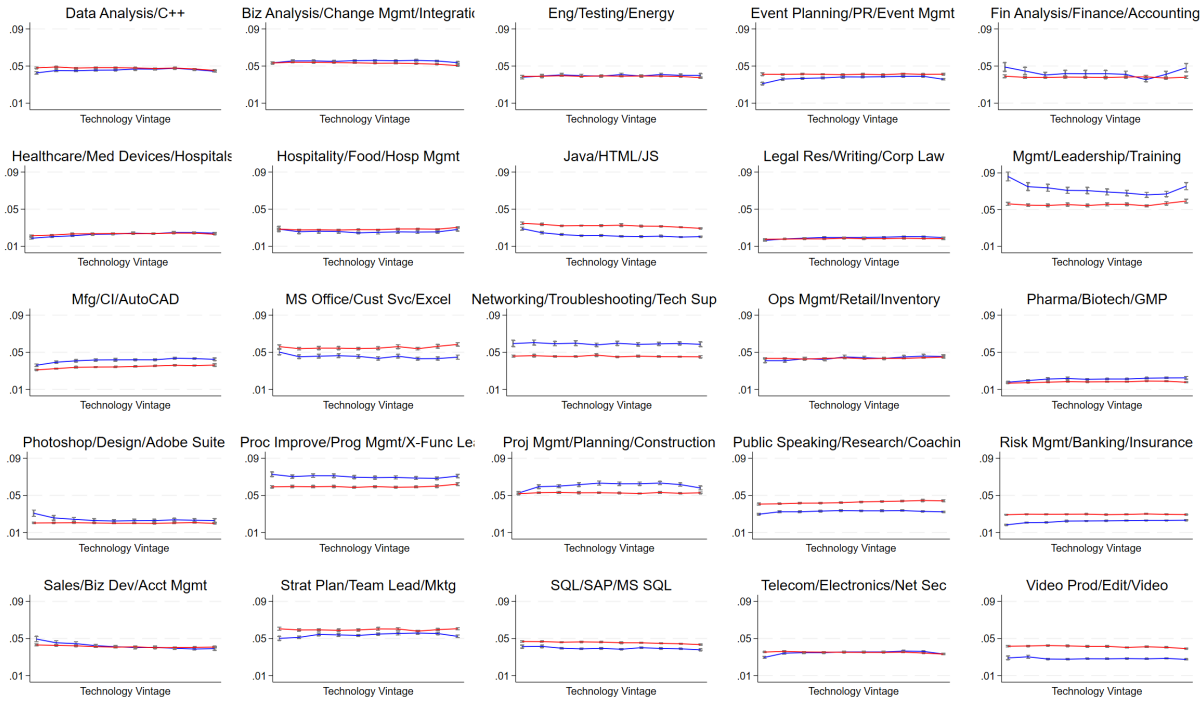


Figure 4: Skill Vector Weights by Skill Category

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

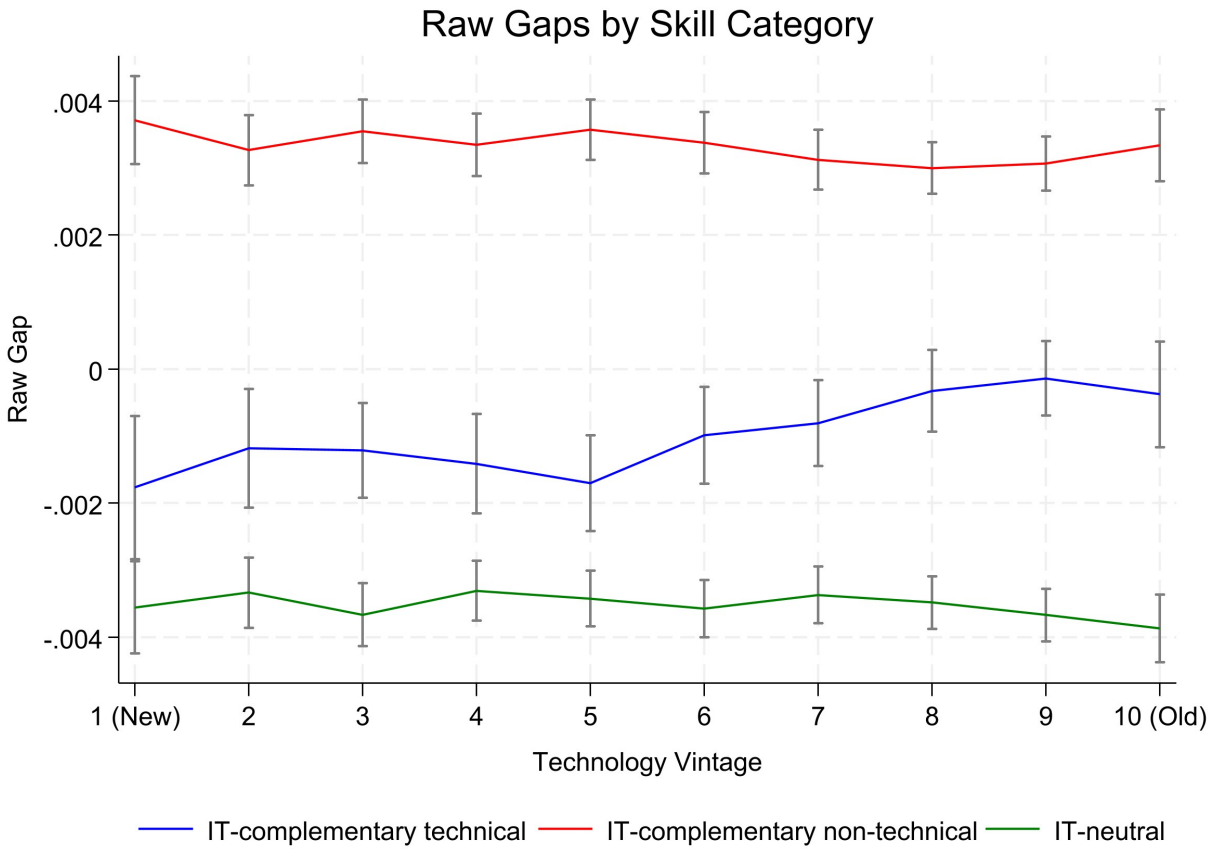


Figure 5: Raw Gap Scores by Skill Category

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

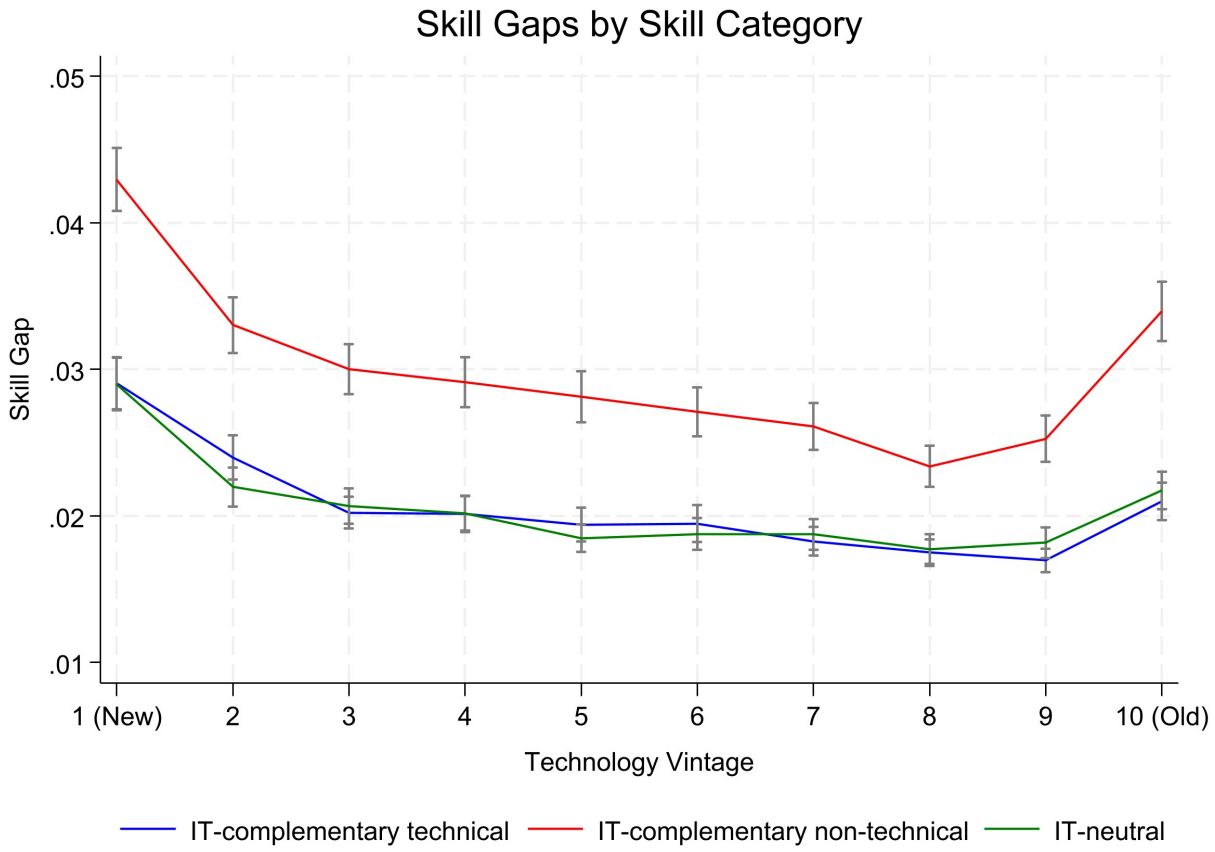


Figure 6: Skill Gaps Across Skill Categories

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

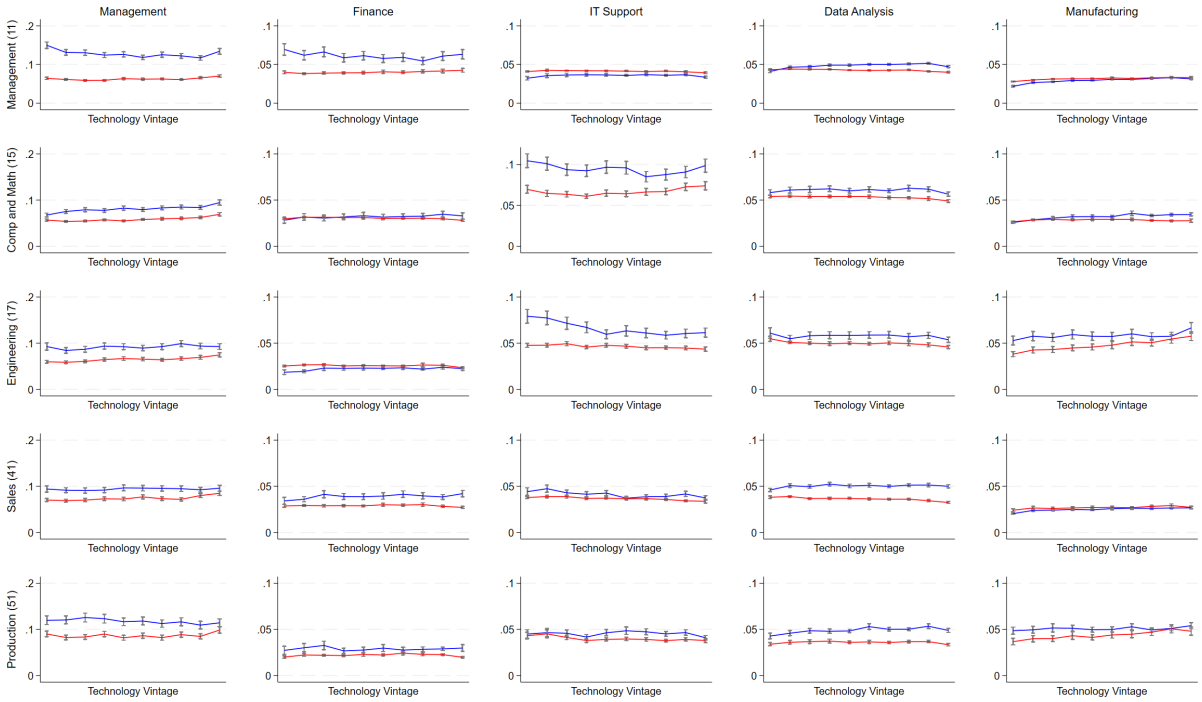


Figure 7: Skill Vector Weights for {Occupation x Skill} Pairs

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

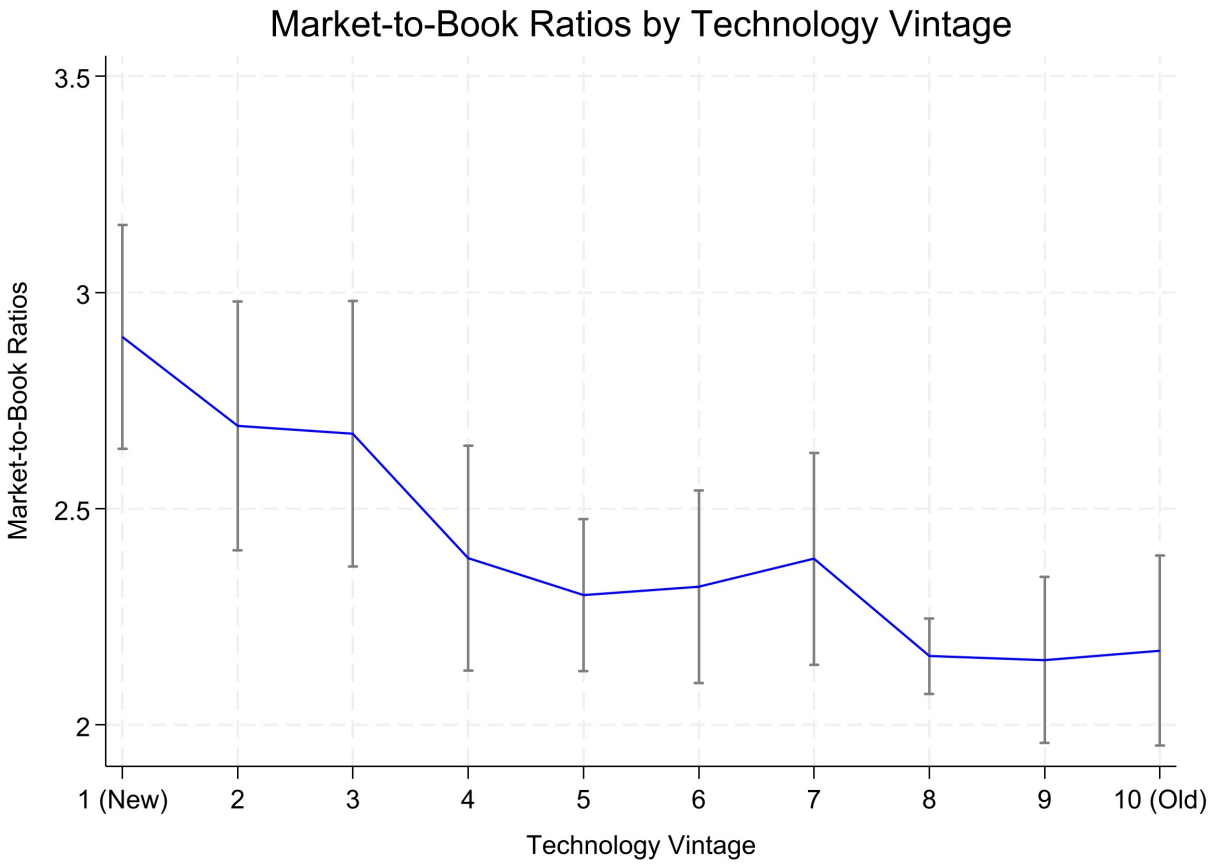


Figure 8: Market-to-Book Ratios across Technology Vintages
This figure depicts average market-to-book ratios for firms that deploy information technologies of different vintages. The X-axis depicts technology vintage bins as defined in Figure 3; the Y-axis depicts the average market-to-book ratio for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

Table 1: Summary Statistics: Sample vs Population

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

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

Table 2: Summary Statistics: Skill Gaps

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

Continued on next page

Summary Statistics: Skill Gaps (continued)

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

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

Table 3: Skill Gaps Across Technology Vintages

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

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

Table 4: Skill Gap and Technology Vintage: IT-Complementary Non-Technical Skills

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

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

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

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

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

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

Table 6: Skill Gap and Technology Vintage: IT-Neutral Skills

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

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

Table 7: Skill Gap and Technology Vintage — Occupation-Level Analysis

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

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

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

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

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

Table 9: Robustness Analysis for Skill Gap and Technology Vintage

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

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

Appendix

A Variable definitions

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

B Additional Tables

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

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

Table A2: Skill Labels and Groupings

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

C Skill Gap Dynamics

This appendix provides a simple microfoundation for the assumption that the skill gap between the firm and its new employees decreases as the number of its skilled employees increases.

C.1 Skill Vectors and Adding-Up Constraints

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

Firm f at time t has a vector of skill requirements

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

and workers hired at time t have skill endowments

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

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

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

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

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

C.2 Skill Gap Measure

Skill gap is defined as the Euclidean distance:

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

For analytical convenience, we work with squared skill gaps:

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

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

C.3 Firm Adjustment Problem

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

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

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

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

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

C.4 Lagrangian

The Lagrangian is:

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

C.5 First-Order Conditions

For each i:

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

Rearranging:

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

C.6 Solving for the Multiplier

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

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

which implies

$$\lambda = 0.$$

C.7 Adjustment Rule

Substituting $\lambda = 0$:

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

Define

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

Then

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

C.8 Skill Gap After Adjustment

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

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

Squaring and summing across skills:

$$\mathcal{G}_{t+1} = \sum_{i=1}^{N}(F_{i,t+1} - W_{i,t+1})^2 \tag{9}$$

$$= \sum_{i=1}^{N}(1 - \alpha(S_t))^2(F_{i,t} - W_{i,t+1})^2 \tag{10}$$

$$= (1 - \alpha(S_t))^2 \sum_{i=1}^{N}(F_{i,t} - W_{i,t+1})^2. \tag{11}$$

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

C.9 Comparative Statics with Respect to Skilled Workers

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

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

First derivative.

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

since the summation term is strictly positive.

Thus skill gaps between the firm and its new employees are decreasing in the number of skilled workers S_t employed by the firm.

Second derivative.

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

Thus the skill gap decreases in S_t at a diminishing rate.